

IBM DB2 9.7  
for Linux, UNIX, and Windows



Version 9 Release 7



**Net Search Extender Administration and User's Guide**  
**Updated November, 2009**





**Note**

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 317.

**Edition Notice**

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at [www.ibm.com/shop/publications/order](http://www.ibm.com/shop/publications/order)
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at [www.ibm.com/planetwide](http://www.ibm.com/planetwide)

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1995, 2009.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

---

## Part 1. Net Search Extender overview and concepts . . . . . 1

### Chapter 1. Net Search Extender key concepts . . . . . 3

SQL scalar search function overview . . . . .	5
Stored procedure search overview . . . . .	6
SQL table-valued function overview . . . . .	7
Additional concepts . . . . .	8
Column transformation function. . . . .	8
DB2 Net Search Extender instance services . . . . .	8
DB2 Net Search Extender instance services on Windows . . . . .	8
Externally stored data . . . . .	9
Administration tables and views. . . . .	9
Partitioned database support . . . . .	9
Indexes on nicknames in a federated database. . . . .	10
Native XML support . . . . .	10
Partitioned table support . . . . .	11
Incremental update based on integrity processing . . . . .	11

### Chapter 2. Key features of DB2 Net Search Extender . . . . . 13

### Chapter 3. Introducing the db2text commands . . . . . 15

---

## Part 2. Installation . . . . . 17

### Chapter 4. Installing Net Search Extender in the DB2 client/server environment . . . . . 19

### Chapter 5. Net Search Extender installation system requirements . . . . . 21

### Chapter 6. Installing NSE on a partitioned DB2 server . . . . . 23

### Chapter 7. Installation on UNIX . . . . . 25

### Chapter 8. Installation on Windows . . . . . 27

### Chapter 9. Directory names and file names . . . . . 29

### Chapter 10. Installing the Outside In libraries . . . . . 31

### Chapter 11. Installation verification . . . . . 33

Installation verification on UNIX . . . . .	33
Installation verification on Windows . . . . .	33

### Chapter 12. Uninstalling Net Search Extender . . . . . 35

Uninstalling Net Search Extender on UNIX. . . . .	35
Uninstalling Net Search Extender on Windows . . . . .	35

---

## Part 3. Migrating to DB2 Net Search Extender Version 9.7. . . . . 37

---

## Part 4. Planning considerations . . . . . 41

### Chapter 13. Directory locations and index storage for Net Search Extender . . . . . 43

### Chapter 14. Resource considerations for a partitioned DB2 server . . . . . 45

### Chapter 15. Stored procedure search memory requirements . . . . . 47

Memory requirements for AIX (64-bit) . . . . .	47
Memory requirements for Windows (32-bit and 64-bit) . . . . .	47
Memory requirements for Solaris (64-bit) . . . . .	47
Memory requirements for Linux (32-bit and 64-bit) . . . . .	48

### Chapter 16. Table, column and index name considerations . . . . . 49

### Chapter 17. Document formats and supported code pages . . . . . 51

### Chapter 18. Outside In filtering software . . . . . 53

### Chapter 19. User roles . . . . . 55

### Chapter 20. Extended text-maintained staging infrastructure for incremental update . . . . . 57

---

## Part 5. Administering Net Search Extender . . . . . 59

### Chapter 21. Net Search Extender instance services . . . . . 61

<b>Chapter 22. Starting and stopping Net Search Extender instance services using the command line . . . . .</b>	<b>63</b>
---	-----------

<b>Chapter 23. NSE locking services . . . . .</b>	<b>65</b>
Using the locking services . . . . .	65
Viewing a lock snapshot . . . . .	67

<b>Chapter 24. Update services . . . . .</b>	<b>69</b>
--	-----------

<b>Chapter 25. Using the DB2 control center . . . . .</b>	<b>71</b>
---	-----------

Starting and stopping Net Search Extender Instance Services using the DB2 Control Center . . . . .	72
Enabling and disabling a database using the DB2 control center. . . . .	72
Text index administration using the DB2 control center . . . . .	73
Creating a text index using the DB2 control center	74
Name panel . . . . .	74
Target panel . . . . .	75
Text properties panel . . . . .	77
Update characteristics panel . . . . .	79
Cache table panel . . . . .	80
Summary panel . . . . .	83
Maintaining a text index using the DB2 control center . . . . .	84
Altering a text index using the DB2 control center . . . . .	84
Dropping a text index using the DB2 control center . . . . .	86
Updating a text index using the DB2 control center . . . . .	87
Showing index events using the DB2 control center . . . . .	87
Activating a text index cache using the DB2 control center. . . . .	87
Deactivating a text index cache using the DB2 control center. . . . .	88
Showing index status using the DB2 control center . . . . .	88

---

**Part 6. Developing: creating and maintaining a text index . . . . . 89**

<b>Chapter 26. Enabling a database . . . . .</b>	<b>91</b>
--	-----------

<b>Chapter 27. Disabling a database . . . . .</b>	<b>93</b>
---	-----------

<b>Chapter 28. Creating a text index . . . . .</b>	<b>95</b>
--	-----------

Creating a text index on binary data types . . . . .	97
Creating a text index on an unsupported data type	97
Creating a text index on a nickname with incremental index update using DB2 Replication . . . . .	98
Creating a text index which the stored procedure search can use . . . . .	99

<b>Chapter 29. Creating text indexes on views . . . . .</b>	<b>103</b>
---	------------

<b>Chapter 30. Creating a text index on range partitioned tables . . . . .</b>	<b>105</b>
--	------------

<b>Chapter 31. Performance considerations for indexing . . . . .</b>	<b>107</b>
--	------------

<b>Chapter 32. Maintaining text indexes</b>	<b>109</b>
---	------------

Updating and reorganizing a text index . . . . .	109
Updating a text index . . . . .	110
Altering a text index . . . . .	111
Clearing index events. . . . .	111
Dropping a text index . . . . .	112
Viewing text index status . . . . .	112
Backing up and restoring indexes. . . . .	113
Removing files from the /tmp directory . . . . .	113

---

**Part 7. Methods for searching text 115**

<b>Chapter 33. Searching for text using SQL scalar search functions . . . . .</b>	<b>117</b>
---	------------

Issuing a query. . . . .	117
Searching and returning the number of matches found . . . . .	117
Searching and returning the score of a found text document . . . . .	118

<b>Chapter 34. Specifying SQL search arguments. . . . .</b>	<b>119</b>
---	------------

Searching for terms in any sequence. . . . .	119
Searching with the Boolean operators AND and OR	119
Searching with the Boolean operator NOT. . . . .	120
Fuzzy search . . . . .	120
Searching for parts of a term (character masking)	120
Searching for terms that contain a masking character . . . . .	121
Searching for terms in a fixed order . . . . .	121
Searching for terms in the same sentence or paragraph . . . . .	121
Searching for terms in sections of structured documents . . . . .	122
Thesaurus search . . . . .	122
Numeric attribute search . . . . .	122
Free-text search. . . . .	123

<b>Chapter 35. Additional search syntax examples . . . . .</b>	<b>125</b>
--	------------

<b>Chapter 36. Searching for text using a stored procedure search . . . . .</b>	<b>127</b>
---	------------

<b>Chapter 37. Searching for text using an SQL Table-Valued Function . . . . .</b>	<b>129</b>
--	------------

Using the highlight function . . . . .	129
--	-----

<b>Chapter 38. Searching on more than one column . . . . .</b>	<b>133</b>
--	------------

<b>Chapter 39. Using text search in outer joins . . . . .</b>	<b>135</b>
---	------------

<b>Chapter 40. Performance considerations during search . . . . .</b>	<b>137</b>
---	------------

<b>Chapter 41. User scenarios . . . . .</b>	<b>139</b>
Simple example with the SQL scalar search function . . . . .	139
Simple example with cache usage and stored procedure search . . . . .	140
Simple example with the SQL table-valued function . . . . .	141

<b>Chapter 42. Using a thesaurus to expand search terms . . . . .</b>	<b>143</b>
The structure of a thesaurus . . . . .	143
Predefined thesaurus relations. . . . .	144
Defining your own relations . . . . .	145
Creating and compiling a thesaurus . . . . .	145
Creating a thesaurus definition file . . . . .	145
Compiling a definition file into a thesaurus dictionary . . . . .	146
Thesaurus support . . . . .	146
Thesaurus supported CCSIDs . . . . .	148
Messages returned by the thesaurus tool . . . . .	149

<b>Chapter 43. Net Search Extender indexing configuration . . . . .</b>	<b>153</b>
Tokenization. . . . .	153
Stop words . . . . .	154
Languages supporting stop words . . . . .	154
Configuration . . . . .	155

---

## **Part 8. Working with structured documents . . . . . 157**

<b>Chapter 44. Searching natively stored XML documents . . . . .</b>	<b>159</b>
Using the default document model . . . . .	159
Using a customized document model . . . . .	160
XQuery support . . . . .	161

<b>Chapter 45. Structured document support . . . . .</b>	<b>163</b>
How a document model describes structured documents . . . . .	163
An example of a document model . . . . .	163
Document models . . . . .	164
Text fields . . . . .	165
Document attributes . . . . .	165
Default document models . . . . .	166
Definition of a document model for structured plain-text documents . . . . .	167

What happens when a GPP document is indexed . . . . .	167
Definition of a document model for HTML documents . . . . .	168
Definition of a document model for XML documents . . . . .	170
Definition of a document model for Outside In filtered documents. . . . .	174
What happens when an Outside In document is indexed . . . . .	175

<b>Chapter 46. Document model reference . . . . .</b>	<b>177</b>
DTD for document models . . . . .	177
Semantics of locator (XPath) expressions . . . . .	178
Limitations for text fields and document attributes . . . . .	180
Outside In tag attribute values . . . . .	180

---

## **Part 9. Reference . . . . . 183**

<b>Chapter 47. Administration commands for the instance owner . . . . .</b>	<b>185</b>
CONTROL command. . . . .	185
START command . . . . .	187
STOP command . . . . .	188

<b>Chapter 48. Administration commands for the database administrator . . . . .</b>	<b>191</b>
ENABLE DATABASE command . . . . .	191
DISABLE DATABASE command . . . . .	193
DB2EXTHL command . . . . .	194

<b>Chapter 49. Administration commands for the text table owner . . . . .</b>	<b>197</b>
ACTIVATE CACHE command. . . . .	197
ALTER INDEX command . . . . .	199
CLEAR EVENTS command. . . . .	203
CREATE INDEX command. . . . .	204
DEACTIVATE CACHE command . . . . .	217
DROP INDEX command . . . . .	219
DB2EXTTH command . . . . .	220
RESET PENDING command . . . . .	221
UPDATE INDEX command. . . . .	222
HELP command . . . . .	225
COPYRIGHT command . . . . .	226

<b>Chapter 50. Net Search Extender install and uninstall command reference for UNIX . . . . .</b>	<b>227</b>
db2nse_install command . . . . .	227
db2nse_deinstall command. . . . .	228
db2nsels command . . . . .	228

<b>Chapter 51. Syntax of search arguments. . . . .</b>	<b>231</b>
Search argument syntax . . . . .	231
Search parameters . . . . .	234

<b>Chapter 52. SQL scalar search function and the SQL table-valued function.</b>	<b>239</b>
CONTAINS scalar function.	239
NUMBEROFMATCHES scalar function.	239
SCORE scalar function.	240
DB2EXT.TEXTSEARCH command.	240
DB2EXT.HIGHLIGHT.	243

<b>Chapter 53. Stored procedure search function.</b>	<b>247</b>
DB2EXT.TEXTSEARCH for stored procedure search	247

<b>Chapter 54. Net Search Extender Messages</b>	<b>251</b>
CTE0000 - CTE0499	251
CTE0500 - CTE0999	270

<b>Chapter 55. Windows system errors</b>	<b>273</b>
--	------------

<b>Chapter 56. Net Search Extender information catalogs</b>	<b>277</b>
Views for database-level information	277
Views for index-level information	280
db2ext.textindexes view	280
db2ext.indexconfiguration view	282
db2ext.textindexformats view	282
Table views for a text index	283
Event view	283
Log tables, views, and nicknames	284

<b>Chapter 57. Text Search Engine reason codes</b>	<b>287</b>
--	------------

---

**Part 10. Troubleshooting . . . . . 293**

<b>Chapter 58. Tracing faults . . . . . 295</b>
---

<b>Chapter 59. Dropping DB2 objects without using the correct Net Search Extender commands</b>	<b>297</b>
Dropping a table	297
Dropping a database	297

<b>Chapter 60. Installation return codes on Windows . . . . . 299</b>
---

<b>Chapter 61. Hints and tips . . . . . 301</b>
---

---

**Part 11. Appendixes . . . . . 305**

<b>Appendix A. Overview of the DB2 technical information</b>	<b>307</b>
DB2 technical library in hardcopy or PDF format	307
Ordering printed DB2 books	310
Displaying SQL state help from the command line processor.	311
Accessing different versions of the DB2 Information Center.	311
Displaying topics in your preferred language in the DB2 Information Center.	311
Updating the DB2 Information Center installed on your computer or intranet server.	312
Manually updating the DB2 Information Center installed on your computer or intranet server	313
DB2 tutorials	315
DB2 troubleshooting information.	315
Terms and Conditions	316

<b>Appendix B. Notices . . . . . 317</b>
--

<b>Index . . . . . 321</b>
----------------------------



---

## **Part 1. Net Search Extender overview and concepts**



# Chapter 1. Net Search Extender key concepts

Net Search Extender offers users and application programmers a fast, versatile, and quick method of searching full-text documents stored in DB2® databases, other databases, and file systems using SQL queries.

To fully understand the capabilities of Net Search Extender, it is necessary to understand key terms and the various options available. It is also necessary to have a basic understanding of DB2 database concepts and terms.

Basically, Net Search Extender searches *text documents* that are held in the column of a database table.

The text documents must be uniquely identifiable. Net Search Extender uses the *primary key* of the table for this purpose.

The documents can be stored in various formats. The formats include either unstructured plain text, structured text such as HTML or XML, or proprietary document formats such as PDF or Microsoft® Office document formats. For the latter, additional filtering software that might have to be licensed separately is required.

Rather than sequentially searching through the text documents at query time that would take a considerable amount of time, Net Search Extender creates a *text index* in order to make documents searchable efficiently.

A text index consists of significant *terms* that are extracted from the text documents.

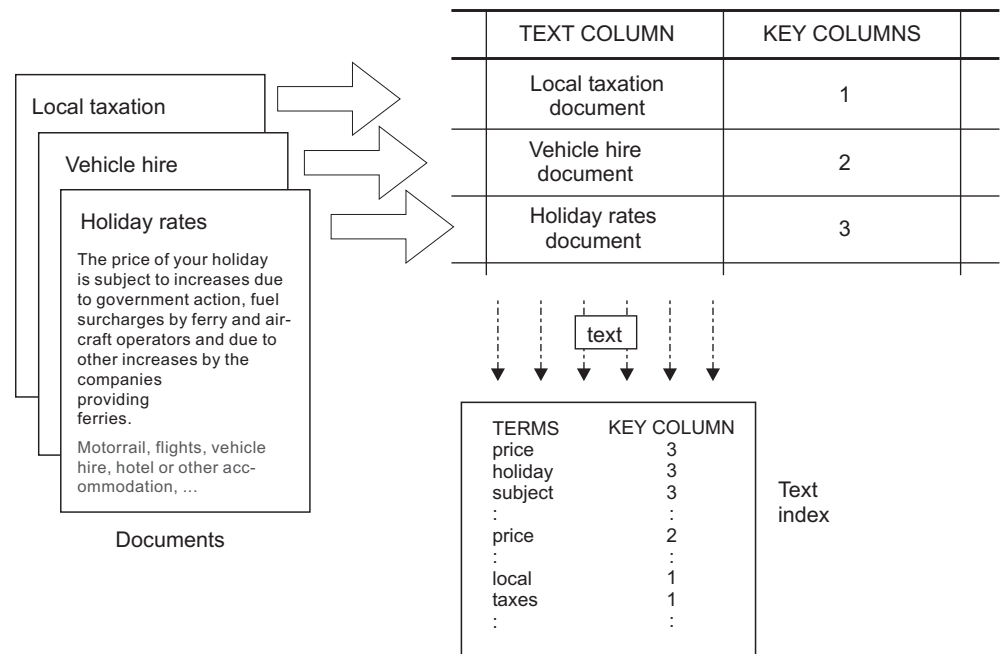


Figure 1. Creating a text index

*Text index creation* is the process of defining and declaring the properties of the index, such as the location of the index. After creation, the text index does not as yet contain any data. *Index update* is the process of adding data about terms and documents to the text index. The first index update adds information about all text documents from the text column to the index. The first update is known as the *initial update*.

By using a text index for searching, there are synchronization issues between the table and the text index that must be taken into account, as any follow-up changes to the table, such as additions, deletions, and updates to the text documents must be reflected in the text index. These changes are applied to the text index with an incremental update.

Net Search Extender supports two options to synchronize the text index with its source table. The basic synchronization in Net Search Extender is based on triggers that automatically store information about new, changed, and deleted documents in a log table. There is one log table for each text index.




Figure 2. Incremental update process with triggers

The basic option is based only on triggers; updates that are not recognized by triggers will therefore be ignored, for example, loading data with the LOAD command, attaching or detaching ranges for a range-partitioned table. An extended synchronization option enables capturing such changes through integrity processing by adding a text-maintained staging table to store information about new and deleted documents, while the log table stores information about changed documents captured through a trigger.

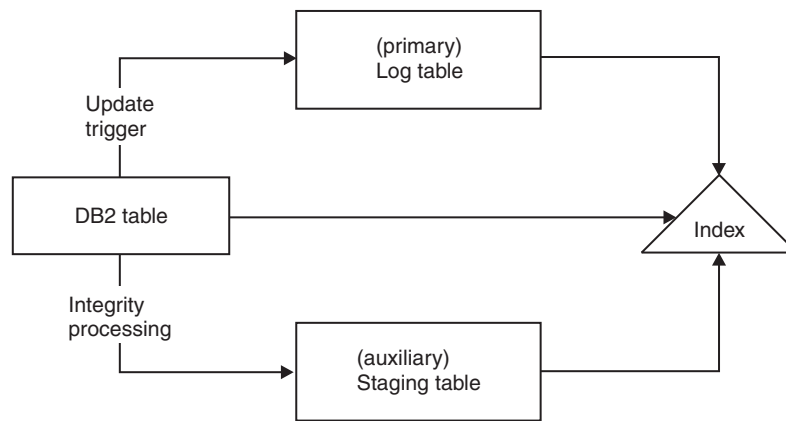


Figure 3. Incremental update with triggers and integrity processing

You can update the text index by using a *manual* or *automatic* option. The automatic option uses an update schedule with specified days and times.

Note that neither of these options synchronizes the text index within the scope of a transaction that updates, deletes, and inserts text documents. Net Search Extender's asynchronous text indexing improves performance and concurrency. The update is applied within a separate transaction to a copy of a very small part of the index. The index is only locked for read access during a very short period of time when the copy is put in place of the original. This is invisible to search operations, see Chapter 21, "Net Search Extender instance services," on page 61 for information.

A text index has certain properties, such as index file location and automatic update properties. If necessary, you can change some of the properties. This is known as *altering* the index. Altering the index does not modify any index data.

One such property is whether the ORDER BY phrase should presort the text index on the table columns. In such a case, the initial update will index the text document in the order specified, and return the search results in this order.

For example, you might specify presorted book abstracts according to the book price. When looking for the least expensive books about relational database systems, you can restrict your text search to return only the first couple of books as these will be the cheapest. However, without presorted indexes, you would have to search for all books and join these with the cheapest books, which would be a more costly operation.

Net Search Extender allows several presorted indexes per text column. For example, one index for presorting books according to the date of publication and, a second presorting books according to the price.

Usually the first update after creating a text index is an initial update, and the following updates are incremental. However, when working with presorted indexes you want to keep the order in case of updates. This is addressed by the RECREATE INDEX ON UPDATE option, which totally rebuilds the index each time an update is performed.

After the text index is updated, you can search using one of the following options:

- An SQL scalar search function
- A stored procedure search
- An SQL table-valued function

As the search options have different operating characteristics, they are explained in the following sections.

---

## SQL scalar search function overview




Figure 4. Using an SQL scalar search function for searching

Net Search Extender offers three scalar text search functions (CONTAINS, NUMBEROFMATCHES, and SCORE) that are seamlessly integrated within SQL. You can use the search functions in the same places that you would use standard SQL expressions within SQL queries. Typical queries are:

```
SELECT * FROM books WHERE CONTAINS (abstract, 'relational databases') = 1
AND PRICE <10
```

```
SELECT ISBN, SCORE (abstract, 'relational databases') as SCORE
from BOOKS
where NUMBEROFMATCHES (abstract, 'relational databases')
>5 AND PRICE <10
order by SCORE
```

The SQL scalar functions in the example return an indicator to how well the text documents matched a given text search condition. Then the SELECT phase of the SQL query determines the information returned to the end user.

Use the SQL scalar search functions as the default search method. These search functions should be suitable for a majority of situations, especially when the text search expression is combined with other conditions.

Note that the DB2 Optimizer is aware of how many text documents can be expected to match a CONTAINS predicate and how costly different access plan alternatives will be. The optimizer will choose the cheapest access plan.

---

## Stored procedure search overview

Usually, presenting search results to the end user involves a call to the search function itself that is followed by a join operation against the user table and possibly a sorting of the result data. This can be a costly operation. However, there are situations in which an application can avoid costly disk operations by performing the join operations on presorted data that is stored in memory. These situations include:

- The subset of data to present to the user is small
- The subset of data is known in advance
- The intended sort order is fixed and known in advance
- A ranked subset of the search results is sufficient

During text index creation, you must specify which columns out of the table or view are to be returned to the end user. The data is stored in a **cache** in main memory. This enables the stored procedure search to return search results extremely quickly. The cache needs to be **activated** before it can be used and there is a corresponding **deactivate** command.




Figure 5. Using a stored procedure search

The **ACTIVATE** command loads data into either a temporary cache (which is created from scratch on activation), or a persistent cache, which is maintained on disk.

The decision to use the stored procedure for search requires careful memory calculations, such as how much memory is required and how much free memory should be left for index updates. For defaults, see Chapter 15, “Stored procedure search memory requirements,” on page 47.

The stored procedure can work on text indexes that are created on views. However, as triggers can not be created on views, any changes are not automatically recognized. You must add the changed information to the log table manually, or work with the **RECREATE** option.

The main use of the stored procedure search is in high performance and high scalability applications that are interested in text-search-only queries, that is, in queries that do not need to join text search results with the results of other complex SQL conditions.

The main functional differences to the SQL scalar search functions are:

- The stored procedure search can not be used in arbitrary SQL queries, but is a query against a predefined cache table.
- The stored procedure search can exploit indexes on views.
- The stored procedure search can exploit multiple presorted text indexes on a column.

---

## SQL table-valued function overview

The SQL table-valued function is a compromise between the SQL scalar search functions and the stored procedure search. With the SQL table-valued function you can also use a `db2ext.highlight` function to get information about why a document qualified as a search result.

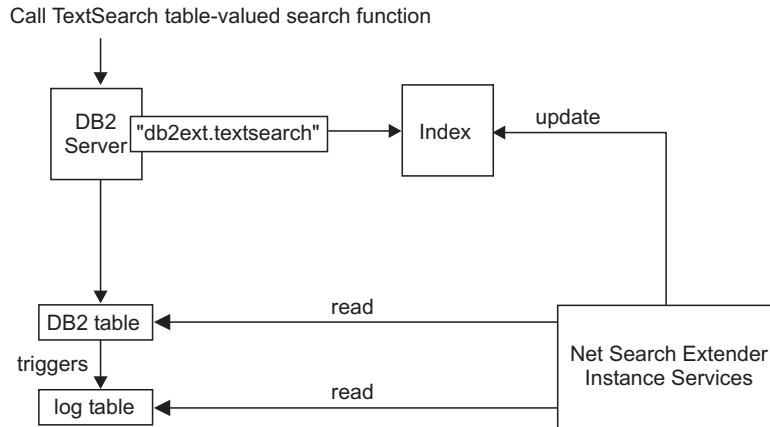


Figure 6. Using an SQL table-valued function for searching

The main functional differences to the stored procedure search are:

- No cache is necessary (and no cache is exploited).
- The table-valued function can be used in arbitrary SQL statements.
- Large amounts of memory are not required to pre-store cache table content.

The main functional difference to the SQL scalar search functions is:

- The SQL table-valued function can exploit indexes on views.

Use the SQL table-valued function in those cases where you would normally use an SQL scalar function, but you want to exploit text indexes on views.

## Additional concepts

### Column transformation function

You can use your own function to convert an unsupported format or data type into a supported format or data type. By specifying a User Defined Function (UDF), you can get the original text document as input. The output from the UDF should be a supported format, that can be processed during indexing.

You can also use this feature for indexing documents that are stored on external data stores that are not directly supported. In this case, the DB2 column contains document references and the function returns the document contents that have the relevant document reference.

### DB2 Net Search Extender instance services

Net Search Extender instance services take care of index-specific locking services and text index update services (both automatic and manual).

### DB2 Net Search Extender instance services on Windows

When a new DB2 instance is created, the DB2EXT service for the first partition is automatically created. DB2EXT services will be added subsequently when the db2nprt command is used to add new partitions. DB2EXT services will also be added when the db2start add dbpartitionnum command is executed. Similarly, the db2ndrop and db2nchg also drop or modify DB2EXT services appropriately.



## Externally stored data

In a majority of cases, the data on which you create a text index is stored within native DB2 table columns, such as in CLOBs or VARCHARs.

However, text documents that are stored externally, such as in other databases, are also supported. For documents that are stored on other databases, use DB2 nickname tables to create a text index.

You can also use the column transformation function for data that are stored in unsupported external data stores.

## Administration tables and views

There are several tables and views available in Net Search Extender. They provide information on the text indexes and their properties.

## Partitioned database support

When enabling and administering NSE in a partitioned database environment, consider the following:

- Ensure that the DB2 setup is complete as described in the DB2 documentation. The NFS mount must be configured with root access and setuid.
- If a problem occurs during `db2text start`, no detailed message is returned stating which of the available partitions is affected. If you issue `db2text start` a second time, the system tries to start the service on each of the partitions. The `db2text start` command is successful if the following message is displayed: CTE0185 The update and locking services are already active.
- The fenced user ID should be the same as the instance owner ID .
- You cannot insert a new partition number or delete an existing one from `db2nodes.cfg` while the NSE Instance Services are running. This applies to any command that might result in changes to `db2nodes.cfg`.
- On Windows platforms, while using NSE with partitioned database environment, `db2nodes.cfg` should not use IP addresses as well as host names for the same host.
- The usage of the DB2 Control Center to administer NSE in a partitioned database environment is not supported.
- To avoid unexpected results and error messages, drop existing text indexes prior to performing data redistribution. If the redistribution operation is executed without first dropping the text indexes in that database, they can still be dropped after the redistribute operation completes.

The search functions of Net Search Extender use partitioned database environment support in the following ways:

- The stored procedure search and the SQL table-valued function can operate only on tables that are local to the coordinator partition on a partitioned database.
- The SQL scalar search functions (`CONTAINS`, `NUMBEROFMATCHES`, and `SCORE`) can be used on tables spanning multiple partitions except in rare cases.

You should be aware of the following considerations when conducting searches in a partitioned database environment:

- The `RESULT LIMIT` is evaluated on every partition during search. This means that if you specify a `RESULT LIMIT` of 3 and use 4 partitions, you might get up to 12 results.

- The SCORE value reflects the document’s relevance when compared to the SCORE value of all documents from a single partition even if the query accesses multiple partitions.

## Indexes on nicknames in a federated database

You can also create a text index on nicknames in a federated database that points to tables in a remote database. In this case, the role of the log table (for incremental index updates) differs from its role for an index on a regular table. Unlike regular tables, DB2 triggers can not be created on nicknames, so that change information about documents cannot be inserted into a log table using triggers, nor can the extended text-maintained staging infrastructure be used to capture changes. Therefore, there are two different ways for incremental updates to create an index on a nickname:

- The log table is created locally in the federated database and the application is responsible for ensuring that the log table contains correct change information on the nickname. For DB2 views, this is similar to the incremental index update. This option is the default option.
- DB2 Replication has been set up so that changes to the table referenced by the nickname are captured in a so-called “Change Data Table” (CD-table) for DB2 remote databases, or a “Consistent Change Data Table” (CCD-table) for non-DB2 relational databases. DB2 Net Search Extender can then use the CD or CCD table instead of creating a log table for an index on a nickname. In this case, you must specify the capture table characteristics in the DB2TEXT CREATE INDEX command.

## Native XML support

By fully supporting the SQL XML data type in databases with codepage UTF-8, all Net Search Extender search functions can be used on XML documents that are stored natively in the database. Note that the text search on SQL XML data types in databases with non-UTF-8 is not supported and might not return any results.

The structural text search on XML documents by sections (see “Search parameters” on page 234 for more information on searching in **sections**) can be extended by powerful XQuery processing of the search results. Net Search Extender text search functionality can be leveraged within DB2’s XQuery language support to provide optimal processing of XML documents.

By using full text search within the `db2-fn:sqlquery()` XQuery input function, it is possible to search in XML documents, and process the resulting XML documents using XQuery:

```
FOR $dept in db2-fn:sqlquery('select Department from MyTable
  where contains(Department, 'sections(/dept/employee/resume) "DB2 XML" ')
  = 1')/dept
RETURN $dept/employee/name
```

In the example, column “Department” is of data type “XML”. See Part 8, “Working with structured documents,” on page 157 for more information.

A sample is available that shows you how you can query XML data. Refer to `<sqllib>/samples/extenders/db2ext`. Call `xmlsample <database>` to populate the database, and create and update the indexes. After you have connected to the database, you can perform searches on the data by issuing `db2 -tvf xmlsearch`.

## Partitioned table support

You can create a text index for range-partitioned tables or tables that use the multidimensional clustering feature in single-partition and multi-partition database environments. Text indexes are supported for any partitioning feature combination.

**Note:** For Version 9.7, the text index will be partitioned according to the partitioning of the table across multiple database partitions. Other partitioning features, like range-partitioning or multidimensional clustering, have no impact on the partitioning of the text index

## Incremental update based on integrity processing

By using the **AUXLOG** option for the Net Search Extender **CREATE INDEX** command, you can control whether an auxiliary log infrastructure (a text-maintained staging table) is used for a text index.

This auxiliary staging table captures information about new and deleted document through integrity processing, while document updates are captured through an update trigger on the base table column referenced in the **INDEX CREATE** command.

The following restrictions apply to using the option:

- The object for which you created the text index must be a base table, not a view or a nickname.
- You cannot specify the **CACHE** option
- You cannot use the **RECREATE INDEX ON UPDATE** option.
- You cannot use replication to control updates.

By default this configuration option is set to **ON** for range-partitioned tables and the configuration option is set to **OFF** for non-partitioned tables.

Capturing changes for an incremental update of the text index through integrity processing might require you to perform additional administrative tasks after performing a database operation on the base table. You might have to perform post-processing tasks for the command or preprocessing tasks for a text index update to set integrity for the base table or its dependent tables.

### Example

```
db2 "create table test.simple (pk integer not null primary key,
comment varchar(48))"
db2 "insert into test.simple values (1, 'blue and red')"
db2text "create index test.simpleix for text on test.simple(comment) index
configuration(auxlog on) connect to mydb"
db2text "update index test.simpleix for text connect to mydb"
db2 "load from loaddata4.sql of del insert into test.simple"
```

After the load operation, the base table is blocked. For example, a select operation will result in **SQL0668N Operation not allowed for reason code "1" on table "TEST.SIMPLE". SQLSTATE=57016.**

The staging table is accessible, but does not yet contain the information about the changed data.

```
db2 "set integrity for test.simple immediate checked"
```

returns SQL3601W The statement caused one or more tables to automatically be placed in the Set Integrity Pending state. SQLSTATE=01586.

At this point the staging table is blocked and modifying operations for the base table are rejected.

```
"insert into test.simple values(15, 'green')"
```

returns DB21034E The command was processed as an SQL statement because it was not a valid command line processor command. During SQL processing it returned: SQL0668N Operation not allowed for reason code "1" on table "SYSIBMTS" ."SYSTSAUXLOG\_IX114555". SQLSTATE=57016.

```
db2text "reset pending for table test.simple for text connect to mydb"
```

```
db2text "update index test.simpleix for text connect to mydb"
```

For more details see Part 4, "Planning considerations," on page 41

---

## Chapter 2. Key features of DB2 Net Search Extender

Net Search Extender Version 9.7 has the following key features:

- Indexing
  - Fast indexing of very large data volumes
  - Dynamic updating of indexes
  - Optional: Storing table columns in main memory at indexing time to avoid expensive physical read operations at search time
  - Support for structured text formats, for example HTML and XML
  - Support of third-party filtering software "Outside In"
  - Nickname table support
  - Support of presorted text indexes
  - Partitioned database support
  - Native XML support
  - Support for range-partitioned tables and clustered tables (MDC)
- Search
  - Boolean operations
  - Proximity search for words in the same sentence or paragraph
  - "Fuzzy" searches for words having a similar spelling as the search term
  - Wildcard searches, using front, middle, and end masking, for whole words and single characters
  - Free-text searches. For documents containing specific text, the search argument is expressed in natural language
  - A highlight function to show why a particular document qualified as a search result
  - Thesaurus support
  - Restrict searching to sections within documents
  - Numeric attribute support
  - High-speed searching through a large number of text documents with many concurrent users
  - Integration into XQuery processing using the `db2-fn:sqlquery()` function.
- Search results
  - You can specify how the search results are sorted at indexing time
  - You can specify search result subsets when searching large data volumes and expecting large result lists
  - You can set a limit on search terms with a high hit count
  - Built-in SQL functionality combined with the DB2 Optimizer automatically selects the best plan according to the expected search results



---

## Chapter 3. Introducing the db2text commands

Here is an example of a Net Search Extender command:

```
db2text ENABLE DATABASE FOR TEXT
```

For every create and index maintenance command, you can specify the database, user, and password.

```
db2text ... connect TO <database> USER <userID> USING <password>
```

**Note:** if you leave out the connect options in the db2text command, the environment variable DB2DBDFT specifies the database.

To display a list of commands, enter the following command:

```
db2text ?
```

To display the syntax of an individual command, enter the following command:

```
db2text ? command
```

For example, to display the syntax of the CREATE INDEX command, use the following command:

```
db2text ? CREATE INDEX
```

db2text returns 0 if the command has been processed successfully, and 1 if the command has not been processed. Note that if there are document errors but the index still updates, the db2text command returns 0 with a warning message. Information about document errors can be found in the event table of the index.

Depending on your operating system and your active command shell, the system interprets special characters such as ?, (, ), \*, !, and ". Therefore, if the command contains these characters, use quotation marks or an escape character.

Here is an example of a UNIX<sup>®</sup> command that uses special characters:

```
db2 "SELECT * FROM sample WHERE CONTAINS (DESCRIPTION, '\"enable\"') = 1"
```





---

## Part 2. Installation



---

## **Chapter 4. Installing Net Search Extender in the DB2 client/server environment**

Net Search Extender search functionality is integrated into SQL and executed on the server. Therefore, you do not need to install Net Search Extender on the client to issue text search queries.

If you plan to administer Net Search Extender from a remote client, you have to have one of the DB2 server editions and Net Search Extender itself installed on the client side.



---

## Chapter 5. Net Search Extender installation system requirements

The minimum hardware and software requirements for Net Search Extender–supported platforms for DB2 Version 9.7.

You need to install DB2 Version 9.7 before you can run Net Search Extender.

Net Search Extender is supported on the following platforms:

- AIX® (64-bit) platforms
- Linux® x86 (32-bit) platforms
- Linux x86-64 (64-bit) platforms
- Linux on zSeries® (64-bit) platforms
- HP-UX on Itanium-based HP Integrity Series (64-bit) platforms
- Solaris UltraSPARC (64-bit) platforms
- Windows® on x86 (32-bit) platforms
- Windows on x86-64 (64-bit) platforms

**Note:** Partitioned database environment is not supported by Net Search Extender for (32-bit) platforms.

The minimum disk space for a typical Net Search Extender installation is 50 MB. Additional hardware requirements might be required depending on the amount of data you are planning to index.



---

## Chapter 6. Installing NSE on a partitioned DB2 server

Ensure the correct installation and configuration of DB2 on every partition. After installing DB2, you need to install Net Search Extender on each partition.

**Note:** A fenced user ID that is different from the instance owner ID does not work with partitioned databases.





---

## Chapter 7. Installation on UNIX

To install on UNIX, follow these steps:

1. Install the product.
2. Update the DB2 instance.

### Step 1 for UNIX: Install the product components

To install on UNIX, follow these steps:

1. Log on at the target machine as the root user.
2. Change to the correct directory for your platform:
  - `cd /<cdrom>` where `<cdrom>` is your CD-ROM driver path.
  - `cd <platform>`
3. Call `./nsetup.sh` and follow the instructions displayed on the screen.

After you accept the license agreement, you are presented with a list of possible installation paths. The installation paths that are eligible for you to use depend on the Net Search Extender version you want to install and the installed copies of DB2. After you select a path for the installation, the product is installed at this path. An installation log file is written to the `/tmp` directory and prefixed `db2nsei`.

### Step 2 for UNIX: Update the DB2 instance

To update the DB2 instance, follow these steps:

1. Ensure that you are active as the root user.
2. Use the following command to change your working directory to the path where you installed Net Search Extender:

```
cd <path>/instance
```

`<path>` is the path of the DB2 copy where you installed Net Search Extender.
3. Run `db2iupdt` using `./db2iupdt <db2instance>`, where `<db2instance>` is the name of an existing DB2 instance user ID that you would like to use with Net Search Extender.



---

## Chapter 8. Installation on Windows

### About this task

This method uses some command-line options. The command-line options that require a parameter must be specified with no space between the option and its parameter.

### Procedure

To install on Windows you must log on with a user ID that has administrative rights, and then follow these steps:

1. Use the `cdrom:\windows\install\setup.exe` to transfer the files from the package to the target machine.

Note that for every DB2 instance, you must enter a user ID and password to create the correct Net Search Extender services.

After you accept the license agreement, you are presented with a list of possible DB2 copy names. The DB2 copy names that are eligible for you to use depend on the Net Search Extender version you want to install. After you select a DB2 copy name, the product is installed at the path where this DB2 copy has been already installed.

For silent install, you can call `setup.exe` in two different modes:

#### RECORDMODE

Creates a silent install response file.

A predefined response file called `setup.iss` is located in the installation source directory (this is not used in the case of silent installation). If you want to create a new response file, run `setup.exe -r`. The new `setup.iss` is created in your Windows directory. Copy this `setup.iss` into your installation source directory. Make sure that you have made a backup of your old response file.

Example: `setup.exe -r -f1"d:\some_directory\setup.iss"`

#### SILENTMODE

Silent installation.

Make sure that the file `setup.iss` is located in your installation source directory. Run `setup.exe -s`. The installation is successful if `ResponseResult` is set to 0 in the `setup.log` file located in your installation source directory.

Example: `setup.exe -s -f1"d:\some_directory\setup.iss" -f2"d:\another_directory\mysetup.log"`

Installation parameters:

- `/r`: Record mode (records a response file for silent installations)
  - `/s`: Silent install (runs installation silently)
  - `/f1`: Specify alternative response file name (full path)
  - `/f2`: Specify alternative setup log file name (full path)
2. Reboot the system after the installation.
  3. Call `db2text start` to start the DB2 Net Search Extender Instance Services.

## Results

Every DB2 instance creates a Windows service. Make sure that the DB2 instance services run under a user account and not under the systems account.

For a partitioned instance, a set of Windows services are created, one per partition. The DB2EXT instance services for every partition must run under the same user account as the DB2 instance services.

Net Search Extender does not support Microsoft Cluster Server.

---

## Chapter 9. Directory names and file names

You must specify the directory names and file names in SBCS characters for all Net Search Extender commands. The maximum length of the path names (including the file name) is 256 bytes.



---

## Chapter 10. Installing the Outside In libraries

The Outside In software from Stellent is available for several platforms. For details please see <http://www.oracle.com>

To use Net Search Extender with the Outside In software from Stellent, you must set up the libraries for each platform:

- On Windows, ensure that the directory where the libraries are located is added to the path environment variable.
- On UNIX, add the Outside In libraries into the DB2 lib install directory.





---

## Chapter 11. Installation verification

---

### Installation verification on UNIX

Complete the following steps to verify that Net Search Extender is correctly installed.

- Follow these steps to call the administration script `nseample` to set up the text indexes:
  1. Change to `<instance_owner_home>/sqllib/samples/extenders/db2ext`
  2. Call `./nseample <yourdb>`. Note that this command creates the database if it does not already exist.
  3. Check the generated output file `nseample.log` in your home directory.
- Then call some sample queries to execute in the same DB2 command window:
  1. Connect to your database using `db2 connect to <yourdb>`
  2. Execute the sample queries using `db2 -tvf search`
  3. Check the results of the queries contained in the script. Note that every query should return one or more hits.

If there are no errors in the `nseample_partitioned.log` file and all the queries are working, Net Search Extender is successfully installed.

In a partitioned database, use the following verification sample:

```
nseample_partitioned database_name [node_number][table_space_filename]
```

---

### Installation verification on Windows

Complete the following steps to verify that Net Search Extender is correctly installed.

#### Non partitioned environment

1. Set up the sample text indexes as follows:
  - a. Call `db2cmd` to open a DB2 command window.
  - b. Change to `<sqllib>\samples\extenders\db2ext`
  - c. From the DB2 command window, call `nseample.bat <yourdb>`. Note that this command creates the database if it does not already exist.
  - d. Check the generated output file `nseample.log` in the current directory.
2. Call the following sample queries to execute in the DB2 command window:
  - a. Connect to your database using `db2 connect to <yourdb>`
  - b. Execute the sample queries using `db2 -tvf search`
  - c. Check the results of the queries contained in the script. Note that every query should return one or more hits.

If there are no errors in the `nseample.log` file and all the queries are working, Net Search Extender is successfully installed.

#### Partitioned database environment

1. Set up the sample text indexes as follows:
  - a. Call `db2cmd` to open a DB2 command window.

- b. Change to <sqllib>\samples\extenders\db2ext
- c. From the DB2 command window, call nsesample\_partitioned.bat <yourdb>.
- d. Check the generated output file nsesample\_partitioned.log in the current directory.

If there are no errors in the nsesample\_partitioned.log file and all the queries are working, Net Search Extender is successfully installed.

---

## Chapter 12. Uninstalling Net Search Extender

To permanently remove Net Search Extender from your system and remove all Net Search Extender indexes, you must first disable each database that contains Net Search Extender indexes and then only remove Net Search Extender.

---

### Uninstalling Net Search Extender on UNIX

Complete the following steps to uninstall Net Search Extender correctly on UNIX.

- For each DB2 instance from which you want to uninstall Net Search Extender:
  1. Switch to the user ID of the DB2 instance
  2. If you no longer intend to use Net Search Extender on this instance, you should drop the indexes and disable the database prior to uninstalling Net Search Extender:

```
db2text disable database for text connect to <databasename>
```
  3. Stop the DB2 Net Search Extender instance
  4. Stop the DB2 instance
- Ensure that you are active as the root user.
- Change your working directory to the DB2 path where you want to remove Net Search Extender. For example, `cd /opt/IBM/db2/V9.5/install`.
- Issue the command `./db2nse_deinstall`. For details on the command syntax, see “`db2nse_deinstall` command” on page 228.

---

### Uninstalling Net Search Extender on Windows

Complete the following steps to uninstall Net Search Extender correctly on Windows.

1. For each database run `db2text disable database for text connect to <databasename>`.
2. Stop the DB2 instance.
3. Select **Settings->Control Panel ->Add or Remove Programs**. From the list, select the Net Search Extender<COPYNAME> entry that correlates with the DB2<COPYNAME> that Net Search Extender was assigned during installation. Click **Remove**.



---

## Part 3. Migrating to DB2 Net Search Extender Version 9.7

Before migrating Net Search Extender to DB2 Net Search Extender Version 9.7, you must have successfully migrated your DB2 server and instance (on Linux and UNIX) including all databases to DB2 Version 9.7

Migration to Net Search Extender Version 9.7 is supported from Net Search Extender Version 8, Version 9.1, or Version 9.5. The term pre-Version 9.7 refers to Version 8, Version 9.1, and Version 9.5.

### Prerequisites

- Back up all text index directories and index subdirectories before migrating.

### Procedure

To migrate to DB2 Net Search Extender Version 9.7:

1. The following drop statements are only necessary when migrating from Version 8. Execute the following calls before you migrate to Version 9.7 because the Data Links Manager is unsupported since DB2 Version 9.1:

```
db2 DROP SPECIFIC FUNCTION DB2EXT.DATALINKCONTENT1;  
db2 DROP SPECIFIC FUNCTION DB2EXT.DATALINKCONTENT2;  
db2 DROP SPECIFIC FUNCTION DB2EXT.DATALINKCONTENT4;  
db2 DROP SPECIFIC FUNCTION DB2EXT.DATALINKCONTENT3;
```

**Note:** If you have text indexes that you created using an earlier version of Net Search Extender that includes the Data Links product features, you can still search these indexes but you cannot update these indexes. If you still need to create an index on text stored outside the database, write a UDF that you can plug into Net Search Extender as a transformation function.

2. Upgrade your DB2 server where Net Search Extender is installed using any of the following tasks:
  - Upgrading DB2 servers (Windows) found in the *Upgrading to DB2 Version 9.7*
  - Upgrading DB2 32-bit servers to 64-bit systems (Windows) found in the *Upgrading to DB2 Version 9.7*
  - Upgrading a DB2 server (Linux and UNIX) found in the *Upgrading to DB2 Version 9.7*

Upgrading your database is part of these tasks. If you have external unfenced routines on Linux or UNIX that have no dependency on the DB2 engine libraries, the UPGRADE DATABASE command redefines your external routines as FENCED and NOT THREADSAFE. See Upgrading C, C++, and COBOL routines in Upgrading to DB2 Version 9.7 for details on how to safely run your routines in the new multithreaded database manager. The Net Search Extender functions with schema name DB2EXT that were altered during the database upgrade are redefined as NOT FENCED and THREADSAFE by the db2extmdb migration script in step 6.

3. Install DB2 Net Search Extender Version 9.7.

Unlike DB2 database, DB2 Net Search Extender does not support an 'upgrade installation'.

If your installed DB2 copy was migrated to Version 9.7 by using the DB2 'upgrade installation' option, the installed DB2 Net Search Extender copy is still at the previous version level.

If you try to install DB2 Net Search Extender Version 9.7 on top of an earlier version of DB2 Net Search Extender, you receive an error message that the existing DB2 Net Search Extender installation has to be removed first. In this case, uninstall the earlier version of DB2 Net Search Extender before installing DB2 Net Search Extender Version 9.7. On Windows operating systems, reboot the machine after uninstalling DB2 Net Search Extender.

4. A DB2 Net Search Extender instance migration is only applicable on Linux and UNIX. This step has to be ignored on Windows operating systems. To migrate the instance, log on as root and run the db2extimigr script using the following syntax:

```
DB2DIR/instance/db2extimigr [-h|-?] InstanceName
```

Where DB2DIR is the directory where you installed your DB2 Version 9.7 copy.

**Note:** If your pre-version 9.7 installation did not have indexes in the default index directory, the db2extimigr tool might end with the error message `"/db2extimigr[24]: OLDSQLLIB: Parameter null or not set."` you can resolve this problem by first manually exporting the OLDSQLLIB variable and then running db2extimigr again.

The export statement depends on the DB2 pre-version you upgraded in step 2.

For pre-version 8.1:

```
OLDSQLLIB=sqllib_v81  
export OLDSQLLIB
```

For pre-version 9.1:

```
OLDSQLLIB=sqllib_v91  
export OLDSQLLIB
```

For pre-version 9.5:

```
OLDSQLLIB=sqllib_v95  
export OLDSQLLIB
```

5. On Linux and UNIX, after a successful DB2 Net Search Extender instance migration, verify the Version 9.7 installation before you continue with the database migration. On Windows operating systems, you can verify the Version 9.7 installation immediately.

**Note:** Do not apply the DB2 Net Search Extender sample scripts to a database that is not migrated to DB2 Net Search Extender Version 9.7. A secure way is to create a database to carry out the verification. See Installation verification on UNIX and Installation verification on Windows for more details.

6. To complete the migration, migrate every database that had been enabled for Net Search Extender in a pre-Version 9.7 release. Step a) and c) are currently needed to resolve a known problem with db2extmdb (duplicate entries in the view DB2EXT.DBDEFAULTS after database migration). To perform the database migration steps:
  - a. Log on to the DB2 server as the instance owner.

You must be able to successfully stop or start the Net Search Extender instance services, you need DBADM with DATAACCESS authority on the database to be migrated. On Windows operating systems, the instance user must be part of the Local Administrator group.

On Windows operating systems, you need to proceed from a DB2 command window that is running with full administrative privileges. See User Access Control feature for more details.

On the Windows 2008 and Windows Vista operating system (and later), the db2extmdb command does not work if it is started from a command window running with standard user rights. ("CTE0228 The user has insufficient access rights at the operating system level").

- b. Save the content of the base table of the DB2EXT.DBDEFAULTS view before you perform the database migration by running the following statements:

```
db2 "CREATE TABLE DBDEFAULTSCOPY LIKE DB2EXT.TDBDEFAULTS";
db2 "INSERT INTO DBDEFAULTSCOPY SELECT * FROM DB2EXT.TDBDEFAULTS";
```

- c. Run the db2extmdb migration script to migrate the database that you enabled for Net Search Extender using the following syntax:

```
db2extmdb <database-name>
```

While you are running this migration script, avoid changing user tables with text indexes. You can repeat the command for every database that has Net Search Extender indexes.

All migration steps are logged in the file called db2extm<database-name>.log located in one of the following directories:

- INSTHOME/sql11ib/db2ext/ on Linux and UNIX operating systems
- DB2PATH\db2ext\ on Windows operating systems

Where INSTHOME is the instance home directory and DB2PATH is the location where you installed your DB2 Version 9.7 copy.

- d. Clean up the DB2EXT.DBDEFAULTS view after completion of the database migration by running the following statements:

```
db2 "DELETE FROM (SELECT ROW_NUMBER() OVER
(PARTITION BY NAME) AS COUNT FROM DB2EXT.TDBDEFAULTS) WHERE COUNT > 1";
db2 "UPDATE DB2EXT.TDBDEFAULTS T1 SET T1.VALUE = (SELECT S1.VALUE FROM
DBDEFAULTSCOPY S1
WHERE T1.NAME = S1.NAME AND T1.VALUE <> S1.VALUE FETCH FIRST 1 ROWS ONLY)
WHERE T1.NAME IN (SELECT DISTINCT S2.NAME FROM DBDEFAULTSCOPY S2,
DB2EXT.TDBDEFAULTS T2
WHERE T2.NAME <> 'VERSION' AND T2.NAME = S2.NAME AND T2.VALUE <> S2.VALUE)";
```

Both statements might return a warning like "SQL0100W No row was found for FETCH, UPDATE or DELETE". In these cases, you can regard the execution as successful.

7. If you upgraded from pre-Version 9.7 DB2 32-bit server to a Version 9.7 64-bit server, you must drop your text indexes and re-create your text indexes. In Net Search Extender, you cannot use in a 64-bit instance text indexes that you created in a 32-bit instance. The search engine returns error CTE0101 Reason code: "17".
8. If you want to use text indexes that you created under the installation directories of your pre-Version 9.7 DB2 copies before DB2 server upgrade on Windows operating systems, restore the text index directories that you backed up. Restoring the text index directories is required if you chose a DB2 copy with the upgrade action in the **Work with existing** window during DB2 Version 9.7 installation or you uninstalled your pre-Version 9.7 DB2 copies after migration.

The text index configuration contains the location of these text index directories before migration. Queries and index administration operations that use these text indexes fail if you do not restore the text index directories.





---

## Part 4. Planning considerations

To use Net Search Extender in the most effective way, it is essential that some planning occurs prior to deployment. Planning might involve several user groups including those in database administration, interface and system designers, system architects, and developers.

The following topics provide a guide to the areas that you should consider:

- Directory locations and index storage
- Table, column, and index names
- Document formats and supported code pages
- Outside In filtering software
- User roles
- Extended text-maintained staging infrastructure for incremental update

For more information on developing Net Search Extender based applications, see the following related topics:



---

## Chapter 13. Directory locations and index storage for Net Search Extender

The disk space you need for a Net Search Extender index depends on the amount and type of data you want to index. As a guideline for indexing single-byte documents, reserve disk space of about 0.7 times the size of the documents you want to index. For double-byte documents, reserve the same disk space as the total size of the documents you want to index. The total size might have to include data stored outside the active database that is retrieved through user-defined functions.

The amount of space needed for the temporary files in the work directory is 1.0 to 4.0 times the amount of space needed for the final index file in the index directory. The default index directory is a subdirectory of the DB2 instance directory which is usually located in the /home partition of the system (for Linux and UNIX operating systems) and below the C: drive (for Windows operating systems). The default index might have size limitations as well. See Views for database-level information for more details.

If you have several large indexes, store them on separate disk devices, especially if you have concurrent access to the indexes during index update or search.

For a partitioned database, a text index still uses a single file system on a physical machine to place the index. For large databases, place the index and work directories on a file system that is on a Redundant Array of Independent Disks (RAID) device. This minimizes the possibility of hitting an I/O bottleneck while using the text indexes.

For creating, updating, and deleting Net Search Extender indexes, you can either use the command line interface, or the DB2 Control Center.



---

## Chapter 14. Resource considerations for a partitioned DB2 server

When running NSE in a partitioned environment, you should keep the following resource considerations in mind:

- When multiple text indexes are updated in parallel, this can lead to considerable burst disk usage during the I/O intensive phases of the index update. This requires careful consideration during the setup of partitions, the associated NSE index storage, and scheduling of index updates.
- For NSE administration commands, there are processes started for every partition to carry out operations such as index create, drop, and update. The update process can be a longer running process that will consume resources. For scheduled updates of text indexes in the system, several index updates running concurrently can lead to as many processes per partition as the number of indexes being updated at a given time. It would be advisable to minimize the number of concurrent index updates by planning the index update schedule accordingly.
- On Linux or UNIX platforms, you can use the command `ulimit` with the appropriate options to view or change the size of the process resource limit. If you are running DB2 Net Search Extender in a partitioned database environment, use the command `db2_all ulimit` with the specific options to see the size that applies to all the database partitions. It is essential to verify the hard operating system ulimits for successfully running Net Search Extender update index commands. For example, insufficient data segment size on any of the partitions can cause failures during index update commands most usually with CTE0105 error logged in the event view of the text index.



---

## Chapter 15. Stored procedure search memory requirements

Using the cache for a stored procedure search requires a large amount of memory and different memory requirements for the following platforms:

- AIX
- Windows
- Solaris
- Linux

---

### Memory requirements for AIX (64-bit)

Configuring the system limits:

- Check the system limits using the command `ulimit -a`
- If there are values other than "unlimited", use the following steps:
  - Log on as root.
  - Back up the file `/etc/security/limits` and then edit the file to raise the hard limits.
  - Set all values to the "unlimited" (value -1) for the DB2 instance owner used.

Configuring the shared memory limits:

- On AIX, there is no need to configure the shared memory limits.

Configuring the swap space:

- Obtain the system RAM size using the command `lsattr -E -l sys0`
- Obtain the size of the swap space using the `lspv -a` command.
- Set the swap space size to at least 1.5 - 2 times of either the RAM amount of your system, or use the `MAXIMUM CACHE SIZE` parameter you provide in the `CREATE INDEX` command. Use the `SMIT` utility to select the larger number.

---

### Memory requirements for Windows (32-bit and 64-bit)

Adjusting the size of the paging file:

- Set the Windows virtual-memory paging file size to at least 1.5 - 2 times of either the RAM amount of your system, or use the `MAXIMUM CACHE SIZE` parameter you provide in the `CREATE INDEX` command. Select the larger number. See the Windows documentation for information on changing the size of the paging file.

On Windows 32-bit, you are recommended to not exceed a maximum cache size of about 1000 MB (1 GB = 1073741824 bytes).

---

### Memory requirements for Solaris (64-bit)

Configuring the system limits:

- Check the system limits using the command: `ulimit -a`
- Then use the following steps:
  - Log on as root.
  - Back up the file `/etc/system` and then edit the file to raise the hard limits.
  - Add or check that the following lines are set to at least the minimum values shown:

```
rlim_fd_cur -> Default 64, recommended >= 1024
rlim_fd_cur_max -> Default 1024, recommended >= 4096
```

Configuring the shared memory limits:

- Check current settings using command `sysdef -i`
- Edit the file `/etc/system` to set the shared memory size limit using: `set shmsys:shminfo_shmmax=0xffffffff`

You may also have to increase the following parameter values:

```
set shmsys:shminfo_shmmni=512
```

```
set shmsys:shminfo_shmseg=128 and then reboot the system.
```

Configuring the swap space:

- Obtain the system RAM size using the command `/usr/sbin/prtconf`
- Obtain the size of the swap space using the `swap -l` command.
- Set the swap space size to at least 1.5 - 2 times of either the RAM amount in your system, or use the `MAXIMUM CACHE SIZE` parameter you provide in the `CREATE INDEX` command. Select the larger number.

Refer to the Solaris system documentation for information on how to add swap space.

You are recommended not to exceed the maximum cache size of about 2000 MB (2 GB = 2147483647 bytes).

---

## Memory requirements for Linux (32-bit and 64-bit)

Check the DB2 documentation for information about recommended kernel parameters on Linux.

The validation status for new Linux kernels and distributions is frequently updated. To obtain the latest information for supported Linux software levels, refer to: <http://www.ibm.com/software/data/db2/linux/validate>

To see your current shared resource limits use `ipcs -l`. To check the system limits, use the `ulimit -a` command.



---

## Chapter 16. Table, column and index name considerations

All table, column and index names are usually case-insensitive. Net Search Extender allows you to specify these names in mixed case too. On Windows, if you want to specify table, column and index names in mixed case, you must enter the name in a backslash (\) double quote (") character sequence. For example, `\DocTxt\`.



---

## Chapter 17. Document formats and supported code pages

Net Search Extender needs to know the format (or type) of text documents that you intend to search. This information is necessary for indexing text documents.

Net Search Extender supports the following document formats:

**TEXT** Plain text (for example, flat ASCII), in general, text without any markup

**HTML**

Hypertext Markup Language

**XML** Extended Markup Language

Document format XML is the default for column data type XML, and is the only supported document format for that data type.

**GPP** General Purpose Format (flat text with user-defined tags)

**Outside In (INSO)**

Use this format if you are using filtering software to extract textual content from PDFs and other common text formatting tools, for example, Microsoft Word.

For the document formats HTML, XML, GPP, and the Outside In filter formats, searching can be restricted to specific parts of a document.

Where Outside In filters can not be used because the format of your document is not supported, you can write a User Defined Function (UDF) that does its own filtering. This UDF must be specified at index creation time and converts the data from the unsupported format to a supported format.

You can index documents if they are stored in one of the supported Coded Character Set Identifiers (CCSIDs). See the DB2 documentation for a list of these code pages.

To check the database code page, use the following DB2 command:

```
db2 GET DB CFG for <dbname>
```

and take the value written for Database code page.

For consistency, DB2 normally converts the code page of a document to the code page of the database. However, when you store data in a DB2 database in a column with a binary data type, such as BLOB or FOR BIT DATA, DB2 does not convert the data, and the documents retain their original CCSIDs.

Note that incompatible code pages might cause problems when creating a text index or searching.



---

## Chapter 18. Outside In filtering software

Net Search Extender supports a third-party document filtering software. Known as Outside In Transformation Technology from Oracle, you can use the software for extracting the textual contents from PDF files, or from documents written in the proprietary format of common text formatting tools without using native applications. Example formats include Microsoft Word and Lotus® Word Pro®.

Net Search Extender loads the Outside In libraries as plug-ins during UPDATE INDEX. The libraries are not part of Net Search Extender and need to be installed separately. You need to ensure that Net Search Extender can find the Outside In libraries.

The Outside In software generates not only textual content but also structural information, for example, fields. Net Search Extender can also customize which part of the Outside In generated document information is to be stored in the index. To do this, you need to apply a specific type of document model, the Outside In document model.

To view a list of filtering formats and supported platforms, visit the Oracle Web site at <http://www.oracle.com>.



---

## Chapter 19. User roles

### DB2 instance owner

The DB2 instance owner user can start and stop the instance services for DB2 Net Search Extender and control the locking services. In addition, the DB2 instance user is granted DBADM authority for each enabled database. This enables a central point of control for all database changes driven by Net Search Extender.

#### Required DB2 authorizations

The SECADM must grant DBADM with DATAACCESS privilege to the instance owner. These authorizations are a prerequisite for the execution of DB2 Net Search Extender administrative commands.

#### Required file system authorizations

Read and write access for all text index directories and read access to model files.

#### Commands for the instance owner

DB2TEXT START, DB2TEXT STOP, DB2TEXT CONTROL and DB2EXTHL

The commands are only allowed on the server. In a partitioned database environment, this can be any of the configured nodes. Each command checks if the user running the command is the DB2 instance owner. If you, as the instance owner, decide to use a fenced user ID to run the stored procedure and UDFs, the fenced user must have read and write access to all files in the index directory (with read access to the entire directory path). Be aware that your fenced user ID and instance user ID must be members of the same primary group to grant the instance user ID correct access to files created by the fenced user ID and vice versa. Assign correct group membership and file permissions.

### Database administrators

Database administrators can enable and disable databases for use with Net Search Extender.

#### Required DB2 authorizations

DBADM

#### Commands for the database administrator

DB2TEXT ENABLE DATABASE and DB2TEXT DISABLE DATABASE.

### Text table owners

The text table owner can create, drop, and change indexes. Note that they must be able to control (by having read and write access) the location of indexes and updates to the full-text indexes.

#### Required DB2 authorizations and privileges

Owner of text table.

#### Commands for the text table owner:

DB2TEXT CREATE INDEX, DB2TEXT DROP INDEX, DB2TEXT ALTER INDEX, DB2TEXT ACTIVATE CACHE, DB2TEXT DEACTIVATE CACHE, DB2TEXT UPDATE INDEX, DB2TEXT CLEAR EVENTS, and DB2EXTTH.

Note that the command implementation partially runs under the user ID of the DB2 instance owner. Therefore, grant the instance owner the necessary file system access before creating or altering the text indexes. For details on required permissions as listed for each command see Chapter 49, “Administration commands for the text table owner,” on page 197



---

## Chapter 20. Extended text-maintained staging infrastructure for incremental update

A configuration option is available in Version 9.7 to add a staging infrastructure that enables capturing changes that are not recognized through the triggers in the regular log table. If this option is enabled, updates are captured through a trigger in the regular log table, inserts and deletes are captured in the text-maintained staging table.

This configuration option is enabled by default for range-partitioned tables and disabled for non-partitioned tables. Adding the text-maintained staging infrastructure has a critical impact on the availability and status of the base table for various database operations.

The impact of the text-maintained staging infrastructure is similar to adding a Materialized Query Table (MQT) with deferred refresh. Even though the text-maintained infrastructure does not maintain data in a MQT, the staging table causes corresponding behavior to that shown for a MQT staging table.

For example, after a LOAD insert, the tables require integrity processing to enable subsequent database operations on the base table.

If tables are only updated with database commands that affect all rows in the table, for example with LOAD REPLACE, adding the extended staging infrastructure does not provide a benefit, instead, the index should be re-created.



---

## **Part 5. Administering Net Search Extender**



---

## Chapter 21. Net Search Extender instance services

DB2 Net Search Extender Instance Services consist of the following services:

- Locking services
- Update services

DB2 Net Search Extender Instance Services on Windows are represented by Windows services. On a non-partitioned DB2 instance there is one such service per DB2 instance with a service name:

DB2EXT - *<instance\_name>*

On a partitioned DB2 instance there is one such service for each partition of the DB2 instance with a service name:

DB2EXT - *<instance\_name>*[-*<nodenum>*]

The following topics explain how to start and stop the DB2 Net Search Extender Instance Services and further discusses Locking Services and Update Services in detail:

- Starting and stopping NSE instance services
- Locking services
- Update services
- NSE information catalogues



---

## Chapter 22. Starting and stopping Net Search Extender instance services using the command line

Before you can maintain text indexes and search your documents, you have to start the Net Search Extender Instance Services.

To start the Instance Services, log on to the DB2 instance owner user ID (UNIX systems only) and enter the following command:

```
db2text start
```

To stop the Instance Services, enter the following command:

```
db2text stop
```

For DB2 instances used with partitioned databases, it is highly recommended that the Net Search Extender Instance Services be started and stopped using `db2text start` and `db2text stop` instead of using the normal Windows methods. This ensures that the Instance Services are started and stopped in the correct order.

**Note:**

- There must be one Net Search Extender Instance Service per DB2 instance. The locking service maintains the locks for all enabled databases for that instance.
- DB2 Net Search Extender Instance Services on Windows are represented by Windows services. On a partitioned DB2 instance there is one such service for each partition of the DB2 instance.





---

## Chapter 23. NSE locking services

When you start Net Search Extender, the locking services automatically start. The locking services are needed to synchronize concurrent access to text indexes in Net Search Extender.

The locking services ensure that no two processes attempt to change a text index simultaneously, or that no process reads text index data while another process is making changes to the same text index data. Therefore, most processes request a lock on a text index before starting, and release it again when processing has completed.

Note that the locking services for Net Search Extender text indexes must not be confused with DB2 locks that control access to DB2 tables.

---

### Using the locking services

In Net Search Extender, there are different types of locks that control concurrent access to an index. The different locks depend on whether the text index is only being read, as in the case of a search request, or if changes to the text index need to be computed and subsequently written to files, as in the case of an index update.

During `db2text` start, the locking services automatically start. There are the following types of locks on a text index:

**S-lock** For shared read-only access. For example, search requests.

**U-lock**

For read and write access while computing changes to an index (update) with concurrent read access.

**X-lock** For exclusive read/write access for a short period during which changes are actually written to the index.

**IX-lock**

For intended exclusive read/write access preventing any new S-locks while the update process is waiting for an X-lock.

There is one Net Search Extender locking service per DB2 instance. The locking service maintains the locks for multiple databases.

The locking services configuration file is `db2ext1m.cfg`. It is stored on `<instance_owner_home>/sql1lib/db2ext` for UNIX systems and on `<sql1lib>\<DB2INSTANCE>\db2ext` for Windows.

Changes to the configuration file only take effect when Net Search Extender Instance Services are started during `db2text` start. The user can set the following values:

- The maximum number of databases
- The maximum number of indexes per database
- The maximum number of allowed locks (concurrent users) per index
- Waiting times and the number of attempts to obtain a lock

The default values of the configuration file are as follows:

```
<default
    maxDbs          = " 8"
    maxIdxPerDb     = " 50"
    maxLocksPerIdx = "100"

    sWait = " 50"
    uWait = " 500"
    xWait = " 500"

    sAttempt = "50"
    uAttempt = "10"
    xAttempt = "60"

    latchTimeout = "80"

/>
```

The syntax is `<default attribute=value.../>` and the attributes have the following meanings:

**maxDbs**

The number of databases the locking services can handle (integer >1).

**maxIdxPerDb**

The number of indexes per database that can be locked (integer >1). This value is the same for all databases.

**maxLocksPerIdx**

The number of locks that can simultaneously exist on an index (integer >1). This value is the same for all indexes.

Shared memory usage is proportional to the product of the above three 'max' values. To avoid excessive shared memory usage, make sure that the values in use match the actual configuration of your DB2 instance. If you increase the values for maxDbs, maxIdxPerDb, or maxLocksPerIdx beyond the default values in the configuration file mentioned above, ensure that you have sufficient memory. Pay particular attention to the values for maxIdxPerDb and maxLocksPerIdx if you are using the partitioned database environment because these settings will be used for every partition. This is especially significant in determining memory requirements if a number of logical partitions are defined for a DB2 instance on a physical machine.

**sWait/sAttempt**

When requesting an S-lock, sAttempt is the number of attempts that are made if the lock is not granted immediately. sWait is the waiting time between these attempts (integer >1). These parameters also apply to IX-locks.

**uWait/uAttempt**

When requesting a U-lock, uAttempt is the number of attempts that are made if the lock is not granted immediately. uWait is the waiting time between these attempts (integer >1).

**xWait/xAttempt**

When requesting an X-lock, xAttempt is the number of attempts that are made if the lock is not granted immediately. xWait is the waiting time between these attempts (integer >1).

### **latchTimeout**

This is additional waiting time for interval locking services. To determine the total waiting time for a lock, use the following calculation:

$$\text{waiting time} = \# \text{ attempts} * (\# \text{ waits} + (2 * \# \text{ latchTimeout}))$$

It is strongly recommended to leave the default values for the wait, attempt and timeout parameters intact. The waiting time is calculated in milliseconds. Note that with each attempt, the latchTimeout value is doubled when added to the overall waiting time.

---

## **Viewing a lock snapshot**

You can view a lock snapshot by using one of the following commands:

- For a single text index:

```
db2text CONTROL LIST ALL LOCKS FOR DATABASE mydatabase INDEX myindex
```

- For all the locked text indexes of a database:

```
db2text CONTROL LIST ALL LOCKS FOR DATABASE mydatabase
```

Note that only indexes that are actually locked are in the list.

The first time a text index is locked, memory is reserved for the database and the text index in the locking services. If further text indexes are locked, memory is also allocated for these indexes in the locking services. This memory is only freed again when the text index is dropped or the database disabled, or whenever the Net Search Extender services are restarted. This means that a text index or database consumes memory in the locking services, even if there are no locks currently set.

The command "db2text CONTROL CLEAR ALL LOCKS" forces the release of all the locks on a database or index. See "CONTROL command" on page 185 for details on how to use this command. Always use the index specifier when using the CLEAR ALL LOCKS command. Only use this command if you have checked very carefully that no index update is active on the index where you want to clear the locks. Clearing locks on an index that is currently updated can result in index corruption, and require a complete rebuild of the index. Note that this command does not free any memory allocated for the database or indexes. To free memory, you must either drop the index or disable the database, or restart the Net Search Extender services. Do not release locks during an active index update process.



---

## Chapter 24. Update services

Table changes and index updates are not synchronous. The index update process can be started manually or can be scheduled to begin automatically at given intervals. The update services provide this functionality and are started during `db2text start`.

During index creation, you can specify how often the update services check if an update of the index is required by using the following command:

```
db2text create index DB2EXT.TITLE for text on DB2EXT.TEXTTAB (TITLE)
        UPDATE FREQUENCY D(1,3) H(0,12) M(0) update minimum 5
```

In this example, every Monday and Wednesday at 12 p.m. and 12 a.m. the Update Services wake up and check if there is some work to be completed on index `db2ext.title`. Note that in this example there need to be at least five changes to `DB2EXT.TITLE` before the automatic index update will start to synchronize the text index with the database.

In a partitioned database environment, separate update services are started for each of the nodes. When multiple text indexes are updated in parallel, this can lead to considerable burst disk usage during the I/O intensive phases of the index update. This requires careful consideration when scheduling and executing index updates.

### Note:

Setting index update processes at very short intervals affects system performance negatively. You must consider the amount of changes that you expect to be processed during each update and the time that this will take, as well as the number of indexes that you want to process during automatic index update. Ensure that the intervals between each index update are large enough to allow for an update to finish before the next scheduled update begins, and that updates on several indexes are not scheduled to begin at the same time.

If the text-maintained staging infrastructure is configured for a text index, ensure that the staging table is not in a pending mode by executing the `RESET PENDING` command.



## Chapter 25. Using the DB2 control center

Use the DB2 Control Center to manage Net Search Extender administration functions, DB2 instances, databases and database objects such as tables, views, and user groups.

**Important:** The Control Center and its associated components have been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “Control Center tools and DB2 administration server (DAS) have been deprecated” topic in the *What’s New for DB2 Version 9.7* book.

You can invoke the commands on different DB2 Control Center objects, for example:

- Instance objects
- Database objects
- Index objects

The main elements of the DB2 Control Center are the menu bar, the tool bar, the object tree, and the content pane.

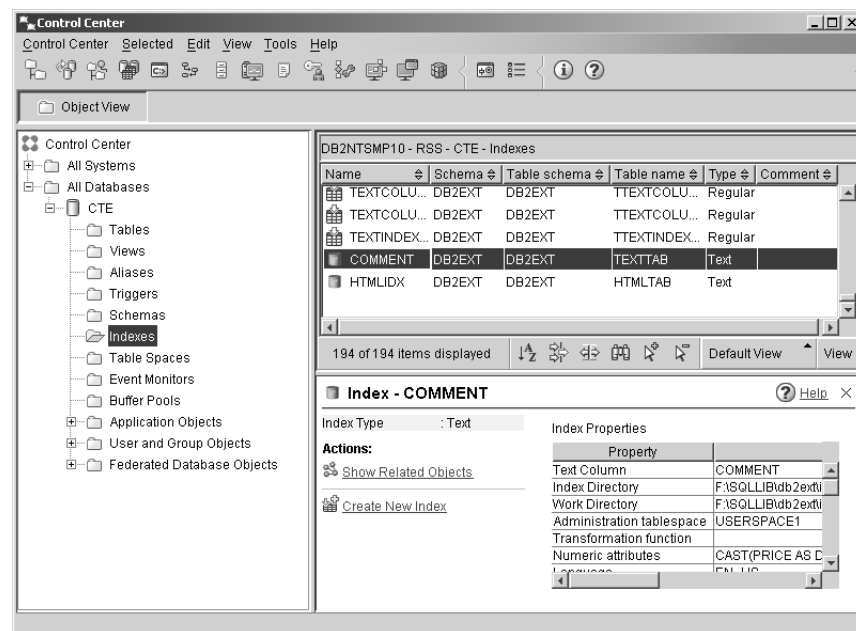


Figure 7. DB2 Control Center

Alternatively, you can use the command line. For more information, see the following sections:

- Part 6, “Developing: creating and maintaining a text index,” on page 89
- Chapter 21, “Net Search Extender instance services,” on page 61

**Note:** Only Net Search Extender indexing and administration functions are found in this chapter. For information on using the DB2 Control Center for other tasks, refer to the DB2 documentation.

---

## Starting and stopping Net Search Extender Instance Services using the DB2 Control Center

From the object tree, click on a system to display the available instances. Highlight the instance and right-click to display the instance object pop-up menu. Highlight **Net Search Extender** and select one of the following commands from the pop-up menu:

### Start Net Search Extender Instance Services

This starts the instance services if they are not already started.

### Stop Net Search Extender Instance Services




Figure 8. Stop Net Search Extender Services dialog

This displays a dialog. Use the check box to stop the instance services and index processes. In the command line syntax, this is known as the **FORCE** option. Click on the **OK** button.

### Instance status

This displays a dialog showing the status of the instance.

---

## Enabling and disabling a database using the DB2 control center

Database administration from DB2 Control Center

In the object tree, click on the instance object to display the available databases. Highlight the database and right-click to display the pop-up menu. Highlight **Net Search Extender** and select one of the following commands from the extended menu:

### Enable the database for text

This displays a dialog if the database is not enabled. Click on the **OK** button to enable the database. If the database is enabled, a message box displays.

### Disable the database for text

This displays a dialog if the database is not disabled. Click on the **OK** button to disable the database. If the database is already disabled, a message box displays.




Figure 9. Disable database for text dialog



Click on the check box if you want to disable the database and drop all the text indexes.

Note that in all the dialogs, the **Show Command** button displays the command line version of the command.

## Text index administration using the DB2 control center

In the object tree, below the database object you can see the index object. Click on the index object to view the indexes in the content pane. You can identify the text index objects by the type 'text' in the Type column.




Figure 10. DB2 Control Center

Right-click on the index object and select one of the following commands from the pop-up menu:

**Create** This displays a dialog. By selecting **Text Index** in the dialog, you start a wizard for creating a text index.

**Filter** This displays a dialog where you can select which index objects display in the control pane view.

**Refresh**

This refreshes information in the object tree and control pane.

**Note:**

To access the instance, database and text index object commands without using the right-click option, click on the **Selected** menu command and highlight **Net Search Extender** to access the relevant commands.

Before creating a text index, ensure that you have considered the prerequisites found in Part 4, "Planning considerations," on page 41.

Other indexing prerequisites include:

- Starting Net Search Extender Instance services
- Enabling the database

## Creating a text index using the DB2 control center

Select the **Create** command and from the extended menu highlight **Text Index**. A Create Text Index Wizard pops up. Use the wizard to specify the configuration options for the text index in a number of panels.

To move between the panels, enter all the mandatory information and click on the **Next** button until the **Finish** button is enabled. Click on the **Finish** button to create the text index.

### Name panel Purpose

This panel allows you to specify the schema and name for the text index. You can also specify a work and index directory for the text index files. Create the administrative tables for the index on the administration table space.

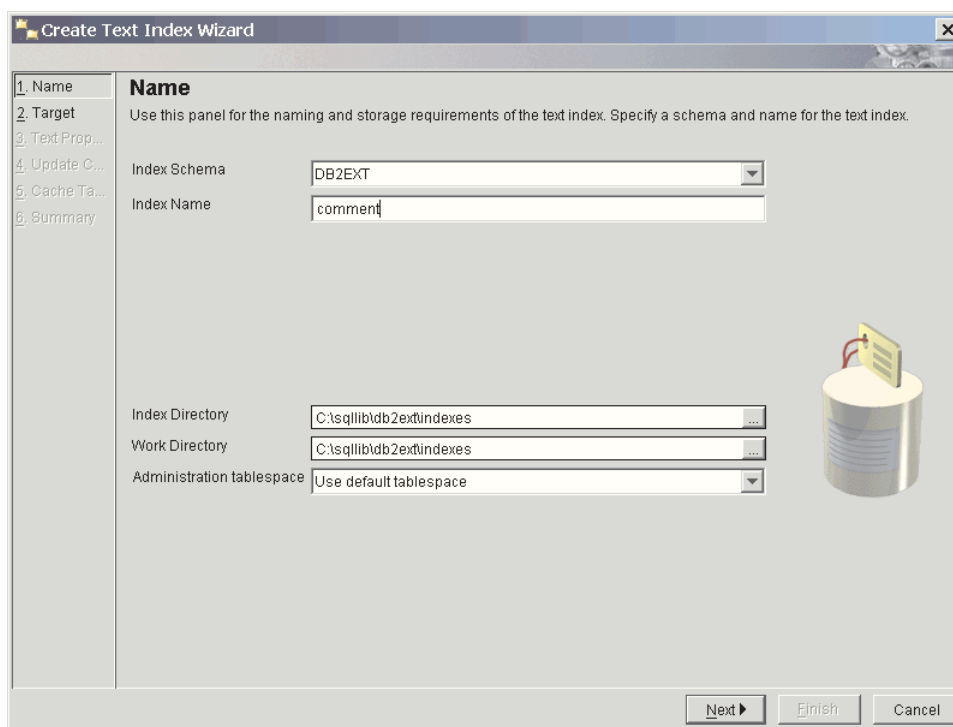


Figure 11. Create Text Index Wizard: Name panel

Here is a description of the fields in the panel:

Table 1. Name panel text fields

Field Name	Mandatory/Optional	Default	Description
Index schema	Mandatory	user ID	Select a schema name of the text index. This is the DB2 schema name for the index-specific administration tables.

Table 1. Name panel text fields (continued)

Field Name	Mandatory/ Optional	Default	Description
Index name	Mandatory	N/A	Enter a valid DB2 index name for the text index. With the index schema, this uniquely identifies a full-text index in the database.
Index directory	Optional	See the path name	Specify the directory path where you will store the text index. The directory must exist with read, write, and execute permissions for the DB2 instance owner user ID.
Work directory	Optional	See the path name	Specify the work directory where you will store temporary files during search and administration operations. The directory must exist with read, write, and execute permissions for the DB2 instance owner user ID.
Administration table space	Optional	Use the default table space	Select a table space name for the text index administration tables. You must define the table space on the same node group as the table space for the user table.

## Target panel

### Purpose

This panel allows you to specify the schema and name of the table or nickname table, and the name of the text column containing the data you want to index. You can use a transformation function to modify the content of the text column. In addition to the text column, you can also specify numeric attributes if you want to add the content of a table column expression to the text index.




Figure 12. Create Text Index Wizard: Target panel

Here is a description of the fields in the panel:

Table 2. Target panel text fields

Field Name	Mandatory/ Optional	Default	Description
Table schema (1)	Mandatory	user ID	Select the schema of the table or nickname table on which you are creating a text index.
Table name (2)	Mandatory	N/A	Select the name of the table or nickname table on which you are creating an index. The table must have a primary key.
Text column (3)	Mandatory	N/A	Select the name of the column used for creating the text index. The column must be one of the following types: CHAR (for bit data), VARCHAR (for bit data), LONG VARCHAR (for bit data), CLOB, DBCLOB, BLOB, GRAPHIC, VARGRAPHIC, LONG VARGRAPHIC, and XML. If this is not the case, the specified transformation function must ensure that it delivers one of the valid data types mentioned here. <b>Note:</b> The LONG VARCHAR and LONG VARGRAPHIC data types are deprecated and might be removed in a future release.
Transformation function	Optional	Disabled	Select to use a transformation function.
Transformation function: Schema	Mandatory (if function selected)	user ID	Select the schema of the UDF used to access the text documents.
Transformation function: Name	As above	N/A	Select the name of a UDF used to access the text documents.

Note that you can only specify the table schema (1), table name (2) and text column (3) in this order.

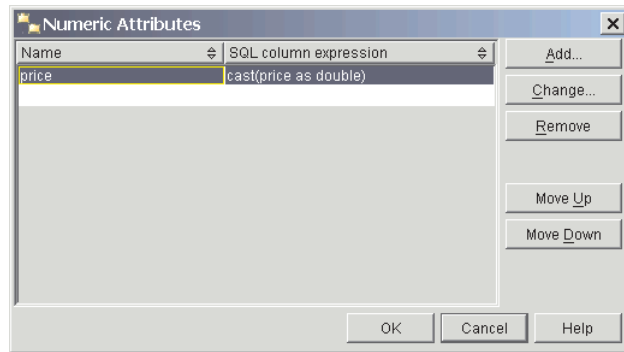


Figure 13. Numeric Attributes dialog

To view or add attributes, click on the **Numeric Attributes** button. A window displays. To add numeric attributes to the index, click on the **Add** button and a further window displays. Specify the SQL column expression and name for the attribute.

Alternatively, select an attribute and press the appropriate buttons to change, move, or remove an entry.

Use Numeric attributes to index numeric column expressions in addition to the text column. For example, if you want to index the column date of type `TIMESTAMP` in addition to the text column, specify a numeric attribute `"cast(julian_day(date) as double)"` and specify a name for the attribute. Numeric attributes must be of data type `DOUBLE`.

Specify a numeric attribute if you want to use a numeric expression inside a search query.

## Text properties panel

### Purpose

This panel allows you to specify the language and format of the text documents. If the documents are not stored in the same `CCSID` as the database, and the text column is of binary type, specify the `CCSID`.

Note that the database `CCSID` is initially selected. If your documents are of a `GPP`, `HTML`, `Outside In`, or `XML` structured format, you can specify a document model.

<p><b>Note:</b> In the format list box, the <code>Outside In</code> filtering format is known as <code>INSO</code>.</p>
---

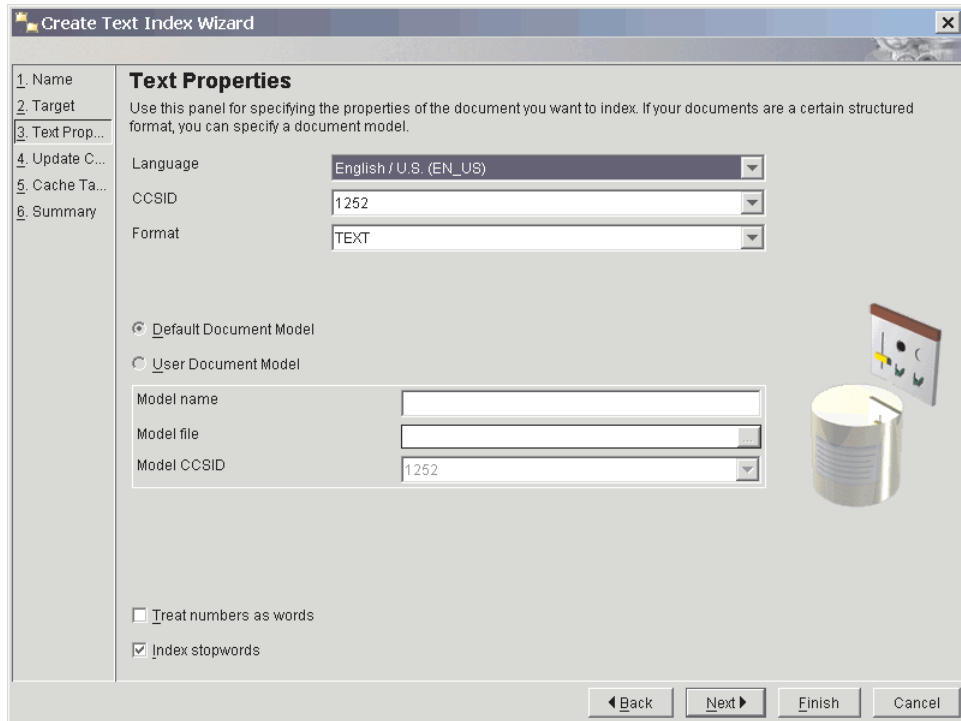


Figure 14. Create Text Index Wizard: Text Properties panel

Here is a description of the fields in the panel:

Table 3. Text Properties panel text fields

Field Name	Mandatory/ Optional	Default	Description
Language	Optional	EN_US	Select a language to determine end-of-sentence and end-of-paragraph delimiters when indexing documents.
CCSID	Optional	CCSID of database	Select the CCSID for indexing text documents.
Format	Optional	TEXT or XML if the column type is XML	Select the text document format: HTML, XML, TEXT, INSO or GPP.
Default Document Model	Optional	Enabled	Use the default document model.
User Document Model	Optional	Disabled	Use your document model.
Model name	Mandatory (if User Document Model selected)	N/A	Enter the name of the document model. For HTML, XML, Outside In, and GPP formats, you can specify a document model. Note that the name is only found in the model file.
Model file	As above	N/A	Specify the document model file. The file must be readable by the DB2 instance owner.
Model CCSID	As above	Database CCSID	Select the CCSID to interpret the contents of the document model file.
Treat numbers as words	Optional	Disabled	Select to interpret sequences of digits as separate words, even if they are adjacent to characters.

Table 3. Text Properties panel text fields (continued)

Field Name	Mandatory/Optional	Default	Description
Index stop words	Optional	Enabled	Select to enable language-specific stop word processing. The <language>.tsw file in the directory sql1lib/db2ext/resources contains the stopword list.

## Update characteristics panel

### Purpose

This panel allows you to specify whether the index updates incrementally or is recreated from scratch. You can specify update settings so that the index automatically updates at the specified time.

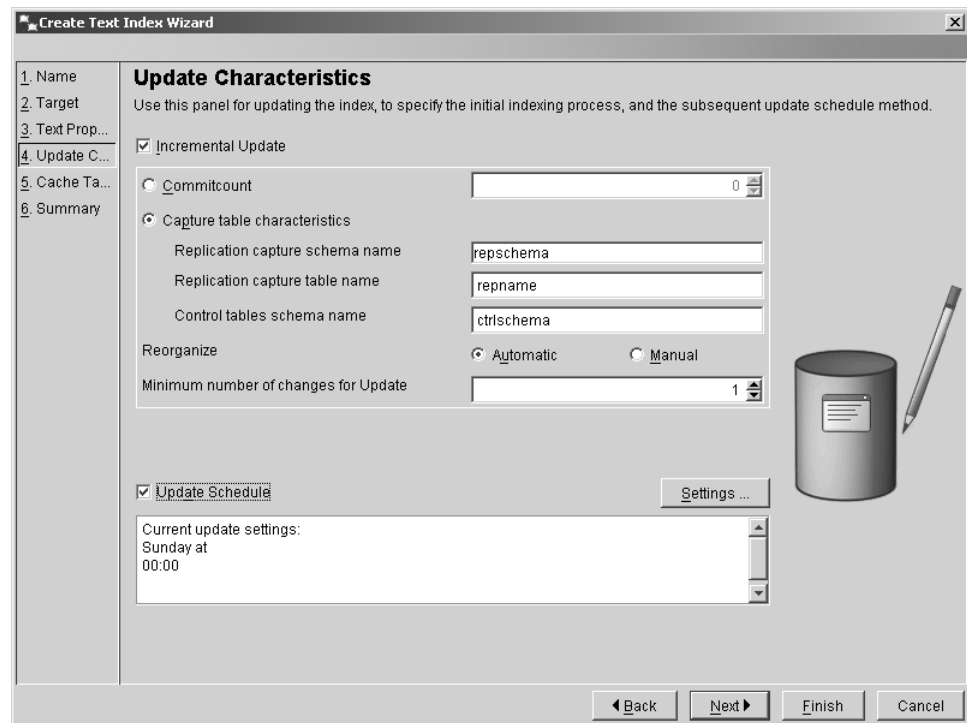


Figure 15. Create Text Index Wizard: Update Characteristics panel

Here is a description of the fields in the panel:

Table 4. Update Characteristics panel text fields

Field Name	Mandatory/Optional	Default	Description
Incremental update	Optional	Enabled	Select for incremental index updates. If you do not enable the check box, you recreate the index when an update operation is performed.
Commitcount	Optional	0	Number of changes processed during an update in one transaction. You are recommended not to change this default.  Using a non-zero Commitcount has implications on performance.

Table 4. Update Characteristics panel text fields (continued)

Field Name	Mandatory/ Optional	Default	Description
Capture table characteristics	Optional	N/A	Select to use a replication capture table for capturing changes on the source table. The replication capture table must either be a Capture Data (CD) table , or a Capture Change Data (CCD) table and replaces the DB2 Net Search Extender generated log table.
Replication capture schema name	Optional	User ID	The schema name of the replication capture table. Note that the table must have been previously created using DB2 Replication.
Replication capture table name	Mandatory, if Capture table characteristics enabled	N/A	The table name of the replication capture table. Note that the table must have been previously created using DB2 Replication.
Control table schema name	Mandatory, if Capture table characteristics enabled	N/A	The control table schema name. Note that the tables must have been previously created using DB2 Replication.
Reorg automatic or manual radio button	Optional or Mandatory	Enabled/ disabled	Completes index reorganization automatically or manually.
Minimum number of changes for Update	Optional	1	Specify the minimum number of changes to the text documents before the index incrementally updates at the specified time.
Update schedule	Optional	Disabled	Select to add automatic update settings.

To add index update settings, click on the **Settings** button. Note that this button is only enabled if you select **Update Schedule**. In the dialog, select the days, hours, and minutes for the update time. Note that if you select multiple days, the update occurs at the same time on all the selected days.

## Cache table panel

### Purpose

This panel allows you to specify a cached table in addition to the index. You can specify the result columns to be cached and you can search the cache using a stored procedure. You can also specify other cache parameters, such as type, maximum size, and the order in which you retrieve the contents of the user table during initial indexing.






Figure 16. Create Text Index Wizard: Cache Table panel

Here is a description of the fields in the panel:

Table 5. Result Cache panel text fields

Field Name	Mandatory/ Optional	Default	Description
Enable table cache	Optional	Disabled	Select to enable the building of a cached table.
Result column table	Mandatory (if Enable table cache selected)	N/A	Displays a list of SQL column expressions specifying the search result columns.
Maximum number of documents	Mandatory	Row count of table	See the following section: Determining cache utilization and cache size.
Average cache row length	Mandatory	N/A	See the following section: Determining cache utilization and cache size.
Percentage of initial cache utilization	Optional	50%	Select the percentage of the cache held free for additional documents.
Maximum cache size	Optional	N/A	Specify a maximum size for the cached table built during index activate. If the number is too small, the activation will fail.
Initial search result order	Optional	Disabled	Select to define the search result order. Documents are returned in the same indexing order as in the cached table. This order can not be ensured after incremental update.
Use persistent cache	Optional	Enabled	This option enables a fast activate execution after a deactivation or system reboot. Note that you must specify a directory path for the persistent cache. Leave disabled if the cache should be temporary.

: Determining cache utilization and cache size

The Percentage of initial cache utilization specifies the percentage of the cache to be held free for additional documents. The Maximum cache size specifies the maximum size of the cached table to be built during activate cache. These options depend on the following factors:

- The actual number of documents in the table.
- The expected number of updates.
- The average size of the SQL expressions you want to cache.

You can enter the recommended values for the Percentage of initial cache utilization and Maximum cache size. Alternatively, you can let them be calculated each time you enter values in the Maximum number of documents or Average cache row length fields.

The Maximum number of documents value is initially set to the row count of the table. Modify this according to the number of documents and expected number of changes. Include all document updates, additions, and deletions.

When you add an SQL expression to the cache table list, the Average cache row length is calculated according to the length of the result. As this is based on the current number of rows in your table, the calculation can take a considerable amount of time. If you know that on average this value is smaller, modify the value.

For example, if your table has 10 entries and the sum of your column expressions is 100, then these values are initially set. If you expect that the maximum number of documents (including deleted ones) is 10 000, enter this figure. If you know that column expressions on average are smaller than the calculated value, such as a VARCHAR(100) and a filled-in text size of 10, use this figure for the average row size.

To define the initial search result order, click on the **Define** button. Note that this button is only enabled if you select the Initial Search Result Order check box. A dialog displays all the specified SQL column expressions. To add a result order, click on the **Add** button and, in the dialog, specify the SQL result order.

To change, move, or remove an entry, select the expression and click on the appropriate buttons.

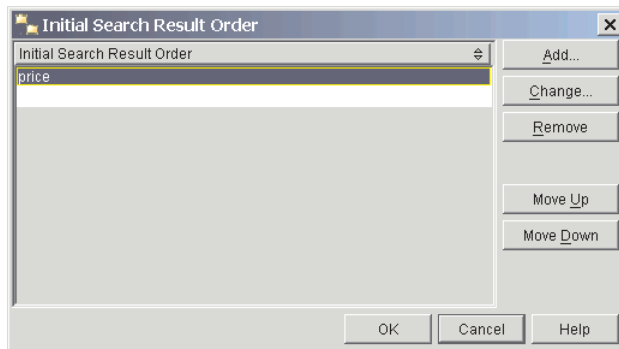


Figure 17. Initial Search Result Order dialog

To add SQL column expressions, click on the **Add** button, next to the Result Column table. In the dialog, specify the result column expression and name.

To change or remove an entry, click on the column expression which enables the appropriate buttons.




Figure 18. Change column expression dialog

## Determining cache utilization and cache size

The Percentage of initial cache utilization specifies the percentage of the cache to be held free for additional documents. The Maximum cache size specifies the maximum size of the cached table to be built during activate cache. These options depend on the following factors:

- The actual number of documents in the table.
- The expected number of updates.
- The average size of the SQL expressions you want to cache.

You can enter the recommended values for the Percentage of initial cache utilization and Maximum cache size. Alternatively, you can let them be calculated each time you enter values in the Maximum number of documents or Average cache row length fields.

The Maximum number of documents value is initially set to the row count of the table. Modify this according to the number of documents and expected number of changes. Include all document updates, additions, and deletions.

When you add an SQL expression to the cache table list, the Average cache row length is calculated according to the length of the result. As this is based on the current number of rows in your table, the calculation can take a considerable amount of time. If you know that on average this value is smaller, modify the value.

For example, if your table has 10 entries and the sum of your column expressions is 100, then these values are initially set. If you expect that the maximum number of documents (including deleted ones) is 10 000, enter this figure. If you know that column expressions on average are smaller than the calculated value, such as a VARCHAR(100) and a filled-in text size of 10, use this figure for the average row size.

## Summary panel

### Purpose

This panel provides an overview of the previously selected parameters.




Figure 19. Create Text Index Wizard: Summary panel

Click on the **Show Command** button to view the commands that are run when you click on the **Finish** button. Clicking **Finish** creates the text index.

---

## Maintaining a text index using the DB2 control center

To maintain the text indexes, select the text index in the content pane and click on the **Select** menu command. You can select one of the following commands from the menu:

1. The ALTER command, for altering a text index.
2. The DROP command, for dropping a text index.
3. The UPDATE command, for updating a text index.
4. The SHOW INDEX EVENTS command, for showing index events.
5. The ACTIVATE INDEX MEMORY command, for activating an index cache.
6. The DEACTIVATE INDEX MEMORY command, for deactivating an index cache.
7. The SHOW STATUS command, for showing the index status.

Note that the Activating and Deactivating commands only display if you create the index with a cache option.

## Altering a text index using the DB2 control center

Select the **Alter** command and a dialog displays a series of panels. These provide an overview of the text index parameters. Note that there are some parameters that you cannot change.

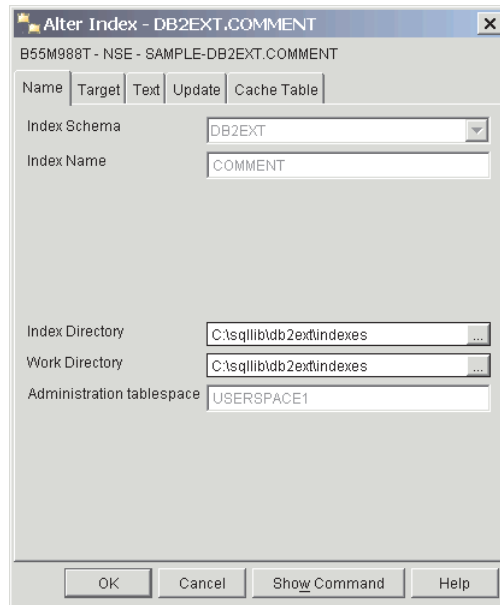


Figure 20. Alter Index dialog: Name tab

The **Name** panel displays the name and storage configurations for the index. You can change the index and work directories.

The **Target** panel displays the target and numeric attribute settings of the index. You cannot change these settings.

The **Text** panel displays the text document configurations. You cannot change these settings.




Figure 21. Alter Index dialog: Update tab

The **Update** panel displays the update characteristics of the index configuration. You can alter the update schedule. If the index was created for incremental

updates, you can also modify the minimum number of changes. If the index was created with the commitcount option, you can also modify the commitcount value.

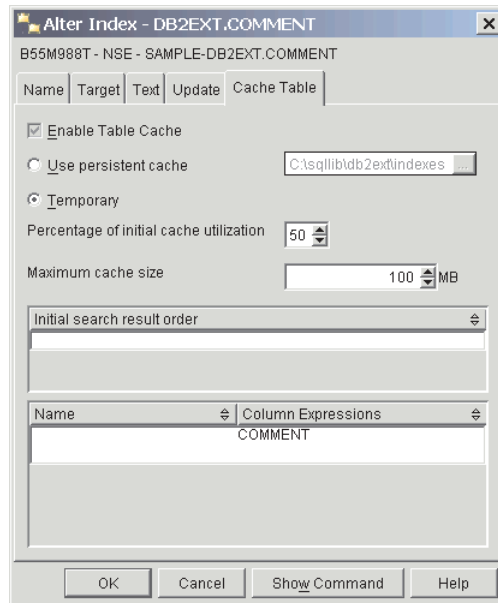


Figure 22. Alter Index dialog: Cache Table tab

The **Cache Table** panel displays the cache option settings. If the result cache is already enabled, you can modify the persistent directory, or make the index cache temporary. You might also change the Maximum cache size and the Percentage of initial cache utilization.

## Dropping a text index using the DB2 control center

Select the **Drop** command and a dialog displays the available text indexes.

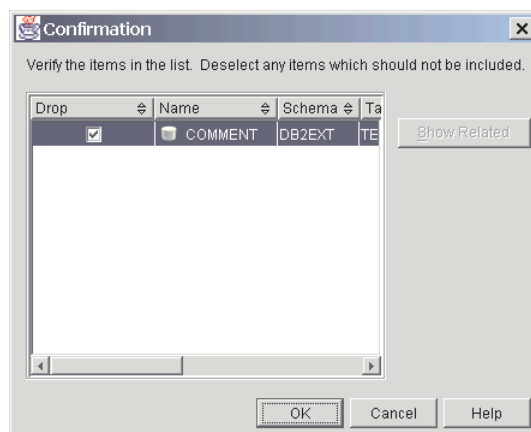


Figure 23. Drop Index dialog

Select the index and click on the **OK** button.

## Updating a text index using the DB2 control center

Select the **Update** command and a dialog displays a number of update options.

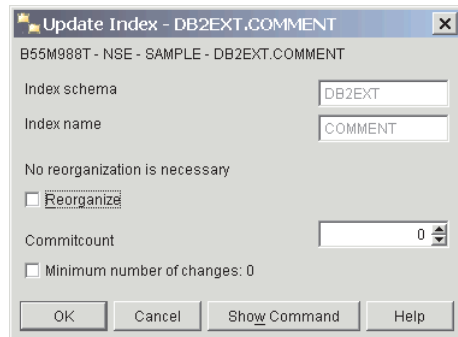


Figure 24. Update Index dialog

You can specify a commitcount for the update operation. If you want to update the minimum specified during create or alter index, select the Minimum number of changes check box. To reorganize the index, select the check box. Note that if reorganization is suggested, the check box is enabled.

## Showing index events using the DB2 control center

Select the **Show index events** command and the contents of the event table display in the dialog. Only the latest 1000 events are listed.

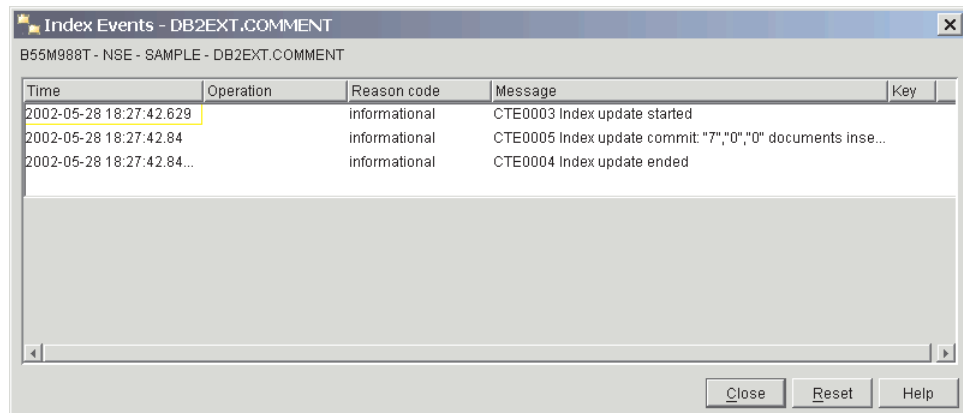


Figure 25. Index Events dialog

To clear the index events, click on the **Reset** button.

## Activating a text index cache using the DB2 control center

Select the **Activate index memory** command and a dialog displays.



Figure 26. Activate the cache dialog

To activate the cache, click on the **OK** button. If you want to build the cache from scratch, select the check box.

## Deactivating a text index cache using the DB2 control center

Select the **Deactivate index memory** command and a dialog displays.



Figure 27. Deactivate the cached table dialog

To release the cache, click on the **OK** button.

## Showing index status using the DB2 control center

Select the **Show status** command and a dialog displays the status of the text index.




Figure 28. Index Status dialog

This includes information on the number of indexed documents, the reorganization suggested flag, and additional index information.



---

## Part 6. Developing: creating and maintaining a text index

This section provides information on creating and maintaining a text index and covers the following areas:

- Introducing the db2text commands
- Enabling a database for text search
- Creating a text index for different data types
- Creating a text index on a nickname with incremental index update using DB2 Replication
- Creating a text index which the stored procedure search can use
- Text indexes on views
- Maintaining an index
- Creating a text index on a range-partitioned table

There is also information on avoiding code page problems which might occur, and performance considerations which you need to take into account.

Before creating a text index, ensure that you have met the prerequisites in Part 4, "Planning considerations," on page 41. Also make sure that you have started the Net Search Extender Instance Services using the db2text start command.

### **Note**

You can also create and maintain a text index by using the DB2 Control Center.



---

## Chapter 26. Enabling a database

**When** Once for each database that contains tables with columns of text to be searched in.

**Command**

ENABLE DATABASE FOR TEXT

**Authorization**

DBADM

This command prepares the database for use by Net Search Extender.

This command also registers Net Search Extender search functions and procedures that are described in Chapter 52, "SQL scalar search function and the SQL table-valued function," on page 239.

When you enable a database, the command additionally creates the following tables and views automatically:

**db2ext.dbdefaults**

Stores the database default values for index, text, and processing characteristics.

**db2ext.textindexformats**

Stores the list of supported formats and the currently active model files used.

**db2ext.indexconfiguration**

Stores the index configuration parameters.

**db2ext.textindexes**

A catalog view that keeps track of all text indexes.

When a database is enabled, it remains enabled until you disable it.



---

## Chapter 27. Disabling a database

**When** When you no longer intend to make text searches on this database

**Command**

```
DISABLE DATABASE FOR TEXT
```

**Authorization**

DBADM on the database

When Net Search Extender prepares the database for use, certain administrative changes are made. This section describes functions that help you to reverse this process.

To disable the connected subsystem, use the following command:

```
db2text DISABLE DATABASE FOR TEXT
```

When you disable a database, the command deletes the following objects and also drops all existing text indexes:

- The Net Search Extender catalog views and tables that were created when the server was enabled.
- The declaration of Net Search Extender's SQL functions (UDFs).

If the `DISABLE DATABASE FOR TEXT` command returns an error, but you want to disable at all costs (even if indexes are still in use) , use the following command:

```
db2text DISABLE DATABASE for text force
```

**Note:** Disabling a database will fail if there are any text indexes defined in the database. It is recommended to remove these indexes one by one and then check if any problems occur. If you use the `disable database for text force` command, it only guarantees that Net Search Extender catalog tables in the database are removed. The `force` option cannot be applied if an existing text index makes use of the text-maintained staging infrastructure.

However, if some of the indexes can not be dropped completely, there may still be resources that need to be manually cleaned up. These include:

- Files in the index, work and cache directory
- Scheduler entries in `ctedem.dat`
- Where an index was created using the replication capture option, the `IBMSNAP_SIGNAL`, `IBMSNAP_PRUNE_SET`, and `IBMSNAP_PRUNCNTL` entries in the tables of the remote database must be deleted manually. These entries can be easily identified using the `APPLY_QUAL='NSEDB2' || <instance name>` and `TARGET_SERVER= <database name>` condition.

In the following example, the instance is `DB2` and the database is `SAMPLE`.

```
DELETE FROM <ccSchema>.IBMSNAP_SIGNAL
WHERE SIGNAL_INPUT_IN IN
      (SELECT MAP_ID FROM <ccSchema>.IBMSNAP_PRUNCNTL
       WHERE APPLY_QUAL= 'NSEDB2' AND TARGET_SERVER= 'SAMPLE');
```

```
DELETE FROM <ccSchema>.IBMSNAP_PRUNCNTL
WHERE APPLY_QUAL= 'NSEDB2' AND TARGET_SERVER= 'SAMPLE';
```

```
DELETE FROM <ccSchema>.IBMSNAP_PRUNE_SET
WHERE APPLY_QUAL= 'NSEDB2' AND TARGET_SERVER= 'SAMPLE';
```



---

## Chapter 28. Creating a text index

**When** Once for each column that contains text to be searched.

**Command**

CREATE INDEX ... FOR TEXT ... (See the following examples)

**Authorization**

One of the following authority levels is required:

- CONTROL privilege on the index table
- INDEX privilege on the table and either IMPLICIT\_SCHEMA authority on database or CREATEIN privilege on index schema
- DBADM authority

You can create a text index on all data types, although there are different requirements for the following data types:

- Binary data types
- Unsupported data types

There are also different requirements for creating a text index for stored procedure search.

When you create a text index, Net Search Extender automatically creates the following objects, depending on whether the extended text-maintained staging infrastructure is enabled for the text index or not.

**Note:** It is mandatory to specify the ADMINISTRATION TABLES IN clause if an index is created on a range partitioned table. See CTE0150E for more information.

**With the regular log infrastructure:**

**A log table**

This keeps track of all changes to rows in the user table. Note that if you select the RECREATE INDEX ON UPDATE option or use replication capture tables, the log table is not created.

**An event table**

This collects information about all updates and potential problems during an update of the text indexes.

**Triggers on the user table (added with initial update)**

These add information to the log table whenever a document in the user table is added, deleted, or changed. The information is necessary for index synchronization during the next scheduled or manual index update.

Note that triggers are only created if you create a log table, and the text index is created on a base table and not on views or nickname tables.

**With the extended log and staging infrastructure:**

**A log table**

This keeps track of updates to the documents.

**An auxiliary staging table**

This keeps track of inserts and deletes.

### **An event table**

This collects information about all updates and potential problems during an update of the text indexes.

### **An update trigger on the user table (added during initial update)**

The update trigger adds the primary key of the affected row to the log table, when a document in the indexed column is updated.

z

To optimize performance and disk space, the CREATE INDEX command has an option to specify a different table space for the tables.

**Note:** If you use the LOAD command to import your documents, the triggers do not fire and incremental indexing of the loaded documents is not possible with the regular infrastructure. In this case, it is preferable to use the DB2 IMPORT command as this activates the triggers.

If the extended text-maintained infrastructure is configured for the text index, documents inserted with a load insert operation are captured in the auxiliary staging table, and incremental indexing is possible.

The following example creates a text index on text column HTMLFILE in table htmltab.

```
db2text create index DB2EXT.HTMLIDX for text on DB2EXT.HTMLTAB
        (HTMLFILE) format HTML
```

A primary key must exist on this table.

The default values for index creation are taken from the db2ext.dbdefaults view.

To reverse the changes made by CREATE INDEX, use the DROP INDEX command. See "Dropping a text index" on page 112 for this information.

To populate the created index with data from the text column, use the following command:

```
db2text update index DB2EXT.HTMLIDX for text
```

Note that you can only search for documents successfully after the text index is synchronized with the table using a db2text update command.

If errors occur during indexing, **index update event** rows are added to the event table. This happens, for example, when a document queued for indexing can not be found or if the document format is invalid. For additional information, see the description of the "Event view" on page 283.

### **Note: Search summary**

Depending on the options selected during index creation, different ways of searching are possible:

- The SQL scalar search functions work on all text indexes, except those created on views.
- The stored procedure search function only works on text indexes that are created with a cache.
- The SQL table-valued function works on all text indexes, including those created on views.



---

## Creating a text index on binary data types

When you store data in a column that has a binary data type, for example BLOB or FOR BIT DATA, DB2 does not convert the data. This means that the documents retain their original code pages (CCSIDs), which can cause problems when creating a text index as you might have two different code pages. Therefore, you need to determine whether you are using the code page of the database, or the code page specified in the CREATE INDEX command.

To avoid this problem, specify the code page when creating the text index:

```
db2text CREATE INDEX db2ext.comment FOR TEXT ON db2ext.texttab (comment)
        CCSID 1252
```

If the code page is not specified, check which CCSID has been used to create the index, by calling:

```
db2 SELECT ccsid FROM db2ext.textindexes WHERE INDSHEMA = 'DB2EXT'
        and INDNAME = 'COMMENT'
```

Note that there is no support for documents with different code pages in one text index. For information on how DB2 converts document code page settings, go to the Globalization Guide.

Note that the problem does not exist when you create indexes on character data types. For character data types, do not specify the CCSID parameter.

---

## Creating a text index on an unsupported data type

To create an index, the text columns must be one of the following data types:

- CHAR
- VARCHAR
- LONG VARCHAR
- CLOB
- GRAPHIC
- VARGRAPHIC
- LONG VARGRAPHIC
- DBCLOB
- BLOB
- XML

If the documents are in a column of a data type that is not contained in this list, such as a user-defined type (UDT), you must provide a conversion function that takes the user type as input and casts it to one of the valid data types as an output type.

Then you must specify the name of this conversion function at index creation time. See “CREATE INDEX command” on page 204 for further information.

Example: You intend to store compressed text in a table.

1. Create a user-defined type (UDT) for the text in an interactive SQL session:  

```
db2 "CREATE DISTINCT TYPE COMPRESSED_TEXT AS CLOB(1M)"
```
2. Create a table and insert the text into it:

```
db2 "CREATE TABLE UDTTABLE (author VARCHAR(50) not null,
                             text COMPRESSED_TEXT, primary key (author))"
db2 "INSERT ..."
```

3. Create a user-defined function (UDF) called, for example, `uncompress`. This receives a value of type `COMPRESSED_TEXT` and returns the corresponding uncompressed text as, for example, a `CLOB(10M)` value.
4. Create your text index in the following way to specify the `uncompress` UDF:

```
db2text "CREATE INDEX UDTINDEX for text ON UDTTABLE
        (uncompress(text))
        ..."
```

---

## Creating a text index on a nickname with incremental index update using DB2 Replication

Before creating a text index on a nickname using a replication capture table, you must perform the following steps:

1. Set up the DB2 federated database with all server definitions and wrapper definitions.
2. Set up the replication control tables and the capture programs at the remote server. This is where the source table for the nickname resides. If DB2 does not automatically create nicknames, you must create nicknames in the federated DB2 database using one schema name for the following tables:

- `IBMSNAP_SIGNAL`
- `IBMSNAP_PRUNE_SET`
- `IBMSNAP_PRUNCNTL`
- `IBMSNAP_REGISTER`
- `IBMSNAP_REG_SYNC` (Non-DB2 remote sources only)

After this step, nicknames for the replication control tables are available as nicknames under one "capture control schema" on the federated DB2 database. This schema name is important for the `DB2TEXT CREATE INDEX` command.

3. Register the table as a replication source.
4. If DB2 does not automatically create a nickname in the registration step, create a nickname for the replication capture table in the federated database. The replication capture table can either be a Change Data (CD) table or a Consistent Change Data (CCD) table. This nickname is a parameter for the `DB2TEXT CREATE INDEX` command.

Note that the column names `IBMSNAP_OPERATION`, `IBMSNAP_COMMITSEQ`, `IBMSNAP_INTENTSEQ`, and the names of the primary key columns must not be changed.

5. If you are using DB2 replication source, ensure that your capture program is running. Do not use a cold start for the capture program: if you do, all rows in the `IBMSNAP_SIGNAL` table for `APPLY_QUAL LIKE 'NSE%'` have to be reinserted. In the following SQL statement you can see how this is done:

```
INSERT INTO <capture control schema>.IBMSNAP_SIGNAL
SELECT CURRENT_TIMESTAMP, 'CMD', 'CAPSTART', MAP_ID, 'P'
FROM <capture control schema>.IBMSNAP_PRUNCNTL
WHERE APPLY_QUAL LIKE 'NSE'
```

6. You can use the following example to create a text index on a nickname using replication:

```
DB2TEXT
CREATE INDEX <indexname> FOR TEXT ON <nickname> (< text column>)
REPLICATION CAPTURE TABLE <capture nickname>
CONTROL TABLE SCHEMA <capture control schema>
```

---

## Creating a text index which the stored procedure search can use

If you know in advance what subset of data from your table you want to present to the user and are only interested in the top ranked search results and not the complete results list, you can use the stored procedure search. For the stored procedure search, you need to specify cache options during the CREATE INDEX command. Working with a cached index, enables high performance at query time, by moving all the specified data into main memory to avoid costly physical read operations from the table.

Before you update the cached index the first time, ensure that your table already contains documents to avoid updating an index on a non-populated table. This provides better indexing performance and a correct estimate of cache memory requirements.

The stored procedure search allows you to quickly return predefined data that is associated with a document. Use the cache table option to define this in the CREATE INDEX command. The ACTIVATE CACHE command then moves the specified data into the memory cache.

When creating a text index for stored procedure search, you must determine and calculate the following parameters:

- The type of cache (temporary or persistent)
- The type of index update (automatic and incremental, or recreation at each update)
- The maximum amount of memory that Net Search Extender can use, by using MAXIMUM\_CACHE\_SIZE.
- The amount of free memory necessary for subsequent document updates, by using PCTFREE. Note that this is only for incremental updates.

The following types of cache are available:

### A temporary cache

This is rebuilt with each DB2TEXT ACTIVATE CACHE command, and requires reloading the data from your DB2 table to memory. Building the cached index from scratch each time Net Search Extender is restarted or the system is rebooted takes longer than reactivating a persistent cache, especially for large tables. Use a temporary cache only if you are dealing with a small amount of fixed data and do not have to consider the time that it takes to build the cached data.

### A persistent cache

This is maintained on disk and can be quickly mapped to memory by using the DB2TEXT ACTIVATE CACHE command. In incremental index update scenarios, the cache must remain activated to allow synchronization between the table and the cached indexed. If this does not occur, the next DB2TEXT ACTIVATE CACHE command recreates the cache from scratch.

The following methods of updating a text index are available:

### Without the RECREATE INDEX ON UPDATE option

If the RECREATE INDEX ON UPDATE option is not set, automatic index update takes place. The process is triggered by the update index command and the update intervals are determined by the update frequency option. The update process is also known as incremental update.

Avoid deleting and reinserting a document in the table because slots for a deleted document cannot be reused in the cache. As a consequence, you should avoid changing key columns on an activated index.

### With the RECREATE INDEX ON UPDATE option

This recreates the index on each update. Use variable data types in the cache column expressions wherever possible. This will save cache space. Use the corresponding cast expressions in the CACHE TABLE clause.

Use this option if your data is not very stable, that is if you expect to insert more than 50% of your documents after the initial index activation.

Net Search Extender provides two SQL functions to help you determine the CREATE INDEX memory parameters. These are MAXIMUM\_CACHE\_SIZE and PCTFREE.

- MAXIMUM\_CACHE\_SIZE specifies the maximum size of the cached index. You can obtain the value of MAXIMUM\_CACHE\_SIZE in megabytes (MB) by using the following UDF function:

```
DB2EXT.MAXIMUM_CACHE_SIZE(maximumNumberDocs INTEGER,  
                           averageRowLength INTEGER, numberOfCacheColumns INTEGER)
```

The following command returns the average row length parameter from your table:

```
SELECT AVG(LENGTH(cache column_1) + ... + LENGTH(cache column_n))
```

Note that the average may change significantly when further documents are inserted into your table. The number of cache columns relates to the number of column expressions used in the CACHE TABLE clause of the DB2TEXT CREATE INDEX command.

For additional information, see Chapter 15, "Stored procedure search memory requirements," on page 47.

- PCTFREE specifies what percentage of the cache that is specified in MAXIMUM\_CACHE\_SIZE to hold free for additional documents. The following UDF function returns the recommended PCTFREE value based on the actual and maximum numbers of documents.

```
DB2EXT.PCTFREE(actualNumberDocs INTEGER, maximumNumberDocs INTEGER)
```

The actual number of documents is the number of rows in your table at the time of the first ACTIVATE CACHE command, which creates the memory cache.

The maximum number of documents is an estimate of the maximum number of documents in your table before the next DB2TEXT ACTIVATE command (for a temporary cache), or DB2TEXT ACTIVATE CACHE RECREATE command (for a persistent cache) is run.

The default is set to 50%. If you are recreating the index on each update, set the PCTFREE value to 0.

Assume that you have 10 000 rows in your table and you do not expect more than 20 000. Use the following call to calculate the PCTFREE value you require:

```
db2 "values DB2EXT.PCTFREE(10000,20000) "
```

Assume that your maximum row size is 20 000 and that you have 2 columns in your cache with an average size of 76. Use the following call to return the size:

```
db2 " values DB2EXT.MAXIMUM_CACHE_SIZE(20000,76,2) "
```

After determining suitable parameters, you can create your cached index by using the following call:

```
db2text CREATE INDEX db2ext.comment FOR TEXT ON db2ext.texttab (comment)
        CACHE TABLE (docid) PCTFREE 10 MAXIMUM CACHE SIZE 5
```

In this example, the docid column is cached, using main memory for fast result table return. Ten percent of the cache memory is reserved for future documents and the cache is limited to a maximum of 5 MB.



---

## Chapter 29. Creating text indexes on views

You can create text indexes on views for use with the stored procedure or table-valued search functions, however, you cannot use any of the scalar functions, for example, CONTAINS.

Another major limitation is that you cannot create triggers on views, so any changes in the underlying base tables are not recognized automatically.

Consequently, for incremental index updates, the user has to know which document has been added, updated, or deleted in order to synchronize the text index with the database. To do this, you must add all the changes to the log table. This process is shown in the following sample:

1. Create a base table using the following command:

```
db2 "create table DB2EXT.TLOGIX140789
(key INTEGER not null PRIMARY KEY,
name VARCHAR(50) not null, comment VARCHAR(90))"
```

2. Add some entries using the following commands:

```
db2 "insert into DB2EXT.TLOGIX140789 values
(1,'Claus','works in room 301')"
db2 "insert into DB2EXT.TLOGIX140789 values
(2,'Manja','is in the same office as Juergen')"
db2 "insert into DB2EXT.TLOGIX140789 values
(2,'Juergen','has the longest way to Raiko')"
db2 "insert into DB2EXT.TLOGIX140789 values
(3,'Raiko','is sitting in the office besides Claus ')"
```

3. Create a view using the following command:

```
db2 "create view sampleview as select key, comment from DB2EXT.TLOGIX140789"
```

4. Use the following commands to create, update, and activate the text index:

```
db2text "create index indexview for text on sampleview(comment)
        cache table (comment) maximum cache size 1 key columns
        for index on view (key)"
db2text "update index indexview for text"
db2text "activate cache for index indexview for text"
```

**Note:** You need to specify the cache table to be able to create a text index on a view. To create the correct log table, you must specify the key columns for the index on a view. If you create an index in this way, you can also search the index using the table-valued function.

When you use the stored procedure search in a partitioned database environment, you must explicitly specify a table space for administration tables on a single partition and explicitly call on this partition. To ensure that you connect to the correct partition, use the DB2NODE environment variable.

5. To update the table, use the following commands:

```
db2 "insert into DB2EXT.TLOGIX140789 values
(4,'Bernhard','is working on the same floor
as Manja, but not as Claus')"
db2 "insert into DB2EXT.TLOGIX140789 values
(5,'Guenter','shares the office with Raiko')"
```

6. Then update the log table. To get the name of the log table, use the following command:

```
db2 "select INDSHEMA,INDNAME,LOGVIEWSCHEMA,LOGVIEWNAME
from db2ext.textindexes"
```

This is the layout of the log table:

sqltype	sqllen	sqlname.data	sqlname.length	
496	INTEGER	4	OPERATION	9
392	TIMESTAMP	26	TIME	4
497	INTEGER	4	PK01	4

To add the entries into the log table, use the following commands:

```
db2 "insert into DB2EXT.TLOGIX140789 values(0,CURRENT TIMESTAMP,4)"
db2 "insert into DB2EXT.TLOGIX140789 values(0,CURRENT TIMESTAMP,5)"
```

The first value describes the operation (0 = insert, 1 = update, 2 = delete). The second should always be the CURRENT TIMESTAMP, and the last value is the primary key of the row that has been inserted, updated or deleted.

7. Use the following command to update the index again:

```
db2text "update index indexview for text"
```

You can now search with the stored procedure on the new values.



---

## Chapter 30. Creating a text index on range partitioned tables

You can create text indexes on range-partitioned tables with and without the extended text-maintained staging infrastructure that supports incremental index updates.

To disable the infrastructure for a text index on a range-partitioned table, specify the CREATE INDEX command with the **AUXLOG** parameter set to OFF, as shown in the following example:

```
db2text create index sampleix for text on sample(comment) administration tables in
mytablespace index configuration(auxlog off) connect to mydb
```

In this case, the primary log table is added, and document changes are recognized through triggers. Note that the ADMINISTRATION TABLES IN clause must be used when creating indexes on range partitioned tables otherwise you will encounter an error.

You cannot use an incremental update to process changes related to attaching or detaching ranges or to process documents that you loaded into an added partition by using the LOAD command with the **INSERT** parameter. You must recreate the text index to synchronize it with the base table.

When the extended text-maintained staging infrastructure is enabled for the text index, document updates are captured through an update trigger into the primary log table, while document inserts and deletes are captured in the auxiliary staging table through integrity processing. This process is shown in the following sample scenarios:

### Scenario 1: Attaching a partition for a table with extended text-maintained staging infrastructure

```
db2 "create table uc_007_customer_archive (pk integer not null primary key,
customer varchar(128) not null, year integer not null,
address blob(1M) not null) partition
by range(year)(starting(2000)ending(2001)every 1)"
db2text "create index uc_007_idx for text on uc_007_customer_archive (address)
administration tables in mytablespace"
db2 "select indexname, logviewname, auxstagingname from db2ext.textindexes"
db2text "update index uc_007_idx for text"
db2 "create table uc_007_customer_2001 (pk integer not null primary key,
customer varchar(128) not null, year integer not null, address blob(1M) not null)"
db2 "import from uc_007_2001.del of del lobs from ./data modified by codepage=1208
insert into uc_007_customer_2001"
db2 "alter table uc_007_customer_archive attach
partition p2001 starting(2001) ending(2002)
exclusive from uc_007_customer_2001"
```

Note that the changes are not visible yet, and an integrity processing is required.

```
db2 "select * from sysibmts.sysauxlog_ix253720"
PK          GLOBALTRANSID          GLOBALTRANSTIME          OPERATIONTYPE
-----
0 record(s) selected.
```

```
db2 "set integrity for uc_007_customer_archive immediate checked"
```

Integrity processing will place dependent tables in a pending mode.

```
db2 "select * from sysibmts.sysauxlog_ix253720"
```

PK	GLOBALTRANSID	GLOBALTRANSTIME	OPERATIONTYPE
-----			
SQL0668N Operation not allowed for reason code "1" on table "SYSIBMTS"."SYSTSAUXLOG_IX253720". SQLSTATE=57016			

Perform integrity processing for the text-maintained staging table(s). The command will process all text indexes for the table

```
db2text "reset pending for table uc_007_customer_archive for text"
```

```
db2 "select * from sysibmts.sysauxlog_ix253720"
```

PK	GLOBALTRANSID	GLOBALTRANSTIME	OPERATIONTYPE
-----			
1	x'000000000002215B'	x'20081020204612500381000000'	1
2	x'000000000002215B'	x'20081020204612500602000000'	1
3	x'000000000002215B'	x'20081020204612500734000000'	1
5	x'000000000002215B'	x'20081020204612500864000000'	1

The incremental update will process the data from the newly attached partition

```
db2text "update index uc_007_idx for text"
```

**Scenario 2:** Detaching a partition for a table with extended text-maintained staging infrastructure

```
db2 alter table uc_007_customer_archive detach partition p2005 into t4p2005  
SQL3601W
```

The statement caused one or more tables to automatically be placed in the Set Integrity Pending state. SQLSTATE=01586

```
db2text "reset pending for table uc_007_customer_archive for text"
```

```
db2text "update index uc_007_idx for text"
```

---

## Chapter 31. Performance considerations for indexing

To enhance performance during indexing, consider the following issues:

- Use a VARCHAR data type to store the text documents instead of LONG VARCHAR or CLOB.
- Use separate physical disks to store the text index and the database files.
- Use small primary key columns, such as TIMESTAMP and INTEGER instead of VARCHAR types.
- Ensure that your system has enough real memory available for all this data. If there is insufficient memory, the operating system uses paging space instead. This decreases indexing and search performance.
- The update commitcount parameter, used during the automatic or manual updating of the index, slows down the indexing performance during incremental indexing. Note that the parameter is not used during the initial update process.
- Performance may decrease during index update if many error and warning messages are written to the event log table.



---

## Chapter 32. Maintaining text indexes

This section describes how to maintain text indexes and get useful information about their status. The maintenance tasks are:

1. Updating and reorganizing a text index
2. Altering a text index
3. Clearing (deleting) index update event information
4. Dropping a text index
5. Viewing index status

This section also includes information on how to back up and restore indexes and enabled databases.

---

### Updating and reorganizing a text index

After creating and updating the text index for the first time, you must keep the text index up to date. For example, when you add a text document to a table, or change an existing document in a table, you must index the document to keep the content of the index synchronized with the content of the table. Likewise, when you delete a text document from a table, its term references must be removed from the index.

If you specify the `RECREATE` option in the `CREATE INDEX` command, the index is rebuilt completely for each update. This option creates no log tables or triggers. Use this option with care if you have large tables because rebuilding the complete index can be costly.

If the text index was created without the `RECREATE INDEX ON UPDATE` option, information about new, changed, or deleted documents is stored through triggers in a log table. If the text index was configured with the extended text-maintained staging infrastructure (`AUXLOG ON`), an update trigger stores information about changed documents in the log table, while information about inserts and deletes is stored in the auxiliary staging table through integrity processing.

Typically you update an index automatically at given intervals. You can change the update frequency for an existing index by using the `ALTER INDEX` command.

You specify the index update frequency in terms of when the update is to be made, and the minimum number of text changes that must be queued in the log table before an index update begins. If there are not enough changes in the log table at the day and time given, the index is not updated.

You should plan periodic indexing carefully; to index large amounts of text documents can be a time- and resource-consuming task. The time it takes depends on many factors. These include the size of the documents, how many text documents have been added or changed since the previous index update, and how powerful the processor is. Two important tips for planning index updates are:

- When working in a partitioned database environment, pay extra attention to the index update schedule to minimize the number of concurrent index updates. This is because for scheduled updates of text indexes in the system, several index updates running concurrently can lead to as many processes per partition as the number of indexes being updated at a given time.

- You should avoid the combination of a large number of indexes and very high automatic update frequencies, as this might result in deadlock situations. For example, 100 indexes with their update frequency set to every 5 minutes, 24 hours a day, and 7 days a week, generates an internal list of  $100 \times 12 \times 24 \times 7 = 201600$  checkpoints in a week that must be administered.

**Note:** On a DB2 table, rollback and deadlock situations might occur in the following cases:

- High update frequencies
- High frequency change transactions
- Long transactions

When a database table is updated, the changes that must be made to the Net Search Extender index are logged in a log table. After these log table entries have been processed, the entries are deleted from the log table. If these delete operations to the log table coincide with updates in the database table that need to be logged, a deadlock situation might result.

If the text-maintained staging table infrastructure is configured for the index, certain database operation on the base table might block access to the auxiliary staging table. Make sure the auxiliary staging table is not in a pending mode before updating the text index.

---

## Updating a text index

The UPDATE INDEX command lets you update an index immediately on request.

**When** When an index must be updated immediately without waiting for periodic indexing to occur.

**Command**

```
UPDATE INDEX
```

**Authorization**

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table on which the index is defined
- DATAACCESS authority

The following command updates the index:

```
db2text UPDATE INDEX comment FOR TEXT
```

This command is useful when you have added several text documents to a database and want to search them immediately.

If you specify AUTOMATIC REORGANIZE during CREATE INDEX, the index will be automatically reorganized when necessary.

If you instead specify MANUAL REORGANIZATION and want to determine if manual reorganization is necessary, query the db2ext.textindexes view by using the following command:

```
db2 "select reorg_suggested from db2ext.textindexes where INDNAME = 'comment'"
```

If you specify MANUAL REORGANIZATION and often update a column, bear in mind that the update process becomes slower. To manually reorganize, use the following command:

---

## Altering a text index

**When** When the update frequency or index and work directories have to be changed.

**Command**

```
ALTER INDEX
```

**Authorization**

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table on which the index is defined
- DBADM authority

Use this command to change the index work directory, the update frequency of an index, or the cache characteristics, principally the MAXIMUM CACHE SIZE or PCTFREE. If you do not specify an update frequency, the current settings are left unchanged. If an index update or search is running, an error message displays. This states that the index is currently locked and no changes can be made.

The following example changes the update frequency for the index.

```
db2text ALTER INDEX comment FOR TEXT
        UPDATE FREQUENCY d(1,2,3,4,5) h(12,15) m(00) UPDATE MINIMUM 100
```

In this example, the index is to be updated at 12:00 or 15:00, on Monday to Friday, if a minimum of 100 text documents are in the queue.

Use the following command to stop the periodic updating of an index:

```
db2text ALTER INDEX comment FOR TEXT
        UPDATE FREQUENCY NONE
```

If you change the index directories using the ALTER INDEX command, the index files are moved from the original index directory to the new location, and the index is locked during this process. For large indexes, and changes across file systems, this can take considerable amounts of time. After the copying process has finished, the index is unlocked and can be used again.

---

## Clearing index events

**When** When you no longer need the messages in an index's event table.

**Command**

```
CLEAR EVENTS FOR INDEX
```

**Authorization**

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table on which the index is defined
- DBADM authority

Information about indexing events, such as the update start and end times, the number of indexed documents, or document errors that occurred during the update, are stored in the index's event table. This can help you determine the cause of the problem. When you no longer need these messages, you can delete them.

The following example deletes messages from the specified text index:  
db2text CLEAR EVENTS FOR INDEX comment FOR TEXT

---

## Dropping a text index

**When** When you no longer intend to make text searches on a text column.

**Command**

```
DROP INDEX FOR TEXT
```

**Authorization**

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table on which the index is defined
- DBADM authority

Example:

```
db2text DROP INDEX comment FOR TEXT
```

When dropping a text index, you also drop the following tables and views:

- The log table and view of the index
- The event table and view of the index
- The log table triggers (if present)
- The text-maintained staging table and view (if present)

**Note:** Always drop the indexes on the table before dropping the table. If you drop the table first, any text-maintained staging table (if present) is dropped as well, but the indexes with their administrative tables and views still exist.

---

## Viewing text index status

To get information about the current text indexes within the database, use the Net Search Extender catalog views. For example, if you want to know about the current database defaults, use the following command:

```
db2 "select * from db2ext.dbdefaults"
```

For information about the currently available indexes, their corresponding tables, and the number of indexed documents, use this command:

```
db2 "select indschema, indname, tabschema, tabname, number_docs  
from db2ext.textindexes"
```

Use this command for information about the formats of a specific index:

```
db2 "select format, modelname from db2ext.textindexformats where  
indschema = 'DB2EXT' and indname = 'TITLE'"
```

If COMMITCOUNT is not set, then the NUMBER\_DOCS parameter from the db2ext.textindexes is not updated during a running update process. To view the current number of documents updated during the update process, use the following command:

```
db2text CONTROL LIST ALL LOCKS FOR DATABASE sample INDEX db2ext.title
```



---

## Backing up and restoring indexes

Use the following steps to **back up** enabled databases and text indexes created by Net Search Extender:

1. To find out which indexes Net Search Extender has created and where they are stored, call a select statement on the `db2ext.textindexes` view:  

```
db2 "select indschema, indname, indexedirectory from db2ext.textindexes"
```
2. Ensure that no index update is running, and then stop Net Search Extender services with the following command:  

```
db2text stop
```
3. After backing up the database, back up the index directories and subdirectories.
4. Restart Net Search Extender services with the following command:  

```
db2text start
```

Use the following steps to **restore** the enabled databases and text indexes created by Net Search Extender:

1. Stop Net Search Extender with the following command:  

```
db2text stop
```
2. Restore the backup copies of the index directories to the same path as before.
3. Restart Net Search Extender with the command:  

```
db2text start
```

---

## Removing files from the /tmp directory

While the Net Search Extender services are running, make sure that no scheduled jobs that clean `/tmp` are accidentally removing these files. The following files must exist in the `/tmp` directory and must not be deleted while the Net Search Extender services are running:

- Semaphore and shared memory files:

```
<instance_owner>.TEXT.0000.LATCH  
<instance_owner>.TEXT.0000  
<instance_owner>.CACHE.0000  
<instance_owner>.SCHEDULER.LATCH  
<instance_owner>.DEMON.SEM  
<instance_owner>.DEMON.MEM
```

**Note:** In a partitioned database environment, there would be additional, similarly named files corresponding to each node:

```
<instance_owner>.TEXT.0001.LATCH, <instance_owner>.TEXT.0001,  
<instance_owner>.CACHE.0001, and so on.
```

- During create index, if the cache is temporary, you might see files in `/tmp` similar to the following:

```
<database_name>.IX123456  
<database_name>.IX123456.data0
```



---

## Part 7. Methods for searching text

Net Search Extender provides the following methods for searching text:

### SQL scalar search functions

Text search sub-queries can be embedded in SQL queries. Net Search Extender provides the SQL scalar search functions as extension to the available SQL functions. By including text search sub-queries in SQL queries, it is possible to combine Net Search Extender search functionality with DB2 XQuery processing. Text search queries on XML documents can be used in the `db2-fn:sqlquery()` XQuery input function, and allow for a direct processing of the resulting XML documents with XQuery.

### A stored procedure search function

This enables you to return predefined cached result tables.

### An SQL Table-Valued Function

You can use this search in a similar way to the stored procedure search.

For SQL scalar search functions, this section describes the following areas:

- Searching for text, using the `CONTAINS`, `NUMBEROFMATCHES`, and `SCORE` functions.

Refer to Chapter 52, “SQL scalar search function and the SQL table-valued function,” on page 239 for a detailed description of the syntax.

- Specifying search arguments by using examples with the `CONTAINS` function. Refer to Chapter 51, “Syntax of search arguments,” on page 231 for a complete description of the syntax.

For the stored procedure search function, this section describes the following areas:

- Searching for text using the stored procedure search.
- For specifying search arguments, refer to Chapter 51, “Syntax of search arguments,” on page 231 for a description of the parameters.

For the SQL Table-Valued Function, this section describes the following areas:

- Searching for text using the SQL Table-Valued Function and the `HIGHLIGHT` function.

Refer to Chapter 52, “SQL scalar search function and the SQL table-valued function,” on page 239 for a description of the syntax.

- For specifying search arguments, refer to Chapter 51, “Syntax of search arguments,” on page 231 for a description of the parameters.

There is also information on search performance considerations that you may need to take into account.

Before searching, ensure that all the appropriate indexing steps, described in Part 6, “Developing: creating and maintaining a text index,” on page 89, involving the different data types have been performed.



---

## Chapter 33. Searching for text using SQL scalar search functions

Using examples, this section describes how to use the SQL scalar search functions in the following ways:

- Using the function CONTAINS to issue a query.
- Using the function NUMBEROFMATCHES to determine how many matches of the search term are found in a text document.
- Using the function SCORE to obtain the relevancy of a found text document.

Refer to Chapter 52, “SQL scalar search function and the SQL table-valued function,” on page 239 for a description of the syntax.

---

### Issuing a query

This example demonstrates how the CONTAINS function searches for text in column comment in table texttab. The function returns 1 if the text satisfies the search argument, otherwise it returns 0.

```
SELECT AUTHOR,TITLE
       FROM DB2EXT.TEXTTAB
       WHERE CONTAINS(COMMENT, '"book"') = 1
```

In this example, you search for the term book in the column COMMENT.

Searching for "" is not supported. Using two consecutive quotes in a search term will result in a syntax error message. Also, a query syntax error occurs if a newline character is used within the search string.

#### Note:

If you know that the text search alone will return a very large result set, it is beneficial to add restrictive search criteria, for example:

```
SELECT AUTHOR,TITLE
       FROM db2ext.texttab
       WHERE CONTAINS(COMMENT, '"book"') = 1 AND PRICE < 20
```

---

### Searching and returning the number of matches found

Use the NUMBEROFMATCHES function to determine how often the search term is found in each text document.

```
SELECT AUTHOR,TITLE,NUMBEROFMATCHES(COMMENT,'"book"')
       FROM DB2EXT.TEXTTAB WHERE
       NUMBEROFMATCHES(COMMENT, '"book"') > 0
```

NUMBEROFMATCHES returns an integer value for every row.

---

## Searching and returning the score of a found text document

SCORE returns a positive number that indicates how well the document meets the search term relative to other found documents in the same index. The value is calculated based on the number of matches that are found in the document in relation to the document's size. In the following example, you can get the score of a found document by using the SCORE function:

```
WITH TEMPTABLE(docid,score)
  AS (SELECT docid,
          SCORE(COMMENT,'"book"')
      FROM DB2EXT.TEXTTAB)
SELECT *
  FROM TEMPTABLE
 WHERE score > 0
 ORDER BY score ASC
```

SCORE returns a DOUBLE value between 0 and 1.

The values returned by SCORE are only meaningful if they are compared to other SCORE values returned for the same index. The values cannot be compared to scores returned for other indexes.

**Note:** You cannot use the CONTAINS, SCORE, and NUMBEROFMATCHES search functions for indexes created on views.

The SCORE values are different depending on the DB2 environment:

- In a non partitioned database environment, all the documents are in a single table. The SCORE value is based on a single table, and a document's relationship to all the other documents in the table.
- In a partitioned database environment, all the documents are located on different partitions. During indexing, only the documents that are local on every partition are used to build the text indexes. In this case, the SCORE value is based on the documents relationship to all documents in only one of the multiple partition.

---

## Chapter 34. Specifying SQL search arguments

The CONTAINS, NUMBEROFMATCHES, and SCORE functions all use search arguments. This section uses the CONTAINS function to show different examples of search arguments in Net Search Extender functions.

Refer to “Search argument syntax” on page 231 for a complete description of the syntax.

---

### Searching for terms in any sequence

You can have more than one term in a search argument. One way to combine several search terms is to connect them together using commas, like this:

```
SELECT AUTHOR, TITLE
       FROM DB2EXT.TEXTTAB
       WHERE CONTAINS(COMMENT,
                      '("kid", "dinosaur")') = 1
```

This form of search argument finds text that contains any of the search terms in any order. In logical terms, an implicit OR operator connects the search terms.

---

### Searching with the Boolean operators AND and OR

You can combine search terms with other search terms using the Boolean operators “&” (AND) and “|” (OR):

```
SELECT AUTHOR, TITLE
       FROM DB2EXT.TEXTTAB
       WHERE CONTAINS(COMMENT,
                      '"author" | "pulitzer"') = 1
```

You can also combine several terms by using Boolean operators:

```
SELECT AUTHOR, TITLE
       FROM DB2EXT.TEXTTAB
       WHERE CONTAINS(COMMENT,
                      '"author" | "pulitzer" & "book"') = 1
```

If you use more than one Boolean operator, these are evaluated from left to right. However, as in regular Boolean logic, the logical AND operator (&) binds stronger than the logical OR operator (|). You can see this evaluation in the following example, which does not include parentheses:

```
"book" & "pulitzer" | "year" & "author"
```

Net Search Extender evaluates the boolean operators in the following way:

```
("book" & "pulitzer") | ("year" & "author")
```

If you want to enforce a different evaluation order of the boolean operators, you must include parentheses:

```
"book" & ("pulitzer" | "year") & "author"
```

You can also combine Boolean operators with search terms that are chained together using the comma separator:

```
("author", "pulitzer") & "book"
```

In this case, the comma is interpreted as a Boolean OR operator:  
("author" | "pulitzer") & "book"

---

## Searching with the Boolean operator NOT

You can use the Boolean operator NOT to exclude particular text documents from the search:

```
SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
                    '("author", "pulitzer") & NOT "book"') = 1
```

This example excludes any text documents containing the term "book" from the search results for "author" or "pulitzer".

---

## Fuzzy search

"Fuzzy" search searches for words that are spelled in a similar way to the search term.

```
SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
                    'fuzzy form of 80 "pullitzer"') = 1
```

In this example, the search could find an occurrence of the misspelled word pulitzer.

The match level, in the example "80", specifies the desired degree of accuracy. Use fuzzy search when misspellings are possible in the document. This is often the case when an Optical Character Recognition device, or phonetic input creates the document. Use values between 1 and 100 to show the degree of fuzziness, where 100 is an exact match and anything below 80 is increasingly "fuzzy".

<b>Note:</b> If the fuzzy search does not provide the appropriate degree of accuracy, search for parts of a term using character masking.
---

---

## Searching for parts of a term (character masking)

Masking characters, otherwise known as "wildcard" characters, offer a way to make a search more permissive. They increase the number of text documents that are found by a search.

Net Search Extender uses two masking characters: percent (%) and underscore (\_). Net Search Extender uses these masking characters the same way the DB2 predicate LIKE uses them.

- % represents **any number of arbitrary characters**. Here is an example of % used as a masking character in the middle of a search term:

```
SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT, '"th%er"') = 1
```

This search term finds text documents containing the word "thriller", "throttle", and "thread-splitter".



- `_` represents **one character** in a search term. The following example also finds text documents containing the word “thriller”.

```
SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT, '"th_i11er"') = 1
```

You can use more than one wildcard character in a phrase (more than one word in the phrase can contain a wildcard), however the terms resulting from the wildcard expansion can only be single terms, not multi-word terms. For example, the wildcard expression “th%er” will not match the phrase “the caller”.

Use wildcard characters sparingly as they can increase the size of your result list significantly, thus decreasing performance and returning undesired search results.

Note that you cannot combine Fuzzy search or Thesaurus search with wildcard character search. Also, make sure that the wildcard characters are expandable within the implicit or explicit EXPANSION LIMIT *number* search parameter. For details on EXPANSION LIMIT *number* search parameter see, “Search parameters” on page 234

## Searching for terms that contain a masking character

If you want to search for a term that contains the “%” character or the “\_” character, you must precede the character with a so-called *escape* character. You must identify the escape character in the query using the ESCAPE keyword.

In the following example, the escape character is a “!”:

```
SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
                    '"!100!%" ESCAPE "!"') = 1
```

## Searching for terms in a fixed order

If you search for “primary key”, you will only find the two terms if they are adjacent and occur in the sequence shown:

```
SELECT AUTHOR,TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT, '"primary key"') =1
```

## Searching for terms in the same sentence or paragraph

Net Search Extender has a limited capability to search for terms in the same sentence or paragraph.

Here is an example of a search argument that finds text documents in which the search term “web” occurs in the same sentence as the term “disk”:

```
SELECT AUTHOR,TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
                    '"web" IN SAME SENTENCE AS "disk"') = 1
```

You can also search for several words occurring together. In the next example, two phrases are searched for that occur in the same paragraph:

```

SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
        "linguistic analysis processing" IN SAME PARAGRAPH AS
        "search algorithms") = 1

```

---

## Searching for terms in sections of structured documents

Here is an example of a search argument that finds text documents where the search term "IBM" occurs in the subsection "H2" of structured documents.

```

SELECT CATEGORY, DATE
      FROM DB2EXT.HTMLTAB
      WHERE CONTAINS(HTMLFILE,
        'SECTIONS ("H2") "IBM"') = 1

```

Note that section names are case-sensitive. Ensure that the section name in the model file and in the query are identical.

---

## Thesaurus search

Thesaurus search is a powerful search-term expansion function in Net Search Extender. The additional terms you search for are taken from a thesaurus that you build yourself, so you have direct control over the terms. For example, a thesaurus search for "database", might find terms like "repository" and "DB2" if you decide that these terms are related.

Use this type of search for specific areas of interest in which you make frequent searches and produce significantly more effective search results.

The following examples demonstrate the syntax for using thesaurus expansion.

This example takes the term "product" and expands it, adding all related terms of this term found in the thesaurus "nseamplethes".

```

SELECT CATEGORY, DATE
      FROM DB2EXT.HTMLTAB
      WHERE CONTAINS(HTMLFILE,
        'THESAURUS "nseamplethes"
        EXPAND RELATED
        TERM OF "product"') = 1

```

The next example takes the search term "product". The search then expands with all the *synonyms* of the search term.

```

SELECT CATEGORY, DATE
      FROM DB2EXT.HTMLTAB
      WHERE CONTAINS(HTMLFILE,
        'THESAURUS "nseamplethes"
        EXPAND SYNONYM
        TERM OF "product"') = 1

```

---

## Numeric attribute search

You can search on numeric attributes that are stored in a the text index using the following syntax:

```

SELECT AUTHOR, TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
        'ATTRIBUTE "PRICE" between 9 and 20') = 1

```

---

## Free-text search

“Free-text search” is a search in which you express the search term as free-form text. A phrase or a sentence describes in natural language the subject to be searched for. The sequence of words in a free-text query is not relevant. However, at least one of the query terms in the free-text query must occur in the documents to be searched.

Note that there is no support for the masking of characters or words for search strings in a free-text argument.

For example:

```
SELECT AUTHOR, TITLE, SCORE(COMMENT,  
    'IS ABOUT EN_US "something related to dinosaur"')  
FROM DB2EXT.TEXTTAB  
WHERE CONTAINS(COMMENT,  
    'IS ABOUT EN_US "something related to dinosaur"') = 1
```



---

## Chapter 35. Additional search syntax examples

To become familiar with additional search syntax examples, run the search script located in the `sql1lib/samples/extenders/db2ext/` directory. This contains examples of Net Search Extender search functions that run against the sample table.

Enter the command as follows:

```
db2 -tvf search
```

There is also a sample that shows you how you can query XML data. After you have connected to the database, you can perform searches on the data by issuing `db2 -tvf xmlsearch`

If the table and indexes have not been created, run one of the following:

- On UNIX platforms: `nsesample` in the `<instance_owner_home>/sql1lib/samples/extenders/db2ext` directory.
- On Windows platforms: `nsesample (.bat)` in the `<sql1lib>\samples\extenders\db2ext` directory.
- For XML searches, call `xmlsample (.bat)<database>` to populate the database, and create and update the indexes.



---

## Chapter 36. Searching for text using a stored procedure search

Use the stored procedure search interface if you only need a ranked subset of the text search results, and high query performance. Do not use the stored procedure if you require all the search results, or you need to index a large number of documents. The main reason for this is that parts of the user table are copied into memory and so a lot of real memory needs to be available.

You can use the stored procedure to first request results from 0 to 20, then 21 to 40, and so on, in a similar way to cursor navigation. Combining this cursor capability with the use of a cache (calculated during indexing), makes searching extremely fast, especially as no join with the user table is necessary.

If you are going to use the stored procedure, ensure that you consider the following:

- The cache-search-result options have been specified during CREATE INDEX.
- The present and future shared memory requirements, possibly involving incremental updates, have been fully considered.
- The cache of the index has been activated using the `db2text activate` command.
- In a partitioned database environment, to be able to use the stored procedure search, the table must use a table space on a single partition and the procedure should be called on the same partition. Otherwise the search will be disallowed and errors will be returned.

The following is an example of a stored procedure search:

```
db2 "call db2ext.textSearch('\book\','DB2EXT','COMMENT',0,2,1,1,?,?)"
```

The first parameter is the search term. The syntax for the search term is the same as in the SQL scalar functions. The next parameters are the index schema and index name. If you have not enclosed the name in quotes, it is translated to uppercase. The following two numeric arguments give you the result slice starting point and the number of results in the slice. The next two integer values specify if score and hit information are requested. The final two values are the function return values.

**Note:** If you request larger result sets, you need a user table space. If there is none available, create a table space. The following example creates a tablespace on a UNIX platform:

```
db2 "create user temporary tablespace tempts managed by system
      using ('/work/tempts.ts')"
```





---

## Chapter 37. Searching for text using an SQL Table-Valued Function

Use the SQL Table-Valued Function if you do not need all of the search results, and if you do not have enough memory to use a cached index as used in the stored procedure search.

There are two SQL table-valued functions available, both called `db2ext.textsearch`. One of them has two additional parameters for use with the `db2ext.highlight` function.

The SQL Table-Valued Function gives you the same cursor interface as the stored procedure to access only parts of the result. However, you still need to join the results with the user table. You can see this in following example:

```
db2 "select docid , author, score from TABLE(db2ext.textsearch('\book\" ,
      'DB2EXT','COMMENT',3,2,cast(NULL as integer))) as t, db2ext.texttab u
      where u.docid = t.primkey"
```

The following are the values you could return from the SQL Table-Valued Function:

```
--> primKey <single primary key type>
the primary key

--> score          DOUBLE
the score value of the found document

--> NbResults      INTEGER
the total number of found results (same value for all rows)

--> numberOfMatches INTEGER
the number of hits in the document
```

### Note:

- Only a single primary key column is allowed.
- In a partitioned database environment, to be able to search using the SQL Table-Valued Function, the table must use a table space on a single partition and the function should be called on the same partition. Otherwise the search will be disallowed and errors will be returned.

---

## Using the highlight function

To use the SQL table-valued `db2ext.highlight` function, you must use the `db2ext.textsearch` function with the additional `numberOfHits` and `hitInformation` parameters.

Note that the highlight function must not be used if you created an index using a transformation function whose implementation might change, else the returned highlight position information might not be correct because of a mismatch between positional information at the indexing time and at the search time.

In this example, call the `db2ext.highlight` function to display the whole document without highlighting any hits found by the `db2ext.textsearch` function.

```

select p.docid,
       db2ext.highlight(p.comment, t.hitinformation, 'WINDOW_NUMBER = 0')
       as highlight
from DB2EXT.TEXTTAB p,
     table (db2ext.textsearch('"bestseller" | "peacekeeping" | "soldiers"
                             | "attention"', 'DB2EXT', 'COMMENT', 0, 20,
                             cast(NULL as INTEGER), 10)) t
where p.docid = t.primkey and p.docid = 2

```

The query returns the following result:

DOCID HIGHLIGHT

```

2      A New York Times bestseller about peacekeeping soldiers called
      "Keepers" who devise a shocking scheme to get the worlds
      attention after their tour of duty ends.

```

1 record(s) selected.

In this example, call the `db2ext.highlight` function to display the whole document and highlight all hits found by the `db2ext.textsearch` function.

```

select p.docid,
       db2ext.highlight(p.comment, t.hitinformation, 'WINDOW_NUMBER = 0,
                       TAGS = ("<bf>", "</bf>" ) ') as highlight
from DB2EXT.TEXTTAB p,
     table (db2ext.textsearch('"bestseller" | "peacekeeping" | "soldiers"
                             | "attention"', 'DB2EXT', 'COMMENT', 0, 20,
                             cast(NULL as INTEGER), 10)) t
where p.docid = t.primkey and p.docid = 2

```

The search argument returns the following result:

DOCID HIGHLIGHT

```

2      A New York Times <bf>bestseller</bf> about <bf>peacekeeping</bf>
<bf>soldiers</bf> called "Keepers" who devise a shocking scheme to
get the worlds <bf>attention</bf> after their tour of duty ends.

```

1 record(s) selected.

In this example, call the `db2ext.highlight` function to display at maximum 10 parts (windows) of the document. Each window size is 24 characters, which is approximately 12 bytes of data on each side of the hit. In addition, hits found by the table function `db2ext.textsearch` are highlighted.

```

select p.docid,
       db2ext.highlight(p.comment, t.hitinformation, 'WINDOW_NUMBER = 10,
                       WINDOW_SIZE = 24, TAGS = ("<bf>", "</bf>" ) ') as highlight
from DB2EXT.TEXTTAB p,
     table (db2ext.textsearch('"bestseller" | "peacekeeping" | "soldiers"
                             | "attention"', 'DB2EXT', 'COMMENT', 0, 20,
                             cast(NULL as INTEGER), 10)) t
where p.docid = t.primkey and p.docid = 2

```

The search argument returns the following result:

DOCID HIGHLIGHT

```

2      York Times <bf>bestseller</bf> about <bf>peacekeeping</bf> ...
<bf>peacekeeping</bf> <bf>soldiers</bf> called "Keepers" ... the
worlds <bf>attention</bf> after their

```

1 record(s) selected.

The first hit found is `<bf>bestseller</bf>` and this hit determines the first window. The second hit, `<bf>peacekeeping</bf>` is only 8 bytes away from the first

hit and is completely taken into the first window. The third hit, **soldiers** is outside the first window and determines a new window. As the second hit **peacekeeping** is only 2 bytes away from the left side of the **soldiers** hit, it is also taken into the second window and highlighted. The fourth hit **attention** is outside the second window and so determines a new window. As no previous or additional hit is contained in the size of this window, only data surrounding the hit is contained in the window.

Additionally, as no WINDOW\_SEPARATOR is specified, the default window separator, " ... " is taken to separate the three hit windows.

<p><b>Note:</b> To ensure high performance when using the <code>db2ext.highlight</code> function, the user should limit the search results in the <code>db2ext.textsearch</code> table-valued function.</p>
---



---

## Chapter 38. Searching on more than one column

In cases where you need to create a text index on more than one column, the easiest way is to use the SQL scalar function and combine the searches on those columns. You can see this in the following example:

```
SELECT AUTHOR,TITLE
      FROM DB2EXT.TEXTTAB
      WHERE CONTAINS(COMMENT,
                    '"book"')=1 and CONTAINS(AUTHOR,'"Mike"')=1
```

For a table-valued function it is more difficult, as you might need to use the union of the returned tables for performance reasons. Another possibility with the table-valued function is to use a view and combine your table columns in a view column to create a single text index on this view column. In this way, you avoid two separate text search calls.

Combining your text columns may provide an improvement in performance. However, this strongly depends on your individual search requirements.



---

## Chapter 39. Using text search in outer joins

If you use an outer join query that uses the `CONTAINS()` search function, the query might fail and result in the reason code CTE0129 NULL values are not allowed to be passed as parameters unless the `CONTAINS()` predicate references the column of a table on the tuple preserving side of the outer join.

For example, T1 is the tuple-preserving side in 'T1 left outer join T2' and T2 is the tuple-preserving side in 'T1 right outer join T2'.





---

## Chapter 40. Performance considerations during search

To enhance performance during search, consider the following issues:

- When searching within SQL:
  - If you notice a decrease in performance, Use the `explain` statement to check the processing plan of the DB2 Optimizer.
  - Parametric search can make searching faster, especially if you use other search predicates to reduce the result size.
  - Use the result limit keyword if you do not require all of the results.
- When searching with the stored procedure:
  - As the specified cache table expression is copied from the database into memory, ensure that your workstation has enough memory available for this data. If there is insufficient memory, paging space is used, which decreases search performance.
- If you use the `NUMBEROFMATCHES` or the `SCORE` function without the `CONTAINS` function, query performance may decrease. Also, to avoid duplicate processing, ensure that the string in the `CONTAINS` function exactly matches the string used in `NUMBEROFMATCHES` or `SCORE` function.



---

## Chapter 41. User scenarios

Use this chapter to learn about Net Search Extender by using the following walk-through examples:

### The SQL scalar search example

This command line example demonstrates the indexing and search functions available.

### The stored procedure example

This command line example uses the index command from the example above. With the addition of a cache however, the example demonstrates the different indexing and search functions available for stored procedure search.

### The SQL table-valued function example

The SQL table-valued function example is a variant of the stored procedure search example.

<b>Note:</b> Before using the examples, ensure that Net Search Extender is installed successfully by using the installation verification procedure.
---

---

## Simple example with the SQL scalar search function

Use the following steps in the DB2 Net Search Extender example:

1. Creating a database
2. Enabling a database for text search
3. Creating a table
4. Creating a full-text index
5. Loading the sample data
6. Synchronizing the text index
7. Searching with the text index

You can issue the sample commands on the command line of the operating system by using an existing database. For the following examples, the database name is `sample`.

### Creating a database

You can create a database in DB2 by using the following command:

```
db2 create database sample
```

### Enabling a database for text search

You can issue DB2 Net Search Extender commands in the same way as DB2 commands on the command line of the operating system. For example, use the following command to start Net Search Extender Instance Services:

```
db2text START
```

Then prepare the database for use with DB2 Net Search Extender:

```
db2text ENABLE DATABASE FOR TEXT CONNECT TO sample
```

You need to do this step only once for each database.

### Creating a table

```
db2 "CREATE TABLE books (isbn VARCHAR(18) not null PRIMARY KEY,  
author VARCHAR(30), story CLOB(100k), year INTEGER)"
```

This DB2 command creates a table called books. It contains columns for the author, story, isbn number, and the year the book was published. Note that the table must have a primary key.

### Creating a full-text index

```
db2text "CREATE INDEX db2ext.myTextIndex FOR TEXT ON books (story)  
CONNECT TO sample"
```

This command creates a full-text index for the column story. The name of the text index is db2ext.myTextIndex

### Loading sample data

```
db2 "INSERT INTO books VALUES ('0-13-086755-1','John', 'A man was  
running down the street.',2001)"  
db2 "INSERT INTO books VALUES ('0-13-086755-2','Mike', 'The cat hunts  
some mice.', 2000)"  
db2 "INSERT INTO books VALUES ('0-13-086755-3','Peter', 'Some men  
were standing beside the table.',1999)"
```

These commands load the isbn, author, story, and publishing year for three books into the table.

### Synchronizing the text index

To update the text index with data from the sample table, use the following command:

```
db2text "UPDATE INDEX db2ext.myTextIndex FOR TEXT CONNECT TO sample"
```

### Searching with the text index

To search the text index, use the following CONTAINS scalar search function:

```
db2 "SELECT author, story FROM books WHERE CONTAINS  
(story, '\"cat\"') = 1 AND YEAR >= 2000"
```

**Note:** Depending on the operating system shell you are using, you might need a different escape character in front of the double quotes surrounding the text search phrase. The above example, uses "\" as an escape character.

This query searches for all books containing the term cat where the year value of the book is greater or equal to 2000. The query returns the following result table:

```
AUTHOR Mike  
STORY The cat hunts some mice.
```

Other functions supported include SCORE and NUMBEROFMATCHES. SCORE returns an indicator on how well the search argument describes a found document. NUMBEROFMATCHES returns how many matches of the query terms are found in each resulting document.

---

## Simple example with cache usage and stored procedure search

Use the following steps in the DB2 Net Search Extender stored procedure search example:

1. Creating a text index with cache option.
2. Synchronizing the index and activating the cache.
3. Searching with the TEXTSEARCH Stored Procedure.

**Note:** The stored procedure example assumes that the steps from the previous example are complete and that the database is still enabled.

### Creating a text index with cache option

As the database is already enabled, use the following command to create a full-text index:

```
db2text "CREATE INDEX db2ext.mySTPTextIndex FOR TEXT ON books (story)
        CACHE TABLE (author, story) MAXIMUM CACHE SIZE 1
        CONNECT TO sample"
```

In this example, the full-text index is for the column `story` and it specifies a cache table containing the columns `author` and `story`. The name of the text index is `mySTPTextIndex`.

### Synchronizing the index and activating the cache

To update the index with the data inserted into the table, use the following command:

```
db2text "UPDATE INDEX db2ext.mySTPTextIndex FOR TEXT CONNECT TO sample"
```

To activate the cache, use the following command:

```
db2text "ACTIVATE CACHE FOR INDEX db2ext.mySTPTextIndex FOR TEXT
        CONNECT TO sample"
```

This loads the content of the columns `author` and `story` into the cache.

### Searching with the TEXTSEARCH Stored Procedure

You can only use the DB2 Net Search Extender stored procedure in certain cases.

```
db2 "call db2ext.textSearch
    ('\"cat\"', 'DB2EXT', 'MYSTPTEXTINDEX', 0, 2, 0, 0, '?', '?")"
```

This query searches for all books about a cat, but only returns the first two results. The result table for a book might be as follows:

Value of output parameters

```
-----
Parameter Name : SEARCHTERMCOUNTS
Parameter Value : 1
Parameter Name : TOTALNUMBEROFRESULTS
Parameter Value : 1
```

```
AUTHOR    STORY
Mike      The cat hunts some mice.
```

Return Status = 0

For more samples about the search syntax, check the following file in the DB2 instance directory: `sqllib/samples/extenders/db2ext/search`

---

## Simple example with the SQL table-valued function

You can use the SQL table-valued function on the text indexes created in the previous examples.

The SQL table-valued function query corresponds to the previously used `CONTAINS` query. See "Synchronizing the text index" under the topic "Simple example with the SQL scalar search function" on page 139 for information.

```
db2 "SELECT author, story FROM books b, table (db2ext.textsearch
      ('\"cat\"','DB2EXT','MYTEXTINDEX', 0, 2, CAST
      (NULL AS VARCHAR(18)))) T where T.primKey = b.isbn
```

In the above example, NULL is cast to the data type of the primary key.

---

## Chapter 42. Using a thesaurus to expand search terms

You can broaden a query by searching not only for a specific search term, but also for terms that are related to it. You can automate this process by using Net Search Extender's functions for looking up and extracting the related search terms from a thesaurus. A thesaurus is a controlled vocabulary of semantically related terms that usually covers a specific subject area.

Net Search Extender lets you expand a search term by adding additional terms from a thesaurus that you have previously created. Refer to Chapter 51, "Syntax of search arguments," on page 231 to find out how to use thesaurus expansion in a query.

To create a thesaurus for using it in a search application requires a thesaurus definition file that has to be compiled into an internal format, the thesaurus dictionary.

This section describes:

- **"The structure of a thesaurus"**

A thesaurus is structured like a network of nodes linked together by relations. This section describes Net Search Extender's predefined relations and how to define your own relations.

- **"Creating and compiling a thesaurus" on page 145**

This is a description of the syntax of a thesaurus definition file, and of the tools that you use to compile it into a thesaurus dictionary.

---

### The structure of a thesaurus

A thesaurus is structured like a network of nodes linked together by relations. Net Search Extender looks up a term in a thesaurus by starting at the term, then following a path through the term relations and delivering the terms found in the process.




Figure 29. An example of the structure of a thesaurus

Thesaurus entries are connected by relations. Relation names, such as `BROADER`, let you restrict an expansion to certain named lines in the relation hierarchy. Some relations are bidirectional, others are unidirectional; `BROADER`, for example, is the name of a unidirectional relation.

## Predefined thesaurus relations

These are the relations that are predefined in the Net Search Extender:

- **Associative relations**

An associative relation is a bidirectional relation between two terms that do not express the same concept but relate to each other.

Predefined associative relation: `RELATED_TO`

Examples:

```
tennis RELATED_TO racket
football RELATED_TO goal (sports)
```

- **Synonym relations**

A synonym relation is a bidirectional relation between two terms that have the same or similar meaning and can be used as alternatives for each other. This relation can, for example, be used between a term and its abbreviation.

Predefined synonym relation: `SYNONYM_OF`

Examples:

```
spot SYNONYM_OF stain
US SYNONYM_OF United States
```

The figure in Figure 29 shows two goal terms in the same thesaurus. One is specified with the comment (sports), the other with the comment (abstract). Even if terms have the same spelling, synonym relations can connect different word groups. You can model this by using different relations when defining the thesaurus.

- **Hierarchical relations**



A hierarchical relation is a unidirectional relation between two terms, one of which has a broader (more global) meaning than the other. Depending on its direction, the relation can be used to look up either more specialized or more global terms.

Predefined hierarchical relations:

- LOWER\_THAN to model narrowing relations

LOWER\_THAN relations are for modelling a sequence of more specialized terms. The deeper you follow a narrowing relation, the more specific the terms become. For example, if you look up the term ball game along a LOWER\_THAN relation, the result could be squash tennis and so on, in a list of increasingly specialized terms.

- HIGHER\_THAN to model broadening relations

HIGHER\_THAN relations are for modelling a sequence of more and more global terms. The deeper you follow such a relation, the less specific the terms become. For example, if you look up the term ball game along a HIGHER\_THAN relation, the result could be game and so on, in a list of increasingly global terms.

## Defining your own relations

Net Search Extender lets you define your own RELATED\_TO, LOWER\_THAN, and HIGHER\_THAN thesaurus relations. Because each relation name must be unique, you must qualify such relations names by the addition of a unique number, like this: RELATED\_TO(42).

You can use the same relation number to define a relationship of a different type, such as LOWER\_THAN(42). The number 0 is used to refer to Net Search Extender's predefined relations.

---

## Creating and compiling a thesaurus

Use the following steps to create a thesaurus that can be used by the Net Search Extender functions:

1. Create a thesaurus definition file.
2. Compile the definition file into a thesaurus dictionary.

## Creating a thesaurus definition file

To create your own thesaurus, your first step is to define its content in a definition file using a text editor.

**Restrictions.** The length of the file name, including the extension, must not exceed 256 characters. You can have several thesauri in the same directory, but it is recommended that you have a separate directory for each thesaurus.

A sample English thesaurus definition file `nsesamplethes.def` is provided. The thesaurus directory for Windows systems is:

```
<sql11ib>\db2ext\thes
```

On UNIX systems, the thesaurus directory is:

```
<instance_owner_home>/sql11ib/db2ext/thes
```

Here are the first few definition groups from that file:

```
:WORDS
  accounting
  .RELATED_TO account checking
  .RELATED_TO sale management
  .SYNONYM_OF account
  .SYNONYM_OF accountant

:WORDS
  acoustics
  .RELATED_TO signal processing

:WORDS
  aeronautical equipment
  .SYNONYM_OF turbocharger
  .SYNONYM_OF undercarriage

:WORDS
  advertising
  .RELATED_TO sale promotion
  .SYNONYM_OF advertisement
:
:
:
```

Figure 30. An extract from the sample thesaurus definition file

For the syntax of each definition group, see “Thesaurus support.”

Each member must be written to a single line. Each associated term must be preceded by the relation name. If the member terms are related to each other, specify a member relation.

The length of the member terms and associated terms is restricted to 64 characters. Single-byte characters and double-byte characters of the same letter are regarded as the same. Uppercase and lowercase letters are not distinct. A term can contain a blank character and either a single-byte character period "." or colon ":" can be used.

The user-defined relations are all based on the *associative* type. They are identified by unique numbers between 1 and 128.

## Compiling a definition file into a thesaurus dictionary

To compile a thesaurus definition file, run the db2extth command.

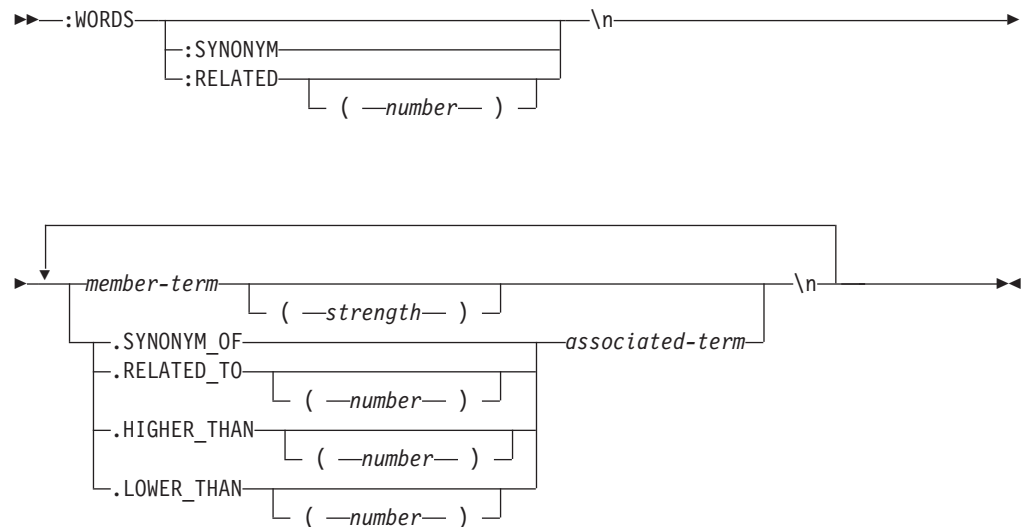
To use a thesaurus dictionary within a partitioned environment, ensure that all the physical nodes can access the created files.

---

## Thesaurus support

Here is the syntax of each definition group when you create your own thesaurus:

### Syntax of a thesaurus definition



Note that \n is not part of the syntax, but represents the end of a line in the thesaurus definition file.

You can insert comment lines in a thesaurus definition file like this:

```
# my comment text
```

**:WORDS**

A keyword that begins a group of related words.

**:SYNONYM, :RELATED [(number)],**

A relation name.

Relation names consist of a relation type and a number. If the number is omitted, zero is assumed, which is the system-provided relation name.

:SYNONYM is always the system-provided relation name.

Relation names that begin with a colon, such as :SYNONYM, precede a list of words that are related to each other by the same relation. For example:

```
:WORDS
:SYNONYM
  air steward
  cabin staff member
  flight attendant
```

*member-term*

A term to be included in the thesaurus dictionary.

- Maximum length is 64 bytes (42 bytes for code page UTF-8).
- Single-byte characters and double-byte characters of the same letter are regarded as the same.
- Uppercase and lowercase characters are not distinguished.
- A term can contain a blank character.
- The single-byte character period "." or colon ":" cannot be used.

This parameter can be useful if you do not want a thesaurus lookup to include words that have a weak relation to the looked up term. Strength is a numerical value from 1 to 100. The default value is 100.

**.SYNONYM\_OF, .RELATED\_TO [(number)], .HIGHER\_THAN [(number)], .LOWER\_THAN [(number)]**

A relation name. The relation name .HIGHER\_THAN corresponds to the BROADER query relation and .LOWER\_THAN to the NARROWER query

relation. Relation names consist of a relation type and a number. If the number is omitted, zero is assumed, which is the system-provided relation name. The relation name .SYNONYM is always the system-provided relation name.

Relation names that begin with a period, such as .SYNONYM\_OF, define the relation between one word and another. For example:

```
:WORDS
  air steward
  .SYNONYM_OF cabin staff member
  .SYNONYM_OF flight attendant
```

The optional *number* identifies a user-defined relation. This must be a unique number from the whole thesaurus definition file (currently 1 to 128). For example: RELATED\_TO(42).

If you want to use symbolic names for thesaurus relations in your application instead of the relation name and number, your application must handle the name-to-number mapping. For example, if you define the relation *opposite\_of* as RELATED\_TO(1), your application must map this name to the internal relation name RELATED\_TO(1).

#### *associated-term*

Each associated term must be preceded by the relation name. The associated term is related to each member term with respect to the specified relation. If all member terms are related to each other, this can be specified using a member relation.

- Maximum length is 64 bytes (42 bytes for code page UTF-8).
- Single-byte characters and double-byte characters of the same letter are regarded as the same.
- Uppercase and lowercase characters are not distinguished.
- A term can contain a blank character.
- The single-byte character period "." or colon ":" cannot be used.

Here is an example of an associated term:

```
:WORDS:SYNONYM
  reject
  decline
  RELATED_TO(1) accept
```

---

## Thesaurus supported CCSIDs

The following CCSIDs are supported by the thesaurus:

819	Latin 1
850	PC Data Latin 1
874	Thai
932	Combined Japanese
943	Combined Japanese
949	Combined Korean
950	Combined Traditional Chinese
954	Japanese
970	Combined Korean
1208	UTF 8

1250	Latin 2
1252	Latin 1
1253	Czech
1254	Turkish
1255	Hebrew
1256	Arabic
1258	Vietnamese
1363	Combined Korean
1381	Combined Simplified Chinese
1383	Chinese (simplified), combined SBCS/DBCS
1386	Chinese (simplified), combined SBCS/DBCS
5039	Japanese (combined SBCS/DBCS)

---

## Messages returned by the thesaurus tool

### ADM\_MSG\_INVALID\_CCSID

Invalid CCSID specified.

The requested code page is not supported.

### ITL\_THES\_MSG\_BUFFER\_OVERFLOW

Buffer overflow.

### ITL\_THES\_MSG\_DICT\_EXIST

Thesaurus dictionary *dictionary name* already exists.

Cannot be overwritten.

### ITL\_THES\_MSG\_DICT\_INTEGRITY\_ERROR

Integrity of dictionary *dictionary name* is lost.

The thesaurus dictionary file is corrupted.

### ITL\_THES\_MSG\_DICT\_NOT\_EXIST

Thesaurus dictionary *dictionary name* does not exist.

### ITL\_THES\_MSG\_DICT\_VERSION\_ERROR

Dictionary *dictionary name* version error.

The thesaurus dictionary was created with an incompatible earlier version.

### ITL\_THES\_MSG\_ERROR\_IN\_FILE

Error in file *file name*.

### ITL\_THES\_MSG\_FILE\_ACCESS\_ERROR

Could not access file *file name*.

### ITL\_THES\_MSG\_FILE\_CLOSE\_ERROR

Could not close file *file name*.

### ITL\_THES\_MSG\_FILE\_EOF\_ERROR

Unexpected end of file in *file name*.

Error in definition file.

**ITL\_THES\_MSG\_FILE\_OPEN\_ERROR**  
Could not open file *file name*.

**ITL\_THES\_MSG\_FILE\_REACHED\_END**  
Unexpected end of file in *thesaurus definition file*.  
There is an error in the definition file.

**ITL\_THES\_MSG\_FILE\_READ\_ERROR**  
Could not read file *file name*.

**ITL\_THES\_MSG\_FILE\_REMOVE\_ERROR**  
Could not remove file *file name*.

**ITL\_THES\_MSG\_FILE\_RENAME\_ERROR**  
Could not rename file *file name 1* to *file name 2*.

**ITL\_THES\_MSG\_FILE\_WRITE\_ERROR**  
Could not write file *file name*.

**ITL\_THES\_MSG\_IE\_BLOCK\_START**  
No block starting line was found in file *file name* at line *line number*.

**ITL\_THES\_MSG\_IE\_EMPTY**  
The thesaurus definition file *file name* is empty.

**ITL\_THES\_MSG\_IE\_NO\_TERM**  
No terms are defined in *file name* at line *line number*.

**ITL\_THES\_MSG\_IE\_REL\_SYNTAX**  
Relationship is specified incorrectly in *file name* at line *line number*.

**ITL\_THES\_MSG\_IE\_STRENGTH\_DOMAIN**  
Strength is out of range.  
Valid values are 1 - 100; the default is 100.

**ITL\_THES\_MSG\_IE\_STRENGTH\_SYNTAX**  
A strength value is specified incorrectly.  
Syntax: After the term, type [ :20 ] for a strength of 20.

**ITL\_THES\_MSG\_IE\_TERM\_LEN**  
A thesaurus term is longer than 64 characters.

**ITL\_THES\_MSG\_IE\_USER\_DEF**  
Relationship is specified incorrectly in *file name* at line *line number*.

**ITL\_THES\_MSG\_IE\_USER\_DEF\_DOMAIN**  
A relationship number is out of range in *file name* at line *line number*.

**ITL\_THES\_MSG\_INPUT\_ERROR**  
Error in the thesaurus definition file *file name* at line *line number*.

**ITL\_THES\_MSG\_INTERNAL\_ERROR**

Internal error.

**ITL\_THES\_MSG\_LOCKED**

Thesaurus dictionary *dictionary name* is in use.

**ITL\_THES\_MSG\_LOCKING\_ERROR**

Could not lock dictionary *file name*.

**ITL\_THES\_MSG\_MEMORY\_ERROR**

Memory error.

**ITL\_THES\_MSG\_NAMELEN\_ERROR**

Parameter error *file name*. The thesaurus definition file name is too long.

**ITL\_THES\_MSG\_NO\_TARGET\_DIR\_ERROR**

Parameter error. No target directory specified.

**ITL\_THES\_MSG\_NONAME\_ERROR**

Parameter error. No thesaurus definition file name specified.

**ITL\_THES\_MSG\_NORMALIZE\_ERROR**

Error in normalizing a term.

Error in the thesaurus definition file.

**ITL\_THES\_MSG\_OUTFILE\_EXIST**

Output file *file name* already exists.

**ITL\_THES\_MSG\_PARAMETER\_ERROR**

Internal parameter error.

**ITL\_THES\_MSG\_PATHLEN\_ERROR**

Parameter error *file name*. The thesaurus definition file path is too long. The path length must not exceed the maximum length supported for directory names in the operating system.

**ITL\_THES\_MSG\_UNEXPECTED\_ERROR**

Internal unexpected error.





---

## Chapter 43. Net Search Extender indexing configuration

This chapter provides information on some configuration options to modify the indexing and search behavior of Net Search Extender .

- Tokenization
- Stop words
- Configuration

---

### Tokenization

During indexing, Net Search Extender processes document text in the following way, breaking the text up into tokens.

#### Words

All alphanumeric characters ("a".."z,"A".."Z", "0".."9") are used to create the full-text index. Separation characters are blank characters and the characters described in the sentence recognition section below. Control characters, such as line feed (also known as a new line character) and blank characters, are interpreted as follows: Control characters (less than 0x20) in the middle of the line are regarded as blank characters. Blank characters and control characters before and after a line feed (0x0A) are ignored. Line feed before and after a 1-byte character are regarded as blank characters and 2-byte characters for the same character are always regarded as the same characters. Capital letters and small letters for the same character, for example, "A" and "a", are regarded as the same characters if nothing is specified during search, or as different characters if exact matching is required during search.

#### Sentences

Net Search Extender recognizes ".", "!", "?" provided the following conditions are met:

- A special character "." has to be followed by a blank space or a new line to be considered the end of a sentence.
- The characters "!" or "?" anywhere in the text will mark the end of a sentence (even without a space or new line).
- If "!", "?", or "." are contained within quotation marks, they are ignored and are not considered to mark the end of a sentence.

#### Paragraphs

Paragraph recognition is dependent on the document format. In Plain Text format, any two consecutive new line characters (possibly with an intervening carriage return) are recognized as a paragraph boundary. In HTML, the paragraph tag <p> is interpreted as paragraph boundary. The other document formats do not support paragraph recognition.

---

## Stop words

Stop words are words with a high frequency and no relevant content for the text retrieval process. Usually, all function words (in a linguistic sense) are considered stop words, for example "and", "or", and "in". Searching an index for stop words can reduce the precision of a text retrieval system significantly.

Net Search Extender provides stop word processing for a list of languages. The configuration parameter `IndexStopWords` can be set at index creation time and determines if stop words are indexed or not. The default is 1 meaning that stop words are indexed.

If you do not want to index stop words, you must set `IndexStopWords` to 0, and specify the language of your input documents using the `language` parameter during create index. If stop words are not indexed, the index is smaller and faster. Do not alter this value in the configuration `.ini` file template after you have created the index, as this leads to documents being treated differently depending on the time that they were indexed and consequently results in incoherent stop word handling.

Ignoring stop words during indexing is effective only if all of the documents in your collection are in the same language.

## Languages supporting stop words

The following languages provide stopword processing.

<b>AR_AA</b>	Arabic as spoken in Arabic countries
<b>CA_ES</b>	Catalan as spoken in Spain
<b>DA_DK</b>	Danish as spoken in Denmark
<b>DE_CH</b>	German as spoken in Switzerland
<b>DE_DE</b>	German as spoken in Germany
<b>EL_GR</b>	Greek as spoken in Greece
<b>EN_GB</b>	English as spoken in the U.K.
<b>EN_US</b>	English as spoken in the U.S
<b>ES_ES</b>	Spanish as spoken in Spain
<b>FI_FI</b>	Finnish as spoken in Finland
<b>FR_CA</b>	French as spoken in Canada
<b>FR_FR</b>	French as spoken in France

**HE\_IL** Hebrew as spoken in Israel  
**IS\_IS** Icelandic as spoken in Iceland  
**IT\_IT** Italian as spoken in Italy  
**IW\_IL** Hebrew as spoken in Israel  
**NB\_NO**  
Norwegian Bokmal as spoken in Norway  
**NL\_BE**  
Dutch as spoken in Belgium  
**NN\_NO**  
Norwegian Nynorsk as spoken in Norway  
**PT\_BR**  
Portuguese as spoken in Brazil  
**PT\_PT**  
Portugese as spoken in Portugal  
**RU\_RU**  
Russian as spoken in Russia  
**SV\_SE**  
Swedish as spoken in Sweden

---

## Configuration

Net Search Extender is able to search for words which may have characters used in different combinations, for example, alphanumeric, numbers, and special characters. To do this, Net Search Extender provides the following configurations:

### Character normalization

Character normalization ensures that words that can be written in two ways can be both searched for. For example, the German word 'Überbau' can also be written as 'Ueberbau'. Normalization ensures that both words can be searched for, by using either 'Überbau' or 'Ueberbau'. The functionality also normalizes accented letters, for example, 'accès' to the matching simple character, for example, 'acces'. Note that the use of this option can have undesired results in languages where for example the character 'Ü' does not have an equivalent standard normalization as 'Ue'

### Using specific characters as part of a word

Using specific characters as part of a word ensures that product names which can involve a series of alphanumeric characters, special characters and numbers can be searched on as a single word. For example, by treating the alphanumeric combination 'DT9' as one word, or by enabling the '/' special character, so that AS/400® or OS/390® are searched on as whole words rather than as 'AS' and '400', or 'OS' and '390'.

For these configuration settings, switches are available. To customize the switches, change the .ini file template before creating an index.

The .ini file template is stored in `sqllib/db2ext/cteixcfg.ini`. As you can also make changes to most of the values in this template file using the CREATE INDEX command, it is recommended that you only change the following values:

AccentRemoval (for character normalization)  
UmlautNormalization (for character normalization)  
TreatNumberAsWords (for treating numeric characters as part of the word)  
AdditionalAlphanumCharacters (for using specific characters as part of a word)

#### **AccentRemoval**

This parameter specifies if accented characters are normalized to the matching simple character. For example, événement is also indexed as evenement. The default is true.

#### **UmlautNormalization**

This parameter specifies if an umlaut character is also indexed as two characters with the same meaning. For example, 'Übersee' is also indexed as 'Uebersee'. The default is true.

#### **TreatNumbersAsWords**

This parameter specifies if numeric characters next to a word are part of the word. For example, 'DT9' is treated as one word and not as one word 'DT' and the number '9'.

#### **AdditionalAlphanumCharacters**

The string value of this parameter defines which characters are treated as part of a word. The string of special characters must be a sequence of one or more characters in UTF-8. The default string contains the characters `" / - @ "`.

You are not allowed to use the wildcard characters `%` and `_` in the list of characters that are treated as being part of a word. This results in problems during query execution.

If you want to change any of these configuration values, edit the `.ini` file before you create your index. To activate inactive switches, remove the comment marker `;"` from the beginning of the line. For further information, see the `cteixcfg.ini` file.

You are recommended not to alter any of the other values in the `.ini` file.

---

## Part 8. Working with structured documents

Net Search Extender allows you to index and search text or numeric fields, such as title, author, or price in a structured document. The documents can be in XML, Outside In, HTML format, or contain user-defined tags (GPP).

Use markup tags and their field names in a *document model* to define which fields in the documents are indexed and, therefore, are available for searching. You can use the name of the field (also known as the section name) in queries against that field.

To be able to search in these fields, you must specify a *FORMAT* and a *MODEL* file when you create the text index containing the documents.



---

## Chapter 44. Searching natively stored XML documents

The following sections address search on natively stored XML documents. You are shown how the concepts of section search can be applied to natively stored XML documents and how to integrate this functionality into XQuery processing.

Generally, when creating an index on an XML data column, you are not required to specify a FORMAT. Net Search Extender selects format XML by default when a text index is created on a column of type XML. The format specifiers TEXT and HTML are not allowed on XML data columns.

In the subsequent samples that illustrate the creation and use of a text index on XML columns, the following XML document is used. It is stored in table t1, column c2 of type XML.

```
<?xml version="1.0">
<purchaseOrder orderDate="2001-01-20">
  <shipAddress countryCode="US">
    <name>Alice Smith</name>
    <street>123 Maple Street</street>
    <city>Mill Hill</city>
    <zip>90999</zip>
  </shipAddress>
  <item partNo="123" quantity="1">
    <name>S&B Lawnmower Type ABC-x</name>
    <price>239.90</price>
    <shipDate>2001-01-25</shipdate>
  </item>
  <item partNo="987" quantity="1">
    <name>Multifunction Rake ZYX</name>
    <price>69.90</price>
    <shipDate>2001-01-24</shipdate>
  </item>
</purchaseOrder>
```

---

### Using the default document model

If no document model is specified in the CREATE INDEX statement, Net Search Extender uses the default document model.

One characteristic of the default document model is that section names are in XPath notation specifying the absolute path to each element and attribute. Note that section names in the search query are not XPath expressions that are evaluated during query execution. Instead they are names referring to specific parts (elements and attributes) within structured documents.

If you are not using a model file, define a text index for XML documents as follows:

```
db2text CREATE INDEX i1 FOR TEXT ON t1(c2) CONNECT TO mydbname
```

As column c2 is of datatype XML, you can omit the FORMAT specification. The FORMAT specification is set to XML by default in this case.

When no document model is specified, each XML element is assigned a name automatically depending on its absolute XPath within the document. For example, element price can be accessed by the section name /purchaseOrder/item/price in

the search query. The attribute `countryCode` can be accessed by using the section name `/purchaseOrder/shipAddress/@countryCode`.

After the index is updated by using the `db2text` update command, a possible SQL expression using `SECTION` search with the scalar search function might be as follows:

```
SELECT c2 FROM t1
WHERE CONTAINS(c2, SECTIONS("/purchaseOrder/item/name") "Rake") = 1
```

The query returns the sample XML document shown above.

---

## Using a customized document model

If you want to define customized section names, you must specify a model file that assigns user defined names to certain parts of a document. An advantage of using a document model is that you can specify which parts of an XML document you want to index and use XPath expressions to specify these parts.

A model file for the above XML document might look as follows:

```
<?xml version="1.0"?>
<XMLModel>
  <XMLFieldDefinition
    name="itemName"
    locator="/purchaseOrder/item/name" />
  <XMLFieldDefinition
    name="customerName"
    locator="//shipAddress/name" />
  <XMLAttributeDefinition
    name="partNumber"
    type="NUMBER"
    locator="/purchaseOrder//item/partNo" />
  <XMLFieldDefinition
    name="none"
    locator="/purchaseOrder/orderDate"
    exclude="yes" />
</XMLModel>
```

Note that the document model assigns the name `itemName` to the element `/purchaseOrder/item/name` which is referenced in the previous search query above.

The index definition, using the model file, is:

```
CREATE INDEX i1 FOR TEXT ON t1(c2) DOCUMENTMODEL XMLModel IN
  /mydir/myfilename/xmlmodel.xml CONNECT TO mydbname
```

The document model name (using the parameter `DOCUMENTMODEL`) specifies the root element in the model file. This is `XMLModel` for XML document models. The path `/mydir/ ...` points to the file that defines the model.

The document model syntax supports a subset of the W3C XPath syntax which allows for the convenient identification of elements.

After creating the text index using the model file above, and updating the index using the `db2text` update command, it is possible to search for the element `/purchaseOrder/item/name` as follows:

```
SELECT c2 FROM t1
WHERE CONTAINS(c2, SECTIONS("itemName") "Rake") = 1
```



Note the difference to the search query where no document model was specified. Both queries return the same sample XML document mentioned above.

The XML document model also defines an attribute `partNumber` on the XML attribute `partNo` of element `item`. The data type of Net Search Extender attribute definitions must always be `NUMBER`.

The attribute definition in the sample model file above allows searching on value ranges like:

```
SELECT c2 FROM t1 WHERE CONTAINS
(c2, ATTRIBUTE "partNumber" BETWEEN 300 AND 500) = 1
```

---

## XQuery support

While searching for XML documents in the database, it is also possible to process the search results by using XQuery. By exploiting DB2's hybrid database engine, an SQL text search query can be combined with XQuery processing.

This is done by using the `db2-fn:sqlquery()` input function in the XQuery context. In order to use the XQuery input function, you must switch from SQL to XQuery by using the set language `XQuery` command, or the query has to be prefixed with the keyword `XQuery`. This is an important indicator to the parser that it is working with an XQuery expression and must follow the case sensitivity rules and syntax rules that apply to the XQuery language.

The `db2-fn:sqlquery()` function takes a string literal that represents a full-select. The `db2-fn:sqlquery()` function returns an XML sequence that represents the concatenation of the XML column values that are selected by the full-select.

The following expression can be used to combine text search and XQuery processing on natively stored XML documents:

```
XQUERY db2-fn:sqlquery('SELECT c2 FROM t1
    WHERE CONTAINS(c2,
        ''SECTIONS ("/purchaseOrder/item/name") "Rake" '')
    = 1 ')//shipAddress/name
```

The above query returns all the name elements under the `shipAddress` element in XML documents that contain a purchasing order item named "Rake". You must explicitly select the XML column (in our case `c2`) in the `SELECT` statement.

The above sample can be extended by a `FLWOR` construct as follows, and embedded in your application:

```
XQUERY FOR $item in db2-fn:sqlquery('SELECT c2 FROM t1
    WHERE CONTAINS(c2, '' SECTIONS ("/purchaseOrder/item/name") "Rake" '')
    = 1 ')
    WHERE $item[@partNo > "800"]
RETURN $item/price
```

Note that the full-select of the `db2-fn:sqlquery()` input function always returns the complete XML document in which a hit occurs.

Consider the following XML document that is stored natively in the database:

```
<?xml version="1.0" ?>
<dept bldg="101">
  <employee id="901">
    <name>Sabine</name>
    <resume>DB2 programmer</resume>
```

```

</employee>
<employee id="902">
  <name>Holger</name>
  <resume>XML expert</resume>
</employee>
</dept>

```

Searching for an employee in your department where the term “XML” is contained in the resume might look as follows:

```
SELECT c2 FROM t1 WHERE CONTAINS(c2, SECTIONS("/dept/employee/resume") "XML")=1
```

The above select returns the complete XML document. Embedding the search query in XQuery as follows:

```
XQUERY db2-fn:sqlquery('SELECT c2 FROM t1
  WHERE CONTAINS(c2,
    ''SECTIONS ("/dept/employee/resume") "XML" '' ) =1') //employee/name
```

returns the following two results:

```

<name>Sabine</name>
<name>Holger</name>

```

Notice that although the employee Sabine does not have the term “XML” in her resume, she appears in the resulting sequence in the above XQuery. This happens because the full-select returns the whole document, that is, it returns the complete XML document that has at least one employee with the term “XML” in the resume.

If you want the query to return only the result <name>Holger</name>, issue the following XQuery statement:

```
XQUERY for $d in db2-fn:sqlquery('SELECT c2 FROM t1
  WHERE CONTAINS(c2,
    ''SECTIONS ("/dept/employee/resume") "XML" '' ) =1')
  return $d/dept/employee/name[contains(parent::employee/resume,"XML")];
```

Net Search Extender filters out all XML documents that have the term XML in the section /dept/employee/resume by using a structure sensitive full-text index on the XML column. Base on the returned subset of XML documents, the return statement return \$d/dept/employee/name[contains(parent::employee/resume,"XML")] returns only those <name> elements that have XML in their sibling element called <resume> by navigating the XML document using the XPath axis.

---

## Chapter 45. Structured document support

---

### How a document model describes structured documents

Documents in HTML or XML format are examples of structured documents; they contain tags that identify text fields and document attributes. Text fields can contain information like the title, author, or a description of the document.

The following is an extract from a structured plain-text document. It contains text that is delimited by HTML-like tags.

```
[head]Handling structured documents
[/head]
```

```
[abstract]This document describes the concept of structured documents
and the use of document models to...
[/abstract]
```

```
:
:
```

When Net Search Extender indexes structured documents, it has to recognize the structure so that it can index the text field and the attributes, and store them together with a unique name. This enables Net Search Extender to selectively search in a particular text field or to find documents that have a particular attribute by using the SECTION or ATTRIBUTE clause.

To enable Net Search Extender to understand the structure of a particular document format, you must pass Net Search Extender a definition of the structure in a *document model*. Alternatively, you can use the default document models provided by Net Search Extender.

You specify the name of the document model as an argument when you call the CREATE INDEX command to index the documents. For example, CREATE INDEX i1 FOR TEXT ON t1(c2) DOCUMENT MODEL GPPModel IN mymodel.xmld CONNECT TO db

The parameter GPPModel refers to the type of document model you are using.

Before you can index documents using a document model, you must first define a document model and then make the document model known to the index.

**Note:** If XML documents use indexes that are not well-formed, the indexing process will stop at the point where the problem is encountered in the document. This means that only a part of the document will be indexed. If you do not correct the document, you will only be able to search in the parts of the document that were indexed. This will occur only if your table column type is not XML.

---

### An example of a document model

You must define one document model for each document format that you intend to index. Here is a simple document model for plain-text structured documents. Note that GPP in the example stands for General Purpose Parser.

```
<?xml version="1.0"?>
<GPPModel>                                - the GPP document model begin here
    <GPPFieldDefinition                    - a field definition begins here
```

```

name="Head"           - the name you assign to this field
  start="[head]"      - the boundary string at the beginning of the field
  end="[/head]"       - the boundary string at the end of the field
  exclude="YES" />

<GPPFieldDefinition  - the next field definition begins here
  name="Abstract"
  start="[abstract]"
  end="[/abstract]"
  exclude="NO" />
:
:
</GPPModel>

```

Document models are specified in the XML language using tags as defined in Chapter 46, "Document model reference," on page 177. A document model consists of text field definitions and attribute definitions. The example above illustrates only text field definitions defined in GPPFieldDefinition elements. In a similar way, you can use GPPAttributeDefinition to define document attributes.

The first line `<?xml version="1.0"?>` in the example specifies that the document model is written using XML tags. Each of the text field definitions specifies boundary strings to identify the start and end of the field definition in the source document. So, whenever a document contains the sequence of characters `[head]` followed by some text and the sequence of characters `[/head]`, the text between those boundary strings is taken to be the content of the text field that is identified by the name `head`.

You assign a field name to each field definition. This field name is the means by which a query can restrict search to the content of a text field using a SECTION clause in the CONTAINS function. The name of the field can be either fixed or can be derived by a rule from the structural unit's content. Such a name could be, for example, the tag name of an XML entity, or the name of an XML attribute.

---

## Document models

A document model primarily controls what parts of a document's structure need to be indexed and how they are indexed. Its purpose is to:

- Identify text fields that should be distinguished in the source document
- Determine the type of such a text field
- Assign a field name to the text field

When the document model identifies text as belonging to a text field, the text is considered to be part of the textual content of the document, and terms are extracted and stored in the index.

The elements of a document model vary depending on the parser used for that document format:

- For HTML format, a document model uses the HTML tag names to define which tags should be indexed, and how to handle meta-tag information.
- For XML format, there is no predefined set of tags, so a document model must first define which tags are of interest. XML elements of the same name can also be distinguished based on what other elements they are embedded in.
- For GPP (general purpose parser) format, the document model interacts even more deeply with the parser, because it has to determine the boundaries of the text fields. Here the field definition must specify strings for detecting the boundaries of fields.

- For Outside In formats, a document model uses tags similar to HTML tag names to define which tags should be indexed, and how to handle meta-tag information. Note that the Outside In Transformation Technology is also known as INSO.

## Text fields

A document model lets you identify document parts or sections as either belonging to a particular text field or as being a document attribute, or both.

The text of a document is fully indexed regardless of whether or not it is part of a text field. Meaningful terms are extracted and stored in the index. This means that unrestricted text searches include a search of that text.

However, by defining text fields, you can search for text selectively in a particular field. For example, you can search for documents that contain the word `structure` in the text field `Abstract`. For example, `SELECT doc from my_docs WHERE CONTAINS (doc, SECTIONS(Abstract) "structure" = 1.`

A text field can occur several times in a document. For instance, you can define a text field that contains all figure captions. A text field may also overlap with another text field.

If you want to avoid indexing the content of certain text fields, you can specify a field definition that contains `exclude="YES"`. You can find a list of limitations for text fields and document attributes in "Limitations for text fields and document attributes" on page 180.

## Document attributes

Document attributes contain short, formatted information of type number. In contrast to text fields, you can use value ranges to search documents containing such attributes.

Attributes are not stored with indexed text, but in a separate item index. So, to search for documents by content of an attribute, you must make an attribute search explicitly on the attribute. For example, `SELECT doc FROM my_docs WHERE CONTAINS (doc,ATTRIBUTE "year" BETWEEN 2001 AND 2005) = 1.`

### Number attributes

Net Search Extender provides a parser that recognizes floating-point numbers. The following are examples of correct and incorrect formats for attribute values.

*Table 6. Supported formats for attribute values*

Correct format	Incorrect format
1000 1 000 1.000 - where the period is a decimal character	1,000
100 000 100 000.00123	1 000 000 - two spaces between 1 and 0

Note that space characters are not allowed in the decimal fraction of a number. For example, 1 000.000 100 is treated as two numbers, 1000.000 and 100.

Language-specific separators and language-specific monetary formats are not supported.

## Default document models

For HTML, XML, and Outside In filtered documents, Net Search Extender provides default document models that are used if you do not specify a document model during index creation. For structured plain text documents, you must provide and specify a document model.

If you use one of the default document models:

- All fields are indexed, and no special information, such as meta information, is extracted.
  - For HTML and INSO formats, each field is assigned the name of the corresponding tag.
  - For XML, all XML nodes of an XML document are mapped to overlapping fields which are identified by the fully qualified element paths of the corresponding nodes. For example, the path `/play/role/name`.
- No numeric attribute is indexed (as no numeric attribute is defined in the default document model).

*Table 7. Behavior of the default document models for the supported document formats*

Document type	Behavior of the default document model
HTML	Accepts these as text fields: <code>&lt;a&gt;</code> <code>&lt;address&gt;</code> <code>&lt;au&gt;</code> <code>&lt;author&gt;</code> <code>&lt;h1&gt;</code> <code>&lt;h2&gt;</code> <code>&lt;h3&gt;</code> <code>&lt;h4&gt;</code> <code>&lt;h5&gt;</code> <code>&lt;h6&gt;</code> <code>&lt;title&gt;</code> . Field name is the tag name, for example "address".
XML	Accepts all tags as text fields. The field name is the fully qualified element path name, for example <code>"/play/title"</code> .
Structured plain text (GPP)	No default document model.
Outside In (INSO)	Accepts as text fields, the document properties shown in "Definition of a document model for Outside In filtered documents" on page 174 as returned by the Outside In filters. The Field name is the name of the document property used by Outside In, for example: "SCCCA_TITLE".

For each type of document, a default document model is defined. As each model is different, an example and explanation is provided for each one in the following sections.

### Note:

Although the default document models do correctly process documents, for better indexing and search you should define your own document models.

With the default document model, the text of a document is fully indexed regardless of whether or not it is part of a text field. This means that unrestricted text searches include a search of that text.

---

## Definition of a document model for structured plain-text documents

These are the parameters of the document model elements:

**name** You assign a name to the text field or document attribute for each definition. The names enable you to restrict a search query to the content of a specific text field or document attribute. Using the above examples, you could search for documents containing the word structure in the text field named Abstract.

**start** A boundary string in codepage UTF-8 that marks the beginning of the text field or document attribute. There are no rules for specifying boundary strings; they can be any arbitrary UTF-8 string. Here are some examples: start="introduction:", start="note!", start="\$\$...."

Nonprintable characters and the special XML characters "<" and "&" must be specified using the default XML character entries ("&lt;" for "<", and "&amp;" for "&").

**end** Optional. A boundary string in codepage UTF-8 that marks the end of the text field or document attribute. If you do not specify an end tag, the next found start tag is assumed to be the end of the field. If no subsequent start tag is found, the field extends to the end of the document, and no further fields are identified.

**type** The type of document attribute must always be "NUMBER". The parameter does not apply to field definitions.

**exclude**

YES or NO. A parameter that determines whether the text in a field definition should be excluded and therefore, not indexed. This parameter does not apply to attribute definitions.

In the example, the field definition "head" would be excluded, but definition "abstract" included.

**Restrictions:**

- There must not be two field definitions or attribute definitions with the same start tag. However, a field definition and attribute definition may have the same start tag and end tags.
- A start tag must not be a proper prefix of another. For example, you cannot have a start tag "author" and a start tag "authority".
- Start tags and end tags must not be empty strings.

## What happens when a GPP document is indexed

The general-purpose parser scans the document looking for one of the start boundary strings. When it finds a start boundary string, it parses the subsequent field until it finds the corresponding end boundary string.

The content of the field is then indexed according to the definition term, that is, as a text field or document attribute. If the text field and document attribute have the same start and end boundary strings, the content of the field is indexed as both a text field and a document attribute.

No nesting of fields is allowed; if a new start boundary string is found in a field before the end boundary string has been reached, the new start boundary string is interpreted as normal text.

If no corresponding end boundary string is found, the field is assumed to extend to the end of the document, and an appropriate reason code is reported.

If no end boundary string is specified in the document model, the new start boundary string signals the end of the previous field.

---

## Definition of a document model for HTML documents

The HTML parser converts the text to codepage UTF-8. It performs HTML tag recognition and classifies them into tag classes:

- Tagged information to be ignored, such as font information
- Tags that provide positional information, such as <p>; for new paragraph
- Tags that provide structural information, such as <Title>

It recognizes all character entity references defined in HTML 4, like "&auml;" (ä) and resolves them to the corresponding code points in UTF-8.

It recognizes meta tags and parses the meta tag text.

Here is an example of an HTML document:

```
<HTML>
<HEAD>
<META NAME="year" CONTENT="2002">
<TITLE> The Firm </TITLE>
</HEAD>
<BODY>
<H1>Synopsis</H1>;
<H1>Prologue</H1>;:
:
</BODY>
```

Here is an example of an HTML document model:

```
<?xml version="1.0"?>
<HTMLModel>

  <HTMLFieldDefinition
    name="subtitle"
    tag="title"
    exclude="NO" />

  <HTMLFieldDefinition
    name="header1"
    tag="h1"
    exclude="YES" /> - This is the start of text field
                    - This is the end of the text field

  <HTMLAttributeDefinition
    name="year"
    tag="meta"
    meta-qualifier="year"
    type="NUMBER" /> - This is the start of the document
                    attribute
                    - This is the end of the document
                    attribute

</HTMLModel>
```

The first line, <?xml version="1.0"?>, specifies that the document model is written using XML tags. Note that this model is not written for XML format documents.



Each field is defined within a `HTMLFieldDefinition` or `HTMLAttributeDefinition` tag, which contain element parameters.

All the text field definitions must be contained within the `<HTMLModel>` tag. The tag name is passed as a parameter during index creation: `CREATE INDEX iA FOR TEXT ON T1(C2) DOCUMENTMODEL HTMLModel IN myModel.xml CONNECT TO db.`

These are the parameters of the document model elements:

- name** You assign a name to the text field or document attribute for each definition. The names enable you to restrict a search query to the content of a specific text field or document attribute. Using the above examples, you could search for documents containing the word `firm` in the text field named `subtitle`.
- tag** Identifies an element whose start and (implied) end tags mark the text field or document attribute. The text that is inside an element of that name makes up the content of the defined field.

The case of the tag is ignored.

Using the above examples, the text following any `H1` tag is indexed as being part of the field `"header1"`. Based on the sample document, `"synopsis"` and `"prologue"` would be indexed.

#### **meta-qualifier**

This tag has to be used with the **tag** element. By specifying `tag="meta"`, the value of the content that matches the meta-qualifier is extracted.

In the HTML document example, the meta tag has the following elements:

```
<META NAME="year" CONTENT="2002">
```

In the document model example, `meta-qualifier="year"`. Therefore, the content `"2002"` is indexed as the value of the attribute `"year"`.

- type** The type of document attribute must be `"NUMBER"`. The parameter does not apply to field definitions.

#### **exclude**

YES or NO. A parameter that determines whether the text in a field definition should be excluded and therefore, not indexed. This parameter does not apply to attribute definitions.

In the example, the field definition `"header1"` would be excluded, but definition `"subtitle"` included.

All other text of a document is indexed, but not as part of any field.

## **Element parameters**

These are the parameters of the document model elements:

- name** You assign a name to the text field or document attribute for each definition. The names enable you to restrict a search query to the content of a specific text field or document attribute. Using the above examples, you could search for documents containing the word `firm` in the text field named `subtitle`.
- tag** Identifies an element whose start and (implied) end tags mark the text field or document attribute. The text that is inside an element of that name makes up the content of the defined field.

The case of the tag is ignored.

Using the above examples, the text following any H1 tag is indexed as being part of the field "header1". Based on the sample document, "synopsis" and "prologue" would be indexed.

### meta-qualifier

This tag has to be used with the **tag** element. By specifying tag="meta", the value of the content that matches the meta-qualifier is extracted.

In the HTML document example, the meta tag has the following elements:

```
<META NAME="year" CONTENT="2002">
```

In the document model example, meta-qualifier="year". Therefore, the content "2002" is indexed as the value of the attribute "year".

**type** The type of document attribute must be "NUMBER". The parameter does not apply to field definitions.

### exclude

YES or NO. A parameter that determines whether the text in a field definition should be excluded and therefore, not indexed. This parameter does not apply to attribute definitions.

In the example, the field definition "header1" would be excluded, but definition "subtitle" included.

All other text of a document is indexed, but not as part of any field.

---

## Definition of a document model for XML documents

A document model for XML documents allows you to define how an element found in an XML document is mapped to a field, a document attribute, or both.

Here is an example of an XML document:

```
<?xml version="1.0"?>
<purchaseOrder orderDate="2001-01-20"> [4]
  <shipAddress countryCode="US"> [1]
    <name>Alice Smith</name> [2]
    <street>123 Maple Street</street>
    <city>Mill Hill</city>
    <state>CA</state>
    <zip>90999</zip>
  </shipAddress>
  <item partNo="123" quantity="1"> [3]
    <name>S&B Lawnmower Type ABC-x</name>
    <price>239.90</price>
    <shipDate>2001-01-25</shipDate>
  </item>
  <item partNo="987" quantity="1"> [3]
    <name>Multifunction Rake ZYX</name>
    <price>69.90</price>
    <shipDate>2001-01-24</shipDate>
  </item>
</purchaseOrder>
```

Here is an example of an XML document model that matches the sample document from above:

```
<?xml version="1.0"?>
<XMLModel>
```

```

<XMLFieldDefinition [1]
name="addresses"
locator="/purchaseOrder/shipAddress" />

<XMLFieldDefinition [2]
name="customerName"
locator="//shipAddress/name"
exclude="yes" />

<XMLAttributeDefinition [3]
name="partNumber"
type="NUMBER"
locator="/purchaseOrder//item/@partNo" />

<XMLFieldDefinition [4]
name="none"
locator="/purchaseOrder/@orderDate" />

</XMLModel>

```

The first line, `<?xml version="1.0"?>`, specifies that the model is written using XML. Each field is defined within a `XMLFieldDefinition` or `XMLAttributeDefinition` tag, which contains element parameters.

Note that all the text field definitions must be contained within the `<XMLModel>` tag. This tag name is passed as a parameter during index creation: `CREATE INDEX i1 FOR TEXT ON T1(C2) DOCUMENTMODEL XMLModel in myModel.xml CONNECT TO db.`

The fields and attributes in the sample are marked with numbers that correspond to the definitions in the example model file.

Nesting of fields is allowed, for example, if the XPath location of one specification selects a node that lies inside an XML element selected by another attribute definition. Nested fields is shown in the sample XML document above. The field `addresses` selects a node in the XML document that dominates the node selected by field `customerName`. The content of that embedded node therefore, logically belongs to both fields. Although text fields may be overlapping, the text inside those fields is indexed only once. In this example, when searching with a field restriction, Alice Smith is found in `addresses` as well as in `customerName`. However, due to the matching semantics of the locator expression, it is not possible to map one and the same XML node to multiple fields.

Net Search Extender does not attempt to detect the code page of an XML document. The DB2 code page is taken.

The content of fields is determined by the following rules:

- For a field whose locator matches a comment, a processing instruction, or an XML attribute, the field content is the actual comment text, processing instruction text, or attribute value text.
- For a field that matches an XML element or the root node, the field content consists of any text from any embedded element except for elements that are matched by fields having the specification `exclude="YES"`.

The document must contain well-formed XML, but it is not necessary for a DTD to be specified in the XML document. No DTD validation or external entity resolution is carried out; Net Search Extender only matches the XML document against the document model. Internal entities are substituted as required by XML.

## Element parameters

These are the parameters of the document model elements:

**name** You assign a name to the text field or document attribute for each definition. These names enable you to restrict a search query to the content of a specific text field or document attribute.

You can use one of the following variables in a name. The variable is replaced by a string generated from the matching element in the source document.

**Variable**  
**Value**

**\$(NAME)**  
The actual qualified name (QName) of the XML element that matched the XPath.

**\$(LOCALNAME)**  
The actual local name (without prefix) of the XML element that matched the XPath.

**\$(PATH)**  
The actual absolute path as a sequence of slashes and tags of the XML element that matched the XPath (used as name in the default document model).

**type** The type of document attribute must be "NUMBER". The parameter does not apply to field definitions.

**locator**  
Expressions in the XPath language that select the parts of the source documents to be used as search fields.  
When writing a XML Document Model file, the qualified names, known as QNames, inside a locator must be identical to some tags in the XML document, otherwise no fields will be recognized and the queries on fields will not return a result.

The following are examples of locators.

**purchaseOrder | salesOrder**  
All purchaseOrder elements and salesOrder elements

**shipAddress**  
All shipAddress elements

**\***  
All elements (this is the abbreviation of `child::*` – see the syntax for further information)

**name/item**  
All item elements that have a name parent

**purchaseOrder//item**  
All item elements that have a purchaseOrder ancestor

**/**  
The root node

**comment()**  
All comment nodes

**processing-instruction()**  
All processing instructions

**attribute::\* (or @\*)**

All attribute nodes

A literal is a string enclosed either in single or double quotes. For an exact definition of terminal tokens see the XML recommendations.

The XPath locators that are supported by the Net Search Extender document model are similar to XML Stylesheet Language Transformation (XSLT) patterns. They comprise exactly the subset of XSLT patterns that do not contain any predicates nor the functions 'id' and 'key' nor the node tests 'text()' and 'node()'.

**ignore** YES or NO. Use the parameter to make exceptions to the locator.

Sometimes you may want to specify a general locator, such as \*, to match the nodes you want to index. But you may also specify that some nodes matching a more specific locator should not be indexed.

To formulate this, include a field definition with the more specific locator for the nodes to be ignored during indexing. You then give this locator a higher priority than the one with the general locator (see below), and specify ignore="yes". This indicates to the indexer that it must not generate field information for the matching nodes.

Note that when such an ignored node is embedded in a field-generating node, the content of the ignored node gets indexed, because it also belongs to the contents of the field-generating node.

**priority**

A floating point number between -1 and +1 that specifies the priority to be given to a definition found by a particular locator.

If you do not specify a priority, the default priorities are used:

- Multiple alternatives separated by | are treated as a set of definitions, one for each alternative.
- Locators that match by a single name; that is, locators of either of the following forms have default priority 0:
  - ChildOrAttributeAxisSpecifier QName
  - ChildOrAttributeAxisSpecifier processing-instruction(Literal))
- Locators of the form ChildOrAttributeAxisSpecifier NCName:\* have default priority -0.25.
- Other locators of the form ChildOrAttributeAxisSpecifier NodeTest have default priority -0.5.
- Any other locator has default priority 0.5.

Note that the more specific the locator is, the higher the default priority. For example, the unspecific locator \* gives a low priority to the found definition, whereas a name is a more specific locator and gives a higher priority.

Also note that when a node is matched by more than one locator, you can determine which of the definitions are chosen by assigning priorities to them. The definition with the highest priority is chosen. If two definitions have the same priority, the latest is chosen.

This conflict resolution is the same as that used in XML Stylesheet Language Transformation (XSLT).

**exclude**

YES or NO. A parameter that determines whether the text in a field

definition should be excluded and therefore, not indexed. This parameter does not apply to attribute definitions.

In the example, the field definition "customerName" would be excluded, but definition "addresses" included.

---

## Definition of a document model for Outside In filtered documents

Document models for the Outside In format are very similar to HTML document models in that they allow you to map structural elements identified by a given set of tags to Net Search Extender text fields and document attributes. Assume you have a set of Microsoft Word documents and want to index the document properties "title", "subject", and "keyword" as fields, and the document properties "author" and "category" as document attributes. The following example for an Outside In document model will achieve this mapping:

```
<?xml version="1.0"?>
<INSOModel>

<INSOFieldDefinition
name="title"
tag="SCCCA_TITLE"/>

<INSOFieldDefinition
name="title"
tag="SCCCA_SUBJECT"/>

<INSOFieldDefinition
name="title"
tag="SCCCA_KEYWORDS"/>

<INSOAttributeDefinition
name="author"
tag="SCCCA_AUTHOR"
type="STRING"/>

<INSOAttributeDefinition
name="category"
tag="SCCCA_CATEGORY"
type="STRING"/>

</INSOModel>
```

### Element parameters

These are the parameters of the document model elements:

- name** A name that you assign to the text field or document attribute. You assign a field name to each field definition, and an attribute name to each attribute definition. These names are the means by which a query can restrict search to the content of a certain text field and can search for documents having a certain attribute.
- tag** Identifies a tag whose begin and end or implied-end elements mark the text field or the document attribute. The text that is inside an element of that name makes up the content of the defined field or attribute. The case of the tag is disregarded. Possible values are described below.
- type** The type of document attribute can be either "NUMBER", "DATE", or "STRING". This parameter does not apply to field definitions.
- exclude** YES or NO. A parameter that determines whether the text in a field

definition should be excluded and therefore, not indexed. This parameter does not apply to attribute definitions.

Outside In document models consist of field and/or attribute definitions that each define a name and a tag. For attribute definitions a type is also required, whereas field definitions have an optional "exclude" flag. As with HTML models, the name attribute of such a definition defines the name of the Net Search Extender field or attribute that the document part is to be mapped on. This can be an arbitrary UTF-8 text string. For additional information, see the Outside In Content Access Specification, Version 7.5.

For a list of possible values for the tag attribute relating to the Outside In begin, end and document property tags.

## **What happens when an Outside In document is indexed**

By default, all the text is indexed as not belonging to any field. Whenever a begin tag that appears in the stream of text matches a definition item in the document model that is currently active, the text between the begin tag and its corresponding end tag is treated according to that definition term. For example, as an indexed field, an excluded field, or as an attribute or both.

If no matching definition exists, the begin tag and its corresponding end tag are ignored.

As Outside In filters automatically recognize the format and code page of the document, the CCSID specification has no effect. If the Outside In filters are unable to determine the correct format and code page, the document is treated as an ASCII file.





---

## Chapter 46. Document model reference

Net Search Extender provides the following reference information for document models:

- The DTD for document models
- The semantics of locator (XPath) expressions
- Limitation for text fields and document attributes
- Outside In tag attribute values

---

### DTD for document models

Here is a formal description of the syntax of document models in the form of a document type definition (DTD):

```
<!ELEMENT GPPModel (GPPFieldDefinition|GPPAttributeDefinition)+>
<!ELEMENT HTMLModel (HTMLFieldDefinition|HTMLAttributeDefinition)+>
<!ELEMENT XMLModel (XMLFieldDefinition|XMLAttributeDefinition)+>

<!ELEMENT GPPFieldDefinition EMPTY>
<!ATTLIST GPPFieldDefinition name CDATA #REQUIRED>
<!ATTLIST GPPFieldDefinition start CDATA #REQUIRED>
<!ATTLIST GPPFieldDefinition end CDATA #IMPLIED>
<!ATTLIST GPPFieldDefinition exclude (YES|NO) NO>

<!ELEMENT GPPAttributeDefinition EMPTY>
<!ATTLIST GPPAttributeDefinition name CDATA #REQUIRED>
<!ATTLIST GPPAttributeDefinition start CDATA #REQUIRED>
<!ATTLIST GPPAttributeDefinition end CDATA #REQUIRED>
<!ATTLIST GPPAttributeDefinition type NUMBER #REQUIRED>

<!ELEMENT HTMLFieldDefinition EMPTY>
<!ATTLIST HTMLFieldDefinition name CDATA #REQUIRED>
<!ATTLIST HTMLFieldDefinition tag CDATA #REQUIRED>
<!ATTLIST HTMLFieldDefinition meta-qualifier CDATA #IMPLIED>
<!ATTLIST HTMLFieldDefinition exclude (YES|NO) NO>

<!ELEMENT HTMLAttributeDefinition EMPTY>
<!ATTLIST HTMLAttributeDefinition name CDATA #REQUIRED>
<!ATTLIST HTMLAttributeDefinition tag CDATA #REQUIRED>
<!ATTLIST HTMLAttributeDefinition meta-qualifier CDATA #IMPLIED>
<!ATTLIST HTMLAttributeDefinition type NUMBER #REQUIRED>

<!ELEMENT XMLFieldDefinition EMPTY>
<!ATTLIST XMLFieldDefinition name CDATA #REQUIRED>
<!ATTLIST XMLFieldDefinition locator CDATA #REQUIRED>
<!ATTLIST XMLFieldDefinition ignore (YES|NO) NO>
<!ATTLIST XMLFieldDefinition priority CDATA #IMPLIED>
<!ATTLIST XMLFieldDefinition exclude (YES|NO) NO>

<!ELEMENT XMLAttributeDefinition EMPTY>
<!ATTLIST XMLAttributeDefinition name CDATA #REQUIRED>
<!ATTLIST XMLAttributeDefinition locator CDATA #REQUIRED>
<!ATTLIST XMLAttributeDefinition ignore (YES|NO) NO>
<!ATTLIST XMLAttributeDefinition priority CDATA #IMPLIED>
<!ATTLIST XMLAttributeDefinition type NUMBER #REQUIRED>
```

---

## Semantics of locator (XPath) expressions

According to the XML data model, XML documents are viewed as trees containing these kinds of nodes:

- The root node
- Element nodes
- Text nodes
- Attribute nodes
- Namespace nodes
- Processing instruction nodes
- Comment nodes

The links between those nodes, in other words the tree-forming relationship, reflect the immediate containment relationship in the XML document.

The **root node** can appear only at the root and nowhere else in the tree. It contains, as its children, the document element and optional comments and processing instructions.

**Element nodes** can contain any kinds of nodes except for the root node. The other kinds of nodes are only allowed as leaf nodes of the tree.

There are three kinds of **containment links**: 'child', 'attribute', and 'namespace'. The 'attribute' and 'namespace' containment links must lead to attribute and namespace nodes, respectively. In other words, to access the children of an element node (in terms of graph theory) you need to follow 'attribute' links to find all contained attributes, follow 'namespace' links to find all contained namespace declarations, and follow 'child' links to find contained elements, text nodes, processing instructions, and comments.

An XPath expression needs to be interpreted with respect to a context node, and denotes a set of nodes. When used as Net Search Extender selector patterns, the context node is free, that is, a relative path pattern *p* is interpreted as *//p*.

These are the Net Search Extender XPath selector patterns:

- Pattern '| ' LocationPathPattern in context N denotes the union of the nodes matched by Pattern and LocationPathPattern, both in context N.
- '/'RelativePathPattern in context N denotes whatever this RelativePathPattern denotes in the context of the root.
- '///'RelativePathPattern in context N denotes the union of the denotations of this RelativePathPattern interpreted in any context that is a descendant (on the child axis) of the root.
- RelativePathPattern '/' StepPattern matches a node in context N, if and only if that node is matched by StepPattern in the context of its parent, and its parent node is matched by RelativePathPattern in context N.
- RelativePathPattern '///' StepPattern matches a node in context N, if and only if that node is matched by StepPattern in the context of its parent, and it has an ancestor node that is matched by RelativePathPattern in context N.
- 'child':::NodeTest (abbreviated syntax: NodeTest) in context N matches a node that is a child of N (on the child axis) and that satisfies NodeTest.
- 'attribute':::NodeTest (abbreviated syntax: @NodeTest) in context N matches a node that is an attribute of N and that satisfies NodeTest.

- NodeType '(' ' )' is satisfied for a node if and only if it is of the specified type.
- 'processing-instruction' '(' Literal ' )' is satisfied for any processing-instruction-type node that has Literal as its name.
- '\*' is satisfied for any element or attribute node (name mask for element name).
- NCName ':' '\*' is satisfied for any element node that has NCName as its name prefix.
- QName is satisfied for any node with the specified name.

## Note

A NodeTest of the form NameTest assumes the node to be of the principal type on the selected axis, which is attribute type on the attribute axis and child type on the child axis. Consequently, NameTest cannot be used to choose comments or processing instruction nodes, but only child and attribute nodes. Moreover, the patterns allow for the selection of any kind of node, except for namespace nodes, because the axis specifier 'namespace' is not allowed.

Examples of patterns:

- chapter | appendix denotes all chapter elements and appendix elements
- table denotes all table elements
- \* denotes all elements (note that this is the abbreviation of child::\*)
- ulist/item denotes all item elements that have a ulist parent
- appendix//subsection denotes all subsection elements with an appendix ancestor
- / denotes the singleton set containing just the root node
- comment() denotes all comment nodes
- processing-instruction() denotes all processing instructions
- attribute::\* (or @\*) denotes all attribute nodes

This is the syntax of the locator element:

```
Locator ::= LocationPathPattern
         | Locator '|' LocationPathPattern
LocationPathPattern ::= '/' RelativePathPattern ?
                    | '//'? RelativePathPattern
RelativePathPattern ::= StepPattern
                    | RelativePathPattern '/' StepPattern
                    | RelativePathPattern '//'? StepPattern
StepPattern ::= ChildOrAttributeAxisSpecifier NodeTest
ChildOrAttributeAxisSpecifier ::=
    ('child' | 'attribute') '::'
    | '@'?
NodeTest ::= NameTest
          | NodeType '(' ' )'
          | 'processing-instruction' '(' Literal ' )'
NameTest ::= '*' | NCName ':' '*' | QName
NodeType ::= 'comment' | 'processing-instruction'
```

NCName and QName are as defined in the XML Names Recommendation:

### NCName

An XML name containing no colons

### QName

A NCName that can be preceded by a NCName followed by a colon. For example: NCName:NCName

---

## Limitations for text fields and document attributes

The following lists the limitations for text fields and document attributes:

- Maximum number of fields in an index: 32767
- Maximum number of values for one attribute of type STRING in one document: 1024
- Maximum number of attributes of type STRING: 253
- Number of characters in a STRING attribute value is truncated to 128
- Maximum number of attributes of types DATE and NUMBER: 32766
- Number of characters in a DATE or NUMBER attribute value is truncated to 128
- For NUMBER attributes, a double precision floating point number is accepted as a value.
- Maximum number of values that can be specified for one attribute of type DATE or NUMBER in one document: unlimited

These are the tags that can be included in an HTML document model:

- <A>
- <ADDRESS>
- <AU>
- <AUTHOR>
- <H1>
- <H2>, <H3>, <H4>, <H5>
- <H6>
- <TITLE>

Tags like <HEAD> and <BODY> that can contain other tags, cannot be specified in an HTML document model as a text field.

---

## Outside In tag attribute values

Possible values for the tag attribute relating to Outside In document property tag types:

SCCCA\_ABSTRACT  
SCCCA\_ACCOUNT  
SCCCA\_ADDRESS  
SCCCA\_ATTACHMENTS  
SCCCA\_AUTHORIZATION  
SCCCA\_BACKUPDATE  
SCCCA\_BASEFILELOCATION  
SCCCA\_BILLTO  
SCCCA\_BLINDCOPY  
SCCCA\_CARBONCOPY  
SCCCA\_CATEGORY  
SCCCA\_CHECKEDBY  
SCCCA\_CLIENT  
SCCCA\_COMPANY  
SCCCA\_COMPLETEDDATE  
SCCCA\_COUNTCHARS  
SCCCA\_COUNTPAGES  
SCCCA\_COUNTWORDS  
SCCCA\_CREATIONDATE  
SCCCA\_DEPARTMENT  
SCCCA\_DESTINATION  
SCCCA\_DISPOSITION  
SCCCA\_DIVISION

SCCCA\_DOCCOMMENT  
 SCCCA\_DOCTYPE  
 SCCCA\_EDITMINUTES  
 SCCCA\_EDITOR  
 SCCCA\_FORWARDTO  
 SCCCA\_GROUP  
 SCCCA\_KEYWORD  
 SCCCA\_LANGUAGE  
 SCCCA\_LASTPRINTDATE  
 SCCCA\_LASTSAVEDBY  
 SCCCA\_MAILSTOP  
 SCCCA\_MANAGERSCCCA\_MATTER  
 SCCCA\_OFFICE  
 SCCCA\_OPERATOR  
 SCCCA\_OWNER  
 SCCCA\_PRIMARYAUTHOR  
 SCCCA\_PROJECT  
 SCCCA\_PUBLISHER  
 SCCCA\_PURPOSE  
 SCCCA\_RECEIVEDFROM  
 SCCCA\_RECORDEDBY  
 SCCCA\_RECORDEDDATE  
 SCCCA\_REFERENCE  
 SCCCA\_REVISIONDATE  
 SCCCA\_REVISIONNOTES  
 SCCCA\_REVISIONNUMBER  
 SCCCA\_SECONDARYAUTHOR  
 SCCCA\_SECTION  
 SCCCA\_SECURITY  
 SCCCA\_SOURCE  
 SCCCA\_STATUS  
 SCCCA\_SUBJECT  
 SCCCA\_TITLE  
 SCCCA\_TYPIST  
 SCCCA\_USERDEFINEDPROP  
 SCCCA\_VERSIONDATE  
 SCCCA\_VERSIONNOTES  
 SCCCA\_VERSIONNUMBER

Possible values for the tag attribute relating to Outside In begin and end tag subtypes:

SCCCA\_ALTFONTDATA  
 SCCCA\_ANNOTATIONREFERENCE  
 SCCCA\_CAPTIONTEXT  
 SCCCA\_CHARACTER  
 SCCCA\_COMPILEDFIELD  
 SCCCA\_COUNTERFORMAT  
 SCCCA\_CUSTOMDATAFORMAT  
 SCCCA\_DATEDEFINITION  
 SCCCA\_DOCUMENTPROPERTYNAME  
 SCCCA\_ENDNOTEREFERENCE  
 SCCCA\_FONTANDGLYPHDATA  
 SCCCA\_FOOTNOTEREFERENCE  
 SCCCA\_FRAME  
 SCCCA\_GENERATEDFIELD  
 SCCCA\_GENERATOR  
 SCCCA\_HYPERLINK  
 SCCCA\_INDEX  
 SCCCA\_INDEXENTRY  
 SCCCA\_INLINEDATAFORMAT  
 SCCCA\_LISTENTRY  
 SCCCA\_MERGEENTRY  
 SCCCA\_NAMEDCELLRANGE  
 SCCCA\_REFERENCEDTEXT  
 SCCCA\_STYLE  
 SCCCA\_SUBDOCTEXT

SCCCA\_TOA  
SCCCA\_TOAENTRY  
SCCCA\_TOC  
SCCCA\_TOCENTRY  
SCCCA\_TOF  
SCCCA\_VECTORSVETAG  
SCCCA\_XREF

Note that the tables include any document property, as well as all the tag subtypes recognized by the INSO filters. There are two subtype exceptions: SCCCA\_DOCUMENTPROPERTY and SCCCA\_BOOKMARK.

---

## Part 9. Reference





---

## Chapter 47. Administration commands for the instance owner

This section describes the syntax of administration commands for the instance owner. Instance owner administration consists of checking the status of Net Search Extender locking and update services, and starting and stopping these services.

The commands are subcommands of the db2text command and allow for the administration of Net Search Extender Services that are specific to a DB2 instance.

Command	Purpose
"CONTROL command"	Lists and deletes full-text index locks. Also lists the cache states.
"START command" on page 187	Starts the Net Search Extender instance services.
"STOP command" on page 188	Stops the Net Search Extender instance services.
DB2EXTHL command	Changes the maximum size of the input parameter of the highlight UDF.

---

### CONTROL command

#### Purpose

This command lets you list and delete full-text index locks managed by the Net Search Extender Instance Services. If the locking and update services are running, you can view their status as well as information on the activated cache.

In a partitioned database environment this only affects the current partition. The user is responsible for invoking the DB2 command db2\_a11 for the desired partitions.


#### Authorization

To issue the command successfully, the user must be the DB2 instance owner with DBADM and DATAACCESS authority.

#### Required connection

This command must be issued from the DB2 database server.

#### Command syntax



#### set-of-locks:



### index-specification:

|—|database-specification|—INDEX—|index-schema-". "—|index-name—|

### Database-specification:

|—DATABASE—*database-name*—|

## Command parameters

### CLEAR

Use CLEAR to force a cleanup for a set of locks. Use this command carefully after checking what leads to a locking problem.

Do not use the CLEAR command if any index administration commands, like an index update, are still active on the index you apply it to, because this could corrupt your index, requiring a complete rebuild.

**LIST** Use LIST to get information about the current locks held for a specific index or database. If there is an update lock, the command also prints information about the number of documents that have been processed so far.

Note this is only for the time that the lock is holding the index.

When using a replication capture table, there are no update operations. Instead, insert operations can be either from an insert or an update operation on the source table the index was created on.

### set-of-locks

Works with locks only in the specified database or index.

### SHOW CACHE STATUS FOR

Shows the activation status for a cached table of the specified index. This can be either: "Not Activated" or "Currently Activated". If the cache is activated, it displays details about cache memory usage. For example, the maximum cache size (in megabytes), the maximum number of documents to insert, and the space left in the cache table (in kilobytes).

### STATUS

By using the STATUS keyword, the command displays whether the locking and update Net Search Extender Instance Services are up and running.

### DATABASE *database-name*

The name of the database on the server that is being used.

### INDEX *index-schema.index-name*

The schema and name of the text index that is currently being used. This is specified in the CREATE INDEX command.

## Usage notes

When an administration command error message indicates that there is a locking problem, ensure that no conflicting task is running. For example, attempting an ALTER command while an UPDATE command is running.

Use SHOW CACHE STATUS FOR for an incremental index update to check that the specified memory size is still large enough to hold all the update information

during the next update, or to check if an activation has been done.

---

## START command

### Purpose

This command starts a demon that controls the locking of full-text indexes and the automatic updating of full-text indexes on the DB2 server.

**Note:** As the command does not activate any temporary cached table for indexes, individual ACTIVATE CACHE commands are necessary before you can start searching with a stored procedure.

### Authorization

The instance owner must hold DBADM with DATAACCESS authority for the current DB2 instance.

### Required connection

This command must be issued from the DB2 database server.

### Command syntax

▶—START—▶

### Command parameters

None

### Usage notes

On Windows, the command starts a service DB2EXT - *<instance\_name>* [- *<nodenum>*]. You can also start the service(s) using the Control Panel or the NET START command. However, you cannot start Net Search Extender through a Terminal Service Client.

For DB2 instances used with partitioned databases, it is highly recommended that the Net Search Extender Instances Services be started using `db2text start` instead of using the normal Windows methods.

Using the services management console you can manually start or stop each of the DB2EXT services for a DB2 instance. However, to keep Net Search Extender in a proper running state it is necessary to start all DB2EXT services and shutdown all DB2EXT services associated with a DB2 instance. Also, during manual start and stop it is necessary that the following sequence must be followed:

#### Starting

Start the NSE (DB2EXT) service beginning with the lowest-numbered partition on a host and working your way up to the highest-numbered partition, before starting other DB2EXT services

#### Stopping

Stop the NSE (DB2EXT) service beginning with the highest-numbered partition on a host and work your way down to the lowest-numbered partition.

However when you use `db2text start` and `db2text stop`, this required sequence is automatically followed by the system.

Startup type "Automatic" is not supported for the DB2EXT services. The DB2EXT services have to be started manually by executing `db2text start` or by executing `net start` for each of the DB2EXT services for the DB2 instance. Do not attempt to run text search operations when the system is not fully started.

If `START` fails, there might still be obsolete entries in the scheduler referring to indexes that no longer exist. Edit the file `../sql11ib/db2ext/ctedem.dat` and remove any obsolete entries. Rerun the `START` command.

After successfully starting Net Search Extender, the process `ctelock` (`ctelock.exe` on Windows) is active on your system. Several shared resources (shared memory and semaphores) are created and stored in the `/tmp` directory on UNIX machines. These files are required by Net Search Extender and must not be deleted as long as the instance is running. However, if after a `STOP`, the `START` command is not successful, make sure that all of the old anchor files in the `/tmp` directory have been removed. The `/tmp` directory must be read, write and execute accessible by all.

---

## STOP command

### Purpose

This command stops the locking and update services of Net Search Extender.

### Authorization

The instance owner must hold `DBADM` with `DATAACCESS` authority for the current DB2 instance.

### Required connection

This command must be issued from the DB2 database server.

### Command syntax

```
➤ STOP [FORCE]
```

### Command parameters

#### FORCE

Stops services even if processes are holding locks or if the cached table is activated for any index. If you do not specify `FORCE`, the command will fail in these cases with a warning about active caches.

### Usage notes

Stopping the Net Search Extender instance services disables any further use of specific Net Search Extender commands. When restarting the services, you must activate the temporary cache again if you previously used an activated cache with your index.

For DB2 instances used with partitioned databases, it is highly recommended that the Net Search Extender Instances Services be stopped using `db2text stop` instead of using the normal Windows methods.

Using the services management console you can manually stop each of the DB2EXT services for a DB2 instance. However, to keep NSE in a proper running state it is necessary to shut down all DB2EXT services associated with a DB2 instance. Also, during a manual stop it is necessary that you stop the NSE (DB2EXT) service beginning with the highest-numbered partition on the host and working your way down to the lowest-numbered partition. However when you use `db2text stop`, this required sequence is automatically followed by the system.

After successfully stopping Net Search Extender, the process `ctelock` (`ctelock.exe` on Windows) is terminated. The shared resources and the anchor files in the `/tmp` directory on UNIX machines are deleted.



---

## Chapter 48. Administration commands for the database administrator

This section describes the syntax of administration commands for the database administrator. Database administration consists of setting up databases for use by Net Search Extender and then disabling this setup.

Only the ENABLE DATABASE and DISABLE DATABASE commands are a variation of the DB2TEXT command, although all these commands allow for administration on the database level.

Command	Purpose
"ENABLE DATABASE command"	Enables the current database to create full-text indexes.
"DISABLE DATABASE command" on page 193	Resets preparation work completed by Net Search Extender for a database.
"DB2EXTHL command" on page 194	The DB2EXTHL command changes the maximum size of the input parameter of the highlight UDF.

### Tip

If no database connection information has been specified as part of the db2text command, the db2text executable causes an implicit connection to be made to the default subsystem specified in the environment variable DB2DBDFT.

---

## ENABLE DATABASE command

The ENABLE DATABASE command enables a database to create and use full-text indexes on text columns. It creates the Net Search Extender infrastructure in the database, such as administration tables, views, user-defined functions (UDFs), and stored procedures for searching.

### Authorization

A user must have DBADM privilege to execute the ENABLE DATABASE command.


### Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the ENABLE DATABASE command.

### Command syntax

```
▶▶—ENABLE DATABASE FOR TEXT—▶▶
    └─┬─ AUTOGRANT ─┬─ connection-options ─┬─▶▶
```

## connection-options:



## Command parameters

### CONNECT TO *database-name*

The name of the database that is a target for this command. You can omit this parameter, if the environment variable DB2DBDFT is set and the user is running the command under a user ID with the necessary DB2 authorizations.

### USER *userid* USING *password*

Use a *userid* and *password* to connect to the database.

### AUTOGRANT

If this option is specified, an attempt is made to grant DBADM with DATAACCESS privileges to the instance owner, in case the instance owner misses these privileges for this database. For a successful grant of privileges, the user must have SECADM privilege for the database and cannot be the instance owner (a user cannot grant privileges to herself/himself).

**Note:** This option is not supported in the stored-procedure interface.

## Usage notes

This command prepares the connected database for use by Net Search Extender. It is a mandatory step before you can create a Net Search Extender index on tables/columns in the database.

You can view the database defaults established after running the command by using the DB2EXT.DBDEFAULTS catalog view.

### Changes to the database

This command grants DBADM authority to the DB2 instance owner associated with the DB2 instance of the enabled database.

The ENABLE DATABASE command creates various database objects in the schema DB2EXT, such as Net Search Extender catalogs, UDFs, and stored procedures. After running the command, the following catalog views are available:

```
db2ext.dbdefaults
db2ext.textindexes
db2ext.textindexformats
db2ext.indexconfiguration
```

These tables are located in the default table space of the database, known as IBMDEFAULTGROUP. This table space is distributed over all the nodes defined in db2nodes.cfg

### Changes to the file system

None.



---

## DISABLE DATABASE command

### Purpose

This command undoes Net Search Extender changes to a database.

### Authorization

A user must have DBADM privilege to execute the DISABLE DATABASE command.

### Prerequisite

Instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. SECADM must explicitly grant DBADM with DATAACCESS to instance owner before running the DISABLE DATABASE command.

### Command syntax

```
➤➤—DISABLE DATABASE FOR TEXT—┬──┴──┬──────────────────────────────────────────┬──┘
```

└──FORCE──┘   └──|connection-options|──┘

#### connection-options:

```
└──┬──────────────────────────────────────────┬──┘
```

└──CONNECT TO—*database-name*──┬──┴──┬──────────┬──┘

└──USER—*userid*—USING—*password*──┘

### Command parameters

**CONNECT TO** *database-name*

The name of the database that is a target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command under a user ID with the necessary DB2 authorizations.

**USER** *userid* **USING** *password*

Use a *userid* and *password* to connect to the database.

**FORCE**

Forces the dropping of all Net Search Extender indexes in the database.

### Usage notes

This command resets the connected database, so that it can no longer be used by other Net Search Extender commands. If full-text indexes exist in the database, this command fails unless the FORCE option is used.

This command does not remove DBADM authority from the DB2 instance owner.

**Note:** Disabling a database will fail if there are any text indexes defined in the database. It is recommended to remove these indexes one by one and then check if any problems occur. If you use the disable database for text force command, it only guarantees that Net Search Extender catalog tables in the database are removed.

However, if some of the indexes can not be completely dropped, there may still be resources that need to be manually cleaned up. These include:

- Files in the index, work and cache directory
- Scheduler entries in ctedem.dat
- Where an index was created using the replication capture option, the IBMSNAP\_SIGNAL, IBMSNAP\_PRUNE\_SET, and IBMSNAP\_PRUNCNTL entries in the tables of the remote database must be manually deleted. These entries can be easily identified using APPLY\_QUAL="NSE" || <instance name> and TARGET\_SERVER= <database name> command.

In the following example, the instance is DB2 and the database is SAMPLE.

```
DELETE FROM <ccSchema>.IBMSNAP_SIGNAL
WHERE SIGNAL_INPUT_IN IN
      (SELECT MAP_ID FROM <ccSchema>.IBMSNAP_PRUNCNTL
       WHERE APPLY_QUAL= 'NSEDB2' AND TARGET_SERVER= 'SAMPLE');
```

```
DELETE FROM <ccSchema>.IBMSNAP_PRUNCNTL
WHERE APPLY_QUAL= 'NSEDB2' AND TARGET_SERVER= 'SAMPLE';
```

```
DELETE FROM <ccSchema>.IBMSNAP_PRUNE_SET
WHERE APPLY_QUAL= 'NSEDB2' AND TARGET_SERVER= 'SAMPLE';
```

#### **Changes to the database**

The following modifications made in the database to enable Net Search Extender are deleted:

- The Net Search Extender catalog views in the database.
- All the database objects created by Net Search Extender.

#### **Changes to the file system and shared memory**

If you use the FORCE option, the index files are deleted.

If you use the FORCE option, the cache is deleted for any activated cache of indexes.

---

## **DB2EXTHL command**

### **Purpose**

The DB2EXTHL command changes the maximum size of the input parameter of the highlight UDF.

By default, the highlight UDF takes as input a document up to a maximum size of 100 KB and returns a 200 KB CLOB. Depending on the size of the largest document in the database, you can increase the input value to a maximum size of 1 GB.

### **Authorization**

To issue this command successfully, the user must be a DB2 instance owner with DBADM and DATAACCESS authority.

### **Required connection**

This command must be issued from the DB2 database server and requires a DB2DBDFT environment variable.

## Command syntax

► `db2exthl` *new-highlight-input-size* ◄

## Command parameters

*new-highlight-input-size*

The new result size of the highlight UDF in kilobytes. This is a positive integer < 1048576.



---

## Chapter 49. Administration commands for the text table owner

This section describes the syntax of administration commands for the text table owner.

The commands are subcommands of the DB2TEXT command. These allow the owner of a table to create and manipulate full-text indexes on columns of the table.

Command	Purpose
"ACTIVATE CACHE command"	Activates the cache so that search operations using the stored procedure are possible
"ALTER INDEX command" on page 199	Changes the characteristics of an index
"CLEAR EVENTS command" on page 203	Deletes index events from an index event table used during index update
"CREATE INDEX command" on page 204	Creates a full-text index
"DEACTIVATE CACHE command" on page 217	Deactivates the cache so that search operations using the stored procedure are no longer possible
"DB2EXTTH command" on page 220	Compiles the thesaurus definition file
"DROP INDEX command" on page 219	Drops a full-text index for a text column
"RESET PENDING command" on page 221	Refreshes text-maintained dependent tables for base tables with the extended staging infrastructure after executing a SET INTEGRITY command.
"UPDATE INDEX command" on page 222	Starts the indexing process based on the current contents of the text columns
"HELP command" on page 225	Displays the list of DB2TEXT command options
"COPYRIGHT command" on page 226	Displays the Net Search Extender product and copyright information

### Tip:

If no database connection information has been specified as part of the db2text command, the db2text executable causes an implicit connection to be made to the default subsystem specified in the environment variable DB2DBDFT.

---

## ACTIVATE CACHE command

### Purpose

This command activates the cached table from either the DB2 user table or the persistent cache files. After completion, search operations using the stored procedure are possible.

This command is only available if the index was created with a CACHE TABLE option.

## Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table or nickname on which the text index is defined
- DBADM authority

## Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the ACTIVATE CACHE command.

## Command syntax

```
▶▶ ACTIVATE CACHE FOR INDEX index-schema-".." index-name FOR TEXT ▶▶
|-----|
| RECREATE | | connection-options |
```

### connection-options:

```
|-----|
| CONNECT TO database-name |
| USER userid USING password |
```

## Command parameters

### *index-schema*

The schema of the text index, as specified in the CREATE INDEX command. If no schema is specified, the user ID of the DB2 connection is used.

### *index-name*

The name of the text index, as specified in the CREATE INDEX command.

### RECREATE

Applies only to indexes using a persistent cache; an existing cache is deleted. If an update without activation has completed, the persistent cache is automatically reconstructed from the database.

### CONNECT TO *database-name*

The name of the database that is target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

### USER *userid* USING *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

## Usage notes

You cannot issue the command if one of the following commands is currently running on the index:

- UPDATE INDEX
- ALTER INDEX
- DROP INDEX
- CLEAR EVENTS
- DEACTIVATE CACHE

**Note:** Activation of a cached table may require its recreation from scratch, even though a persistent cache was used. This occurs if an update operation was performed while the persistent cache was deactivated.

The amount of memory used to build the cache is dynamically calculated from the current number of documents and the size of the result columns. Use the PCTFREE value to increase the calculated minimal amount of memory by a factor of 100/(100-PCTFREE). The PCTFREE value is specified in the CREATE INDEX or ALTER INDEX command.

Thereby, PCTFREE describes the percentage of the allocated cache that is reserved for insert operations while the cache is activated. Note that for each ACTIVATE CACHE command, the actual memory size is re-evaluated.

#### Changes to the file system

Files for implementing the persistent cache are created.

## ALTER INDEX command

### Purpose

The command changes the characteristics of a full-text index, for example, the update options and the storage options.

### Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:


- CONTROL privilege on the table or nickname on which the text index is defined
- DBADM authority

### Prerequisite


The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the ALTER INDEX command.

### Command syntax


```
▶▶ ALTER INDEX [index-schema-"."] index-name FOR TEXT →
└─┬─┘ └─┬─┘
└─┬─┘ └─┬─┘
|update-characteristics| |storage-options|
```




**storage-options:**




**update-characteristics:**



**update-frequency:**



**connection-options:**



**Command parameters**

*index-schema*

The schema of the text index as specified in the CREATE INDEX command. If no schema is specified, the user ID of the DB2 connection is used.

*index-name*

The name of the text index as specified in the CREATE INDEX command.

**INDEX DIRECTORY** *directory*

The directory path where the text index is stored. As the directory will contain index data, ensure that the directory has read, write, and execute permissions for the DB2 instance owner user ID.



Note that in a partitioned database environment, this directory has to exist on every partition. A subdirectory, `NODE<nr>`, is created under the directory to distinguish indexes on logical partitions of a server. Any index files from the previous index directory are deleted.

**WORK DIRECTORY** *workdirectory*

Stores temporary files during search and administration operations. You can change the separate work directory independently of a new index directory.

If the directory does not exist, it is created using the DB2 instance owner user ID. If it exists, ensure that the directory has read, write and execute permissions on UNIX platforms for the instance owner.

Note that in a partitioned database environment, this directory has to exist on every partition. A subdirectory, `NODE<nr>`, is created under the directory to distinguish indexes on logical partitions of a server. Any temporary index files from the previous index directory are deleted.

**CACHE TABLE PERSISTENT IN** *directory*

Specifies that the cached table in `CREATE INDEX` is persistent even after a deactivation or system reboot. In either case, this allows for a fast `ACTIVATE CACHE` execution. The persistent cache is stored in the specified directory.

The previously created persistent cache is moved to a new location. This operation always requires a deactivated index.

**CACHE TABLE TEMPORARY**

Specifies that the cached result table is now temporary and any previously existing persistent cache has been deleted. Note that this change operation requires a deactivated index.

**MAXIMUM CACHE SIZE** *memsize*

Specifies the new maximum size of the cached table to be built during `ACTIVATE CACHE`. Specify the `memsize` parameter in megabytes as a positive integer.

If *memsize* is too small, the `ACTIVATE CACHE` command fails. The actual cache size is calculated during the `ACTIVATE CACHE` command. This change requires a deactivated index.

**PCTFREE** *percentage*

Specifies the percentage of the cache held free for additional documents. The percentage must be an integer value less than 100 and greater or equal to 0. Note that the previous persistent cache is deleted and that this change requires a deactivated index.

**UPDATE FREQUENCY**

Using the following parameters, the index update frequency determines when the update occurs:

- **D.** The day(s) of the week when the index is updated: \* (everyday) or 0..6 (0=Sunday)
- **H.** The hour(s) when the index is updated: \* (every hour) or 0..23
- **M.** The minute(s) when the index is updated: 0..59
- **NONE.** No further index updates occur. This is intended for a text column in which no further changes are made, or where only manual update are run in the future.

If you do not specify the UPDATE FREQUENCY keyword, the frequency settings are left unchanged.

**UPDATE MINIMUM** *minchanges*

The minimum number of changes allowed for text documents before the index is incrementally updated. If you do not specify the UPDATE MINIMUM keyword, the setting does not change.

Note that you can only change the UPDATE MINIMUM if you did not create the index using the RECREATE ON UPDATE option.

**COMMITCOUNT FOR UPDATE** *count*

For update processing, you can specify a commitcount. This applies to both the UPDATE command and the UPDATE FREQUENCY specification, which schedules update processing.

The COMMITCOUNT value is ignored during initial update.

Note that you can only change COMMITCOUNT if you did not create the index using the RECREATE ON UPDATE option.

Also note that you cannot change COMMITCOUNT if you did create the index with the REPLICATION clause.

**CONNECT TO** *database-name*

The name of the database that is target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

**USER** *userid* **USING** *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

## Usage notes

You cannot issue the alter index command if one of the following commands is running on the index:

- ALTER INDEX
- CLEAR EVENTS
- ACTIVATE CACHE
- DROP INDEX
- UPDATE INDEX
- DEACTIVATE CACHE

If you create the index with a cache option, you can not use the ALTER INDEX command for the index directory when the index is activated. You must first deactivate the cache.

In a partitioned database environment, a text index with cache options is only allowed on a single-partition tablespace.

### Changes to the database

Change Net Search Extender catalog views.

### Changes to the file system

- Creation of NODE<nr> subdirectories in the index, and work directories
- Moving of index files
- Creation of persistent cache directories

- Moving of persistent cache files

---

## CLEAR EVENTS command

### Purpose

This command deletes indexing events from an index's event view. The name of the event view is found in the EVENTVIEWNAME column of the DB2EXT.TEXTINDEXES view.

### Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table or nickname on which the text index is defined
- DBADM authority

### Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the CLEAR EVENTS command.

### Command syntax

```

▶▶ CLEAR EVENTS FOR INDEX index-schema-".." index-name FOR TEXT
▶
┌ COMMITCOUNT count ─┘ ┌ |connection-options| ─┘

```

#### connection-options:

```

┌ CONNECT TO database-name ─┘ ┌ USER userid USING password ─┘

```

### Command parameters

#### *index-schema*

The schema of the text index as specified in the CREATE INDEX command. If no schema is specified, the user ID of the DB2 connection is used.

#### *index-name*

The name of the text index as specified in the CREATE INDEX command.

#### COMMITCOUNT *count*

An INTEGER value  $\geq 0$  displays the number of rows deleted in one transaction by DB2.

#### CONNECT TO *database-name*

The name of the database that is target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

**USER** *userid* **USING** *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

## Usage notes

When you schedule regular updates using the UPDATE FREQUENCY option in the CREATE INDEX or ALTER INDEX commands, regularly check the event table. Use CLEAR EVENTS to clean up the event tables, after you have checked the reason for any events indicating an error and removed the source of the error that is mentioned in the event table.

You cannot issue the clear events command if one of the following commands is running on the index:

- UPDATE INDEX
- ALTER INDEX
- ACTIVATE CACHE
- DEACTIVATE CACHE
- DROP INDEX

---

## CREATE INDEX command

This command creates a full-text index on a text column for use in Net Search Extender full-text queries.

### Purpose

In a partitioned database environment, a full-text index is created on every partition of the table space the user table is defined on. Subsequent changes to the distribution of the table space are not allowed and will lead to unexpected behavior in the administration commands and during the search process.

### Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:

One of:

- DBADM authority
- CONTROL privilege on the table or nickname on which the text index is defined
- INDEX privilege on the table or nickname on which the text index is defined and one of following:
  - IMPLICIT\_SCHEMA authority on the database, if the implicit or explicit schema name of the index does not exist
  - CREATEIN privilege on the schema, if the schema name of the index refers to an existing schema

### Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the CREATE INDEX command.

## Command syntax

► CREATE INDEX `[index-schema-"."]` index-name FOR TEXT ON

► `[table-schema-"."]` table-name

► (text-column-name) `[function-schema "."]` function-name-(-text-column-name-)

► `[attribute-list]` `[text-default-information]`

► `[update-characteristics]` `[storage-options]`

► `[cache-search-result-options]` `[index-configuration-options]`

► `[connection-options]`

### attribute list:

► ATTRIBUTES ( SQL-column-expression `[AS-attribute-name]` )

### text-default-information:

► `[CCSID-ccsid]` `[LANGUAGE-language]`

► `[FORMAT-format]` `[model-information]`

### model-information:

► DOCUMENTMODEL `documentmodel-name` IN model filepath


► `[USING CCSID-ccsid]`

### update-characteristics:


► `[UPDATE FREQUENCY-NONE]` `[update-frequency]`

► `[incremental-update-characteristics]` `[RECREATE INDEX ON UPDATE]`


**incremental-update-characteristics:**




**capture-table-characteristics:**




**update-frequency:**



**storage-options:**



**cache-search-results-options:**



▶ | INITIAL SEARCH RESULT ORDER—(—SQL-order-by-list—) |

▶ | KEY COLUMNS FOR INDEX ON VIEW—(SQL-columnname-list)— |

### index-configuration-options:

| INDEX CONFIGURATION—(—option-value—) |

### connection-options:

| CONNECT TO—database-name— |  
| USER—userid—USING—password— |

## Command parameters

### *index-schema*

The schema of the text index. Use this as a DB2 schema name for the index-specific administration tables. If no schema is specified, the user ID of the DB2 connection is used. Note that the index schema must be a valid DB2 schema name.

### *index-name*

The name of the index. Together with the index schema, this uniquely identifies a full-text index in a database.

Note that the index name must be a valid DB2 index name.

### *table-schema*

The schema for which of the table, nickname, or view the index is created. If no schema is specified, the user ID of the DB2 connection is used.

### *table-name*

The name of the text table, nickname, or view in the connected database that contains the column the full-text index is created for.

Note that when the table name does not refer to a DB2 base table, there are the following restrictions:

- A view only allows a stored procedure or table-valued function search. Therefore, you must specify the key columns for the index or views using the KEY COLUMNS FOR INDEX ON VIEW clause.
- For incremental index updates on nicknames without capture tables, a log table is created. If any changes occur to the data in the nickname table or view, you must manually fill in the log table. With base tables this is done automatically, so the user **must not** touch the log table.
- The DB2 predicates CONTAINS, SCORE, and NUMBEROFMATCHES are only allowed for indexes on base tables or nicknames, but not on views.
- Indexes on views are only allowed if you specify cache-search-result options in the command.

- The extended text-maintained staging infrastructure complementing the trigger-based log to support incremental updates can be applied to partitioned and non-partitioned base tables, however, it is not supported for views or nicknames.

*text-column-name*

The name of the column containing the text used for creating the full-text index. The column must have one of the following data types:

- CHAR (FOR BIT DATA)
- VARCHAR (FOR BIT DATA)
- LONG VARCHAR (FOR BIT DATA)
- CLOB
- DBCLOB
- BLOB
- GRAPHIC
- VARGRAPHIC
- LONG VARGRAPHIC
- XML

If the column data type is none of these, specify a transformation function using *function-schema.function-name* to convert the column type to a supported one.

Note that several indexes on the same columns are allowed, but **only** under either one of the following conditions:

**The index is created on a view**

Therefore, you can not use the index in the CONTAINS, SCORE, or NUMBEROFMATCHES search arguments.

**The index is created on a table**

If all the indexes are synchronized, they have identical properties on the same column in the following CREATE INDEX command details:

- Function name and schema
- ATTRIBUTES
- CCSID
- LANGUAGE
- FORMAT
- DOCUMENTMODEL
- INDEX CONFIGURATION

Therefore, it does not matter which index is chosen by the CONTAINS, SCORE, or NUMBEROFMATCHES arguments.

*function-schema.function-name*

The schema and the name of a user-defined function used to access text documents that are in a column of an unsupported type. The function performs a column type conversion, using as input parameter an arbitrary column type. It returns the value as one of the Net Search Extender supported types.

**ATTRIBUTES (SQL-column-expression AS Attribute-name, ...)**

Ensures that the content of a column expression is indexed in addition to the text column. This content can then be searched by the ATTRIBUTES



clause in a search statement. The SQL-column expressions have to be defined using unqualified column names of the table on which the index is created. The only data type allowed is DOUBLE. Cast operators can be used in the column expressions, but implicit casting of DB2 is **not** possible. The attribute-names must follow the rules for attribute-names in document models and must be distinct from all attribute names in the indexes model-definition file.

Determine the attribute names for expressions by using the following rules:

- If explicitly named by the SQL AS clause in the column expression, use the specified name. An example would be:

```
ATTRIBUTES (C1+C2 AS myname)
```

- If a column of the specified table is used without AS, the name of the column is used. For example:

```
ATTRIBUTES (C1)
```

- If an expression is used without AS and it does not refer to a named column, CREATE INDEX reports an error.

For example:

```
ATTRIBUTES (CAST(JULIAN_DAY(date) AS DOUBLE) as day, (price1+price2)/2 as avg_price)
```

Note that attributes that are not enclosed by single quotes are mapped to uppercase, and must be specified in uppercase during search.

#### **CCSID** *ccsid*

The Coded Character Set Identifier is used when indexing text documents. The default value is from the DB2EXT.DBDEFAULTS view where DEFAULTNAME='CCSID'.

Only set a CCSID if the data type of the column is binary.

#### **LANGUAGE** *language*

The language parameter specifies the language of the stop word dictionary that is selected if the index configuration value IndexStopWords is set to 0 (ignore stop words during indexing). This parameter must always be set for Thai (TH\_TH) to enable Thai word breaking, and Turkish to correctly distinguish the use of the dotted and dotless "i".

#### **FORMAT** *format*

The format of text documents in the column, for example, HTML. This information is necessary for indexing documents.

For structured document formats, you can specify information in a document model file. If no document model is specified, the text of the document is indexed using a default document model.

If the format keyword is not specified, the default value is taken from the DB2EXT.DBDEFAULTS view where the DEFAULTNAME='FORMAT'. The initial default set by Net Search Extender is TEXT. For data type XML, the FORMAT XML specifier is mandatory.

#### **DOCUMENTMODEL** *documentmodel-name* **IN** *modelfilepath*

The modelfilepath specifies the location of a model file. The modelfilepath must be a fully qualified path. The model file contains a model definition for the format in the FORMAT clause. It must be readable by the DB2 instance owner. A document model enables you to index and search specific sections of a document. You can define markup tags and section names in a document model. A document model is bound to a document format that supports HTML, XML, or GPP structures. You can only specify one document model in a model file.

Note that because the document model is only read during the CREATE INDEX command, any later changes are not recognized for this index.

Note that in a partitioned database environment, a shared file system must be used to ensure that the model file path is accessible on every node on Linux or UNIX platforms. However, on Windows platform, the model file path for document model files must be set to local path available on each node.

#### **USING CCSID** *ccsid*

Specify a CCSID to interpret the contents of the model file. The default value is from the DB2EXT.DBDEFAULTS view where DEFAULTNAME='MODELCCSID'.

#### **UPDATE FREQUENCY**

The index update frequency determines when the update occurs. If changes to the user table are less than that specified by the UPDATE MINIMUM option, the index is not updated. If you do not specify the UPDATE FREQUENCY, the default NONE is used, so that no index updates are made. This is useful when you expect no further changes to a text column or want to have manual control over the update process.

- **D.** The day(s) of the week when the index is updated: \* (everyday) or 0..6 (0=Sunday)
- **H.** The hour(s) when the index is updated: \* (every hour) or 0..23
- **M.** The minute(s) when the index is updated: 0..59
- **NONE.** No further index updates are made. The update must be started manually.

The default value is from the DB2EXT.DBDEFAULTS view where DEFAULTNAME='UPDATEFREQUENCY'.

If you decide not to use the UPDATE FREQUENCY parameter to schedule automatic index updates, you can use operating system functions, for example crontab, instead.

#### **UPDATE MINIMUM** *minchanges*

The minimum number of changes required to text documents before the index is updated based on the UPDATE FREQUENCY settings. Only positive integer values are allowed. The default value is taken from the DB2EXT.DBDEFAULTS view, where DEFAULTNAME='UPDATEMINIMUM'.

Note that this value is ignored if the DB2TEXT UPDATE command is executed manually. This option cannot be used with the RECREATE INDEX ON UPDATE option, as the number of changes is not available without a log table and triggers for incremental update.

For partitioned database environments, the UPDATE MINIMUM is checked on every partition.

#### **REORGANIZE AUTOMATIC | MANUAL**

Updates performed based on the update frequency settings will only reorganize the index if REORGANIZE AUTOMATIC is specified. This step is completed automatically according to the value of select REORG SUGGESTED from DB2EXT.TEXTINDEXES after the update.

REORGANIZE MANUAL can only be performed with a manual UPDATE command, using the REORGANIZE option.

If the REORGANIZE clause is omitted, the default is taken from the DB2EXT.DBDEFAULTS view, where DEFAULTNAME='AUTOMATICREORG'.

**REPLICATION CAPTURE TABLE** *capture-table-schema.capture-table-name*  
**CONTROL TABLE SCHEMA** *capture-control-schema*

For incremental update processing, the specified replication capture table is taken instead of a log table or text-maintained staging table that is otherwise created for the index. Therefore, schemaname, tablename, and the replication capture table name relate to objects in the local DB2 (federated) database.

The capture-control-schema is the schema name of the replication control tables, for example IBMSNAP\_PRUNE\_SET on the local DB2. The replication control tables must be available as nicknames on the local DB2 system after setting up the replication.

At minimum, there must be nicknames available for the following capture control tables:

- IBMSNAP\_SIGNAL
- IBMSNAP\_PRUNE\_SET
- IBMSNAP\_PRUNCNTL
- IBMSNAP\_REGISTER
- IBMSNAP\_REG\_SYNC (Non-DB2 remote sources only)

As DB2 Replication Center does not automatically guarantee to create local nicknames for a remote capture table and capture control tables, this can be a manual task. The task is similar to creating a nickname for the table that the text index is to be created on.

The column names of primary key columns in the user table nickname and the capture table nickname must match. In addition, the names of the columns IBMSNAP\_OPERATION, IBMSNAP\_COMMITSEQ and IBMSNAP\_INTENTSEQ must not be changed in the capture table nickname.

After index creation, the column names DB2EXT.TEXTINDEXES(LOGVIEWNAME) and DB2EXT.TEXTINDEXES(LOGVIEWSCHEMA) both refer to the local name of the replication capture table.

As Net Search Extender does not require all the functionality of the DB2 Replication Center, the Change Data table (CD) or the Consistent-Change Data (CCD) table must obey following rules:

- Use change capture registration and not the full refresh copying option.
- No horizontal subsetting of capturing changes is allowed. For example, by triggers. See Chapter 6, "Subsetting data in your replication environment" in the *DB2 Replication Guide and Reference, Version 8*.
- Registering changes for a subset of columns is only allowed if the primary key columns, the text column, and all columns involved in the attribute and cache table expressions of the DB2TEXT CREATE INDEX command are included.
- The primary key columns must be included in the capture table. Note that the after-image is sufficient.
- The capture tables must not be condensed. For each primary key there must be one entry with the latest data. However, Net Search Extender requires a full history to be available.

- The table must use the D/I option. This enables updates to primary keys on the source table to be transformed into a pair of inserts/deletes.

For supported remote source versions in DB2 V9.7 see the Technical Support pages.

**Note:**

Ensure that the correct source table name is inserted into the registration table. Depending on the type of remote DBMS, the remote tablename or the local nickname must be used:

- DB2: remote table name (the tablename on the remote server)
- Non-DB2: local nickname (the corresponding nickname in the federated DB2 database)

A user mapping must exist for the local user to access the remote data source via nicknames and the remote user must have control privilege on the tables.

If the DB2 instance owner user ID is different from the local user ID, an additional user mapping for the DB2 instance owner user ID is needed.

The specified base table name must not be a view on a nickname. This is because a view can be over several nicknames and several CD and CCD tables can also be involved. As only one CD or CCD table can be specified in the replication capture clause, a view on nicknames can not be supported. In addition, nicknames on a remote views can not be supported because the primary key is missing.

The CD or CCD table must be a nickname and can not be a view or an alias.

**COMMITCOUNT FOR UPDATE** *count*

For **incremental** update processing a commitcount can be specified. If not specified, a default value is taken from the DB2EXT.DBDEFAULTS view, where DEFAULTNAME='COMMITCOUNT'.

The COMMITCOUNT FOR UPDATE value for the index can be found in DB2EXT.TEXTINDEXES.COMMITCOUNT. This can be changed for each index using the ALTER INDEX command. It also applies to the scheduled update processing according to the UPDATE FREQUENCY specification. A value of 0 means that the update is completed in one transaction, with values >0 specifying the number of documents to process in one transaction. You are recommended not to use a non-zero COMMITCOUNT or, if you must, to set the value high enough so that the number of immediate commits during a single incremental update remains very small. Each commit involves moving index files and cleaning up the index log file which takes a considerable amount of time particularly if performed repeatedly. The COMMITCOUNT applies across all changes for documents as listed in the log table as well as the extended text-maintained staging infrastructure, if configured for the index.

If COMMITCOUNT is not set, then the NUMBER\_DOCS parameter from the db2ext.textindexes is not updated. Therefore, to view the number of documents already processed during the updating process, use the CONTROL LIST command.

## RECREATE INDEX ON UPDATE

This does not allow incremental index updates, but recreates the index when an update operation is performed (by command or scheduled update).

**Note:** If you specify RECREATE INDEX ON UPDATE:

- no triggers are created on the user table,
- no log table is created, and
- you cannot configure the extended text-maintained staging infrastructure for the index.

## INDEX DIRECTORY *directory*

The directory path in which the text index is to be stored. As the directory will contain index data, ensure that the directory has read/write and execute permissions for the DB2 instance owner user ID.

The default value is taken from the DB2EXT.DBDEFAULTS view, where DEFAULTNAME=INDEXDIRECTORY'. A subdirectory, NODE<nr>, is created under the directory to distinguish indexes on logical nodes of a server.

Note that in a partitioned database environment, this directory has to exist on every physical partition.

## WORK DIRECTORY *directory*

A separate work directory may be specified optionally, that will be used to store temporary files during index search and administration operations. The directory must exist and have read/write and execute permissions for the DB2 instance owner user ID.

The default value is taken from the DB2EXT.DBDEFAULTS view, where DEFAULTNAME='WORKDIRECTORY'. A subdirectory, NODE<nr>, is created under the directory to distinguish indexes on logical nodes of a server.

Note that in a partitioned database environment, this directory has to exist on every physical partition.

If a WORK DIRECTORY is not specified, a directory called work is created under INDEX DIRECTORY.

If a WORK DIRECTORY is specified, it is strongly recommended that it is located on the same physical file system as the index directory. Not following this recommendation causes massive index update performance degradation, because files from the work directory have to be copied physically into the index directory instead of being able to rename them.

**Note:** Files created in the INDEX DIRECTORY and WORK DIRECTORY will follow umask restriction set for the instance owner. Make sure that these umask restrictions must allow for group read/write access for the fenced user.

## ADMINISTRATION TABLES IN *tablespace-name*

The name of the regular table space for administration tables created for the index. The table space must exist. If not specified, the table space of the user table is chosen, if the index is created on a base table.

In case of a nickname or a view, a default table space is chosen by DB2.

When creating text indexes on views, nicknames, or text indexes for stored procedure search in a partitioned database environment, the table space for administration tables must be specified on a single node and must be called explicitly on this node.

To ensure that you connect to the correct node, use the DB2NODE environment variable. Note that the ADMINISTRATION TABLES IN clause is mandatory for creating indexes on range partitioned table. Else, the CREATE INDEX command returns an error. See CTE0150E for more information on the error message.

#### **CACHE TABLE** (*SQL-column-expression-list*)

A cached table is built in addition to the index, which consists of the specified column expressions. This cache is used to return the result set via a stored procedure search without joining full-text search results with a DB2 table. Note that a regular DB2 search using the full-text index with the CONTAINS function is always possible.

Define the SQL-column expressions using unqualified column names of the table the index is created on. The allowed SQL-column expression types are all built-in and user-defined distinct types. The column names in the result set are determined using the following rules:

- If explicitly named by the SQL AS clause in the column expression, the specified name is used. For example:  
CACHE TABLE (C1+C2 AS myname)
- If a column of the specified table is used without the AS clause, the name of the column is used. For example:  
CACHE TABLE(C1)
- If an expression is used without AS, and which does not refer to a named column, CREATE INDEX reports an error.
- No duplicate column names are allowed.

CLOB data types are not supported as cache data types. You need to cast these to VARCHARS.

**Note:** Note that if the column names of the result set are not disjunct, the CREATE INDEX command returns an error. Also note that the cached table is not implicitly activated after creation, for example search by stored procedure is not possible until DB2TEXT ACTIVATE CACHE is performed.

This option may be used in a partitioned database environment only if the user table is stored in a single-partition table space. It cannot be used if the default configuration for the text-maintained staging infrastructure is set to ON for the table and the configuration is not manually disabled for the text index.

#### **PERSISTENT IN** *directory*

Specifies that the cache is also created persistently. A persistent cache can be reactivated more rapidly after a deactivation or a system restart than a non-persistent cache. The persistent cache is stored in the specified directory.

Note that if the directory is not specified, the default is taken from the db2ext.dbdefaults view, where DEFAULTNAME='CACHEDIRECTORY'.

#### **TEMPORARY**

Specifies that the cache is not stored persistently. If neither PERSISTENT or TEMPORARY is specified, the default is taken from the



DB2EXT.DBDEFAULTS view, where  
DEFAULTNAME='USEPERSISTENTCACHE'.

#### **MAXIMUM CACHE SIZE** *memsize*

Specifies the maximum size of the cached table to be built during DB2TEXT ACTIVATE CACHE. The memsize parameter must be specified in megabytes as a positive integer. There is no default value for memsize. If the integer is too small, the ACTIVATE CACHE command will fail. The actual cache size is calculated during the ACTIVATE CACHE command.

The limit for the maximum cache size on the different 32-bit platforms is:

- Windows: 1024 MB (1 GB = 1073741824 bytes)
- Linux: 2048 MB (2 GB = 2147483647 bytes)

On 64-bit installations, the maximum cache size limit depends on the available memory.

#### **PCTFREE** *percentage*

Specifies the percentage of the cache to be held free for additional documents. The percentage must be an integer value lower than 100 and greater or equal to 0. If not specified, the default is taken from the db2ext.dbdefaults view, where DEFAULTNAME='PCTFREE' The default is 50%.

#### **INITIAL SEARCH RESULT ORDER** (*SQL-order-by-list*)

Specifies the order used for retrieving the user table contents during initial indexing. When using this option and skipping the dynamic ranking of full-text search results, the documents are returned in their indexing order, as stored in the cached result table.

You can work with presorted indexes and predefined search result ordering only if you are using the Stored Procedure search interface. For example: INITIAL RESULT ORDER(length(column1) asc, column2+column3 desc).

Predefined search result ordering is not implemented for the SQL scalar search functions and the table valued function.

**Note:** The index order can **not** be ensured for the new or changed documents after incremental update.

#### **KEY COLUMNS FOR INDEX ON VIEW** (*SQL-columnname-list*)

If indexes on views are created, the KEY COLUMNS FOR INDEX ON VIEW clause must be specified, otherwise it **MUST NOT** be specified. The list of column names specifies the columns that **UNIQUELY** identify a row in the view.

As this uniqueness cannot be checked by DB2 as in case of primary keys, the user is responsible to ensure the equivalent uniqueness. The specified columns are used to fill the PK01 column of the log table for the index.

#### **INDEX CONFIGURATION** (*option-value*), ...

These are the index configuration values. The default values are underlined.

Option	Values	Description
TreatNumbersAsWords	<u>0</u> or 1	Interprets sequences of digits as separate words if set to 1, even if they are adjacent to characters. The 0 default means that, for example, tea42at5 is considered as one word.

Option	Values	Description
IndexStopWords	0 or <u>1</u>	Ignores or considers stop words during indexing. The default 1 indexes all text including stop words. Currently, the stopword list is an UCS-2 file <language>.tsw in directory <instance>/sql1lib/db2ext/resources. Changes to this file have <b>no effect</b> after index creation. Also note that <language> is the LANGUAGE value from the CREATE INDEX command.
UpdateDelay	seconds	Specifies the duration in seconds for incremental update without capture tables. Only entries older than this duration will be taken from the log table. This is to avoid lost updates. For example, document changes that are not reflected in the index in transaction scenarios where user transactions interfere with update commands. Therefore, the UpdateDelay parameter should be set to the maximum duration of a user write transaction on the table the index was created on.
IgnoreEmptyDocs	<u>0</u> or 1	If IgnoreEmptyDocs is set to 1, empty documents (with content length 0, or a null value) are not represented in the index. If this option is used, and the document content is null (empty), the next incremental update will delete the documents from the index.
AuxLog	ON or OFF	Controls the creation of the additional log infrastructure to capture changes that are not recognized by a trigger. The default setting for range partitioned tables is ON, else it is OFF. The default value can be changed in the default table with a setting for AuxLogNorm for non-range-partitioned tables and AuxLogPart for range-partitioned tables.

**CONNECT TO** *database-name*

The name of the database that is target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

**USER** *userid* **USING** *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

**Changes to the database**

- Change Net Search Extender catalog views.
- Create an index log table in the specified table space. This is only if the RECREATE INDEX option is not specified and the capture table is not specified.
- Create an index event table in the specified table space.
- Deferred to first update: Creation of triggers on the user text table (only if RECREATE INDEX is not specified and no capture table is used)
- If a replication capture table is used, the following change is made to the capture control tables:
  - an insert into the IBMSNAP\_PRUNCTNL and IBMSNAP\_PRUNE\_SET tables



The entries in these tables are uniquely identified by the columns:

- APPLY\_QUAL='NSE' || <DB2 instance running NSE>
- SET\_NAME= <internal index identifier>
- TARGET\_SERVER=<DB2 database name target to DB2TEXT operation>
- If the text-maintained staging infrastructure is configured for the index, an anchor table and a staging table are created to capture change information for documents.

#### Changes to the shared memory

Deferred to ACTIVATE execution: If CACHE TABLE clause is used, a cache for the result table is built in *shared memory*.

#### Changes to the file system

- Subdirectories NODE<nr> are created under index, work and cache directories.
- The directory <internal index name> is created under <indexdirectory>/NODE<nr> where indexdirectory refers to the corresponding parameter of this command and NODE<nr> is related to the partition number in a partitioned database environment.

### Usage notes

Creation of a full-text index requires a primary key on the user table. Since DB2 Net Search Extender Version 9.1, a multi-column DB2 primary key can be used without type limitations. However, to use the table-valued search, no compound primary key is allowed.

The number of primary key columns is limited to 62, the total length of all primary key columns is limited to 1007 bytes for table spaces with pagesize 4k, 2031 bytes for table spaces with pagesize 8k, 4079 bytes for table spaces with pagesize 16k and 4096 bytes for table spaces with pagesize 32k. Note that if the primary key consists of more than one column, the mentioned limits must be reduced by 2 bytes for each additional column.

- The total size of the SQL expressions for ATTRIBUTES, CACHE TABLE and INITIAL SEARCH RESULT ORDER must not exceed 24 K-bytes.
- Initial index updates are always done as one logical transaction, there is no commitcount in this case.

**Note:** After creating the index, the length of primary key columns or the view key columns must not be changed by ALTER TABLE commands.

The synchronization between user table, full-text index and the cached result table is completed during the update index command.

---

## DEACTIVATE CACHE command

### Purpose

This command releases a cached table. A persistent cache is kept to be reused on the next ACTIVATE command. Until the next activation, search operations via the stored procedure are no longer possible on the deactivated cache.

## Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table or nickname on which the text index is defined
- DBADM authority

## Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the DEACTIVATE CACHE command.

## Command syntax

```
DEACTIVATE CACHE FOR INDEX index-schema-"." index-name FOR TEXT
```

Diagram illustrating the command syntax: `DEACTIVATE CACHE FOR INDEX` followed by a bracketed field for `index-schema-"."`, then `index-name`, and finally `FOR TEXT`. A long arrow below indicates the full command structure, with a bracketed field for `connection-options` at the beginning.

### connection-options:

```
CONNECT TO database-name USER userid USING password
```

Diagram illustrating the connection options syntax: `CONNECT TO` followed by a bracketed field for `database-name`, then `USER` followed by a bracketed field for `userid`, and finally `USING` followed by a bracketed field for `password`.

## Command parameters

### *index-schema*

The schema of the text index as specified in the CREATE INDEX command. If no schema is specified, the userid of the DB2 connection is used as schema name.

### *index-name*

The name of the text index as specified in the CREATE INDEX command.

### CONNECT TO *database-name*

The name of the database that is target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

### USER *userid* USING *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

## Usage notes

Note that this command could not be issued when one of the following commands is running on the index:

- ACTIVATE CACHE
- DEACTIVATE CACHE
- UPDATE INDEX

- ALTER INDEX
- DROP INDEX
- CLEAR EVENTS

**Note:** After deactivation of a persistent cache, the cache is made inaccessible for search by the stored procedure. However, this can be used for fast ACTIVATE, unless an update was performed in the meantime.

In this case, the persistent cache is automatically recreated from scratch using the ACTIVATE CACHE command.

## DROP INDEX command

### Purpose

This command drops a full-text index for a text column. If the cache for the index is activated, it is deleted using this command.

### Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table or nickname on which the text index is defined
- DBADM authority

### Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the DROP INDEX command.

### Command syntax

```

▶▶ DROP INDEX [index-schema-"."] index-name FOR TEXT

```

|connection-options|

#### connection-options:

```

|
| CONNECT-TO database-name
| USER userid USING password
|

```

### Command parameters

#### *index-schema*

The schema of the text index as specified in the CREATE INDEX command. If no schema is specified, the userid of the DB2 connection is used as the schema name.

*index-name*

The name of the index as specified in the CREATE INDEX command. With the index schema, it uniquely identifies the full-text index in a database.

**CONNECT TO** *database-name*

The name of the database that is target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

**USER** *userid* **USING** *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

## Usage notes

The index is deleted, irrespective of the activation status of its cached table.

Note that the command must not be issued when one of the following commands is running on the index:

- UPDATE INDEX
- CLEAR EVENTS
- ALTER INDEX
- ACTIVATE CACHE
- DEACTIVATE CACHE
- DROP INDEX

**Note:** Indexes must be manually dropped before or after the user table in DB2 is dropped. If not, the index directories are not correctly cleaned up.

### Changes to the database

- Change Net Search Extender catalog views
- Drop the DB2 index
- Drop the index log, staging and event tables
- Delete triggers on the user text table

When using the replication capture tables, entries in the IBMSNAP\_PRUNE\_SET and IBMSNAP\_PRUNCTRNL tables are removed.

### Changes to the shared memory

The cached table is deleted.

### Changes to the file system

- The directory <internal index name> is deleted in the index and the work directories of the dropped index
- Delete a persistent cache for the index

---

## DB2EXTTH command

### Purpose

This independent utility compiles a thesaurus definition file. After running the thesaurus compiler, the THESAURUS-related features of the search argument syntax can be used.

## Authorization

None. This command is not necessarily restricted for the table owner, but makes only sense in the context of querying.

## Command syntax

```
db2extth -ccsid code page -f definition-file-name
  -quiet
  -h
  -H
  -?
  -copyright
```

## Command parameters

### **-f** *definition-file-name*

The name of the file containing the thesaurus definition. The file name must contain either the absolute path or the relative path to the file. The file name is restricted to 8+3 characters, the extension being optional.

The thesaurus dictionary is generated in the same directory as the definition file and has the same name. The only difference is that the dictionary has the following extensions: wdf, wdv, grf, grv, MEY, ROS, NEY, SOS, and lkn, where n is a digit. Note that if existing thesaurus files have the same name, they are overwritten.

### **-ccsid** *code page*

The code page in which the thesaurus definition file is written.

**-quiet** Output information is not displayed.

### **-copyright**

Returns the internal build number of the product. Use this number when reporting problems.

### **-h | -H | -?**

Displays help information.

## Usage notes

Use this command to compile a thesaurus definition file into a binary thesaurus definition format.

---

## RESET PENDING command

### Purpose

When you use the extended text-maintained staging infrastructure, certain commands cause the staging table to go into pending mode, which blocks other database or text search operations. You can use the RESET PENDING command to execute a set integrity for all text-maintained staging tables associated with a given table. You do not have to find all text indexes and associated staging tables to execute a SET INTEGRITY command for each table.

After detaching a range partition, you must run the RESET PENDING command to update the staging table content.

## Authorization

You must have the CONTROL privilege on the table.

If the SECADM modifies the access privileges for the text-maintained staging table or tables associated with a base table, you still require the privilege to execute set integrity statement for the text-maintained staging table or tables.

## Required connection

This command must be issued from the DB2 database server.

## Command syntax

```
▶▶—RESET PENDING FOR TABLE table-schema.table-name FOR TEXT————▶▶  
|_____|  
|connection-options|
```

### Connection-options:

```
|_____|  
|CONNECT TO—database-name_____|  
|_____|  
|USER—userid—USING—password|
```

## Command parameters

### Table name

The name of the table for which the text-maintained staging infrastructure has been added and which requires integrity processing.

### Table-schema

The schema of the table for which a command was executed that results in a pending mode for dependent tables.

## Usage notes

Use the RESET PENDING command after running a command that causes dependent tables to be put into pending mode, such as a LOAD command with the INSERT parameter, or a command that requires a set integrity statement to refresh dependent tables, such as ALTER TABLE ... DETACH.

---

## UPDATE INDEX command

### Purpose

This command immediately starts the indexing process by bringing the index up to date to reflect the current contents of the text columns with which the index is associated.

While the update is being performed, search using the CONTAINS predicate is possible. For an index with an activated cached result table, search by stored procedure is also possible during update. However, columns in the cached table may show new values, even though the changed text is not yet committed to the full-text index.

Using the RECREATE INDEX ON UPDATE option in the CREATE INDEX command will clear the index before recreation. Until completion of the update, empty results will be returned.

## Authorization

The privileges held by the authorization ID of the statement must include at least one of the following:

- CONTROL privilege on the table or nickname on which the text index is defined
- DATAACCESS authority

## Prerequisite

The instance owner must hold DBADM with DATAACCESS authority. The SYSADM no longer holds SECADM nor DBADM privilege in Version 9.7. The SECADM must explicitly grant DBADM with DATAACCESS to the instance owner before running the UPDATE INDEX command.

## Command syntax

```

▶▶ UPDATE INDEX index-schema-".." index-name FOR TEXT REORGANIZE
▶
┌ COMMITCOUNT count ─┐ ┌ USING UPDATE MINIMUM ─┐ ┌ connection-options ─┐

```

### connection-options:

```

┌ CONNECT TO database-name ─┐
┌ USER userid USING password ─┐

```

## Command parameters

### *index-schema*

The schema of the text index. This is specified in the CREATE INDEX command. If no schema is specified, the user ID of the DB2 connection is used.

### *index-name*

The name of the text index. This is specified in the CREATE INDEX command.

### REORGANIZE

If a text column is frequently updated, then subsequent updates to the index can become inefficient. To make the update process efficient again, reorganize the index. Use the DB2EXT.TEXTINDEXES view to determine if an index needs to be reorganized.

Use the REORGANIZE AUTOMATIC option of the CREATE INDEX command to avoid manually checking and reorganizing the index.

**Note:** The reorganization process takes place after a regular update.

### USING UPDATE MINIMUM

Uses the UPDATE MINIMUM settings from the CREATE INDEX command and starts an incremental update only if the specified number of

changes was reached. If the text-maintained staging infrastructure is configured, the number of changes is combined for this determination. The default is to unconditionally start the update.

For distributed databases, the UPDATE MINIMUM is checked on every partition.

**COMMITCOUNT** *count*

An INTEGER value  $\geq 0$  determines the number of documents processed in one transaction by the search engine and by DB2 for incremental index updates. You can change this value using the ALTER INDEX command.

However, for initial updates, such as the first update after the CREATE INDEX command, or any update with RECREATE INDEX ON UPDATE option, there is only one logical transaction which ignores COMMITCOUNT.

Using a non-zero COMMITCOUNT is not recommended due to its massive negative impact on indexing throughput.

**CONNECT TO** *database-name*

The name of the database that is the target for this command. You can omit this parameter, if DB2DBDFT is set and the user is running the command on the server. Note that the user ID must have the required DB2 authorizations.

**USER** *userid* **USING** *password*

Use a *userid* and *password* to connect to the database. If not specified, a connection is attempted from the current user ID without a password.

## Usage notes

This command runs synchronously. It starts the update processing on all required DB2 logical/physical partitions in a partitioned database environment. The duration depends on the number of documents to indexed and the number of documents already indexed. The status of the update can be seen through a view that is created for each index. The name of this view can be retrieved from DB2EXT.TEXTINDEXES in column EVENTVIEWNAME.

There are two options to view the number of committed documents that have been processed. To determine how many documents have been committed to the index, use the DB2EXT.TEXTINDEXES (NUMBER\_DOCS) view. Use the event view associated with the index for information on starting, committing changes, and finishing update processing.

To view the number of documents that have been processed, while an index update is still active, use the CONTROL LIST ALL LOCKS FOR INDEX command.

**Note:** The views only display information from the connected partition.

For incremental updates on a base table with multiple physical partitions, the time on each partition must be synchronized. If the times are not synchronized, updates may be lost or may not occur at all.

You cannot issue the command if one of the following commands is running on the index:

- CLEAR EVENTS
- ALTER INDEX



- DROP INDEX
- ACTIVATE CACHE
- DEACTIVATE CACHE
- UPDATE INDEX

After updating an index with a deactivated persistent cached result table, the persistent cache is deleted, such that the next ACTIVATE CACHE command recreates it based on the database content.

If the user interrupts this command, all processes involved in the update function will stop. If COMMITCOUNT was used in an incremental update, some updates might have been committed and are visible in the index, while others may require a new update command. Forcefully interrupting the update index process has the potential to corrupt the index.

To stop the automatic updating of an index, look for the DB2 instance owner process running the update index command on the partition used for update services. Stop this process and the update processing on all the partitions.

**Note:** As the command works in two separate phases for index creation on all partitions and initial index updates, issue a db2text drop index command to ensure that the index is not partly available. If this command is not issued, the next update, which can be triggered by a manual update command or the update frequency option, would perform a complete re-indexing to ensure a consistent state.

#### Changes to the database

- Inserts to the event table
- Delete from the index log table, and text-maintained staging table for the index, if configured

When using the replication capture tables, the following changes are made to the database.

- A signal is added to the IBMSNAP\_SIGNAL table before starting the initial update
- The synchpoint of IBMSNAP\_PRUNE\_SET is changed after the incremental update

## HELP command

### Purpose

This displays the list of available DB2TEXT commands, or the syntax of an individual DB2TEXT command.

### Authorization

None

### Command syntax



## Command parameters

### HELP | ?

Provides help for the specified command or reason code.

### *command*

The first keywords that identify a DB2TEXT command:

- ENABLE
- DISABLE
- CREATE
- DROP
- ALTER
- UPDATE
- CLEAR
- START
- STOP
- CONTROL
- ACTIVATE
- DEACTIVATE
- RESET PENDING

### *reasoncode*

Reason code from a Net Search Extender command such as CTE0192.

## Usage notes

If more than the first keyword is specified, the rest are ignored and the syntax of the identified command is displayed.

If no *command* parameter is specified after ? or HELP (or no parameter at all), DB2TEXT lists all available DB2TEXT command parameters.

---

## COPYRIGHT command


### Purpose

Provides Net Search Extender product and copyright information.

### Authorization

None

### Command syntax



## Command parameters

### COPYRIGHT | LEVEL

Provides the version copyright statement, version number, and build information for the product.

---

## Chapter 50. Net Search Extender install and uninstall command reference for UNIX

This section describes the syntax of the Net Search Extender install and uninstall commands for UNIX. It also includes the `db2nse1s` command that displays the installed copies of the Net Search Extender product starting with version 9.

---

### db2nse\_install command

#### Purpose

This command installs a new version of Net Search Extender on a UNIX system.

#### Authorization

You must run this command as a root user.

#### Command syntax

```
db2nse_install package-file path [-s] [-f] [-v]
```

#### Command parameters

*package-file*

The name of the file containing the Net Search Extender product.

*path*

The DB2 path where you now want to install Net Search Extender.

**-s**

Silent installation. Checks are carried out and a log file is written to the `/tmp` directory.

**-f**

Force installation. No checks are done.

**-v**

Displays the program version and exits.

#### Usage notes

This command installs Net Search Extender contained in the package file passed as a parameter. If you do not specify any additional parameters, the installation program checks the system for eligible copies of DB2 where Net Search Extender can be installed. All DB2 installation paths are listed but only a viable subset is selectable. You are asked by the system to select an installation path.

If you select to install silently, there is no user interaction. The outcome of the installation process is passed as a return code to the calling program which might be a command shell or a shell script. It is up to you to handle the return code correctly. A log file that documents the installation is written to the `/tmp` directory.

If you select to force install, there is no user interaction and there are no additional checks. The outcome of the installation process is passed as a return code to the

calling program which might be a command shell or a shell script. It is up to you to handle the return code correctly. A log file that documents the installation is written to the /tmp directory.

Both the parameters -s and -f require that a path is passed. No default path exists for these parameters.

If the -v parameter is passed, only the version of the program is displayed and the program exits without any further action.

---

## db2nse\_deinstall command

### Purpose

This command removes Net Search Extender on a UNIX system. The db2nse\_deinstall command is located in each install subdirectory of a DB2 copy where Net Search Extender is installed.

### Authorization

You must run this command as a root user.

### Command syntax

▶▶—db2nse\_deinstall -v▶▶

### Command parameters

-v      Displays the program version and exits.

### Usage notes

This command removes Net Search Extender beginning with DB2 V9. The command removes the copy of Net Search Extender in the directory where it is issued. It does not remove other copies of Net Search Extender installed elsewhere.

---

## db2nsels command

### Purpose

This command displays a table containing all of the installed copies of Net Search Extender. It only displays versions starting with DB2 Version 9.7. The db2nsels command is located in the /usr/local/bin directory.

### Authorization

You must run this command as a root user.

### Command syntax

▶▶—db2nsels -c -v▶▶

## Command parameters

- c Displays the installed Net Search Extender versions as a simple compact list separated by colons, and exits.
- v Displays the program version, for example 9.7, and exits.

## Usage notes

This command lists all of the installed copies of Net Search Extender starting with DB2 V9. The result is displayed either in a table or as a compact list where the items are separated by colons. The results include the path, the version and the fix pack number of the Net Search Extender installation.

For example, if you call `db2nse1s` with no parameter, the output might be as follows:

```
db2nse1s
```

Install path	Level	FP
/opt/ibm/db2/V9.7	9.7.0.0	0
/test/V9.7	9.7.0.0	0

If you use the `-c` parameter, the returned output is a compact list where each information item is separated by a colon. This type of output is easy to handle using a program or a shell script. For example:

```
db2nse1s -c
```

```
#PATH:VRMF:FIXPACK  
/opt/ibm/db2/V9.7:9.7.0.0:0  
/test/V9.7:9.7.0.0:0
```



---

## Chapter 51. Syntax of search arguments

A search argument is the condition that you specify when searching for terms in text documents. It consists of search parameters and one or more search terms.

Examples of search arguments are given in Chapter 34, "Specifying SQL search arguments," on page 119, and in a file called search in the Net Search Extender samples directory.

The SQL scalar search functions that use search arguments are:

### CONTAINS

This function uses a search argument to search for text in a particular text document. It returns the INTEGER value 1 if the document contains the text, or any relation specified in the search argument. Otherwise, it returns 0.

### NUMBEROFMATCHES

This function uses a search argument to search in text documents and returns an INTEGER value indicating how many matches resulted per document.

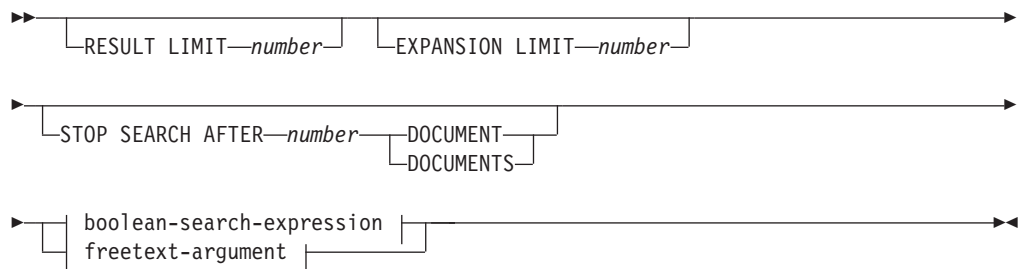
### SCORE

This function uses a search argument to search in text documents. It returns a value for each document found, indicating how well the found document is described by the search argument compared to other documents in the same index.


**Note:** The same syntax is used in the search arguments of the stored procedure search and the SQL table-valued function.

---


## Search argument syntax



### Boolean-search-expression:




### search-term:




**Search-factor:**



**Positive-search-factor:**



**Search-primary:**




**Operator-and:**



**Operator-or:**




**Operator-accum:**



**Operator-minus:**



**Context-condition:**







### Context-argument:




### Text-literal-list:




### Context-unit:




### Text-literal:




### thesaurus-invocation:



### Attribute-factor:



### freetext-argument:



## Examples

Examples are given in Chapter 34, "Specifying SQL search arguments," on page 119.

---

## Search parameters

### Parameters

#### **RESULT LIMIT** *number*

A keyword specifying the maximum number of results to be returned by the full-text search.

The RESULT LIMIT should be used together with the SCORE function to ensure that the returned results are scored and only the best matching results are processed.

#### **EXPANSION LIMIT** *number*

A keyword specifying the maximum number of terms that a wildcard term can be expanded to for searching. For example, to determine how many times you can expand the search term 'a\*'. If your index is very large and you are using many wildcard terms, you must adjust the value of this keyword if you want to obtain a larger result set. The expansion order depends on the internal organization of the text index and cannot be predetermined. If your wildcard expression is too general, and can be expanded into more search terms than specified by 'EXPANSION LIMIT', the search returns with an error, indicating that the search result has been truncated due to this limit exhaustion.

#### **STOP SEARCH AFTER** *number* **DOCUMENT | DOCUMENTS**

A keyword specifying the search threshold. The search is stopped when the given number of documents is reached during the search, and an intermediate result is returned. A lower value will increase the search performance, but may lead to fewer results and omit documents with a potentially high rank.

Note that there is no default value and the *number* value must be a positive integer.

#### **boolean-search-expression**

The search-terms and search-factors can be combined using the boolean operators NOT, AND, OR, ACCUM, and MINUS according to the syntax diagrams. The operators have the following precedence order (with the strongest first): NOT > MINUS = ACCUM = AND > OR. This can be seen in the following example:

```
"Pilot" MINUS "passenger" & "vehicle" | "transport" & "public"
```

is evaluated as:

```
((("Pilot" MINUS "passenger") & ("vehicle"))) | ("transport" & "public")
```

The operator ACCUM evaluates to true, if one of the boolean arguments evaluates to true (which is comparable to the OR operator). The rank value is computed by accumulating rank values from both operands. The ACCUM operator has the same binding (precedence) as AND. The operator MINUS evaluates to true, if the left operand evaluates to true. The rank value is computed by taking the rank value for the left operand and subtracting a penalty, if the right operand evaluates to true.

### search-primary

A search-primary, consisting of a text-literal-list evaluates to true, if any of the text-literals is found in the (specified section of the) document. A search-primary consisting of a thesaurus-invocation evaluates to true, if any of the expanded text-literals is found in the (specified section of the) document.

### SECTION | SECTIONS *section-name*

A keyword specifying one or more sections in a structured document that the search is to be restricted to. The section name must be specified in a model file specified at index creation time or be expressed in XPath notation.

Section names are case sensitive. Ensure that the case of the section name in the model file and query is identical.

This model describes the structure of documents that contain identifiable sections, so that the content of these sections can be individually searched. Section names cannot be masked using masking characters. The *positive-search-factor* using the SECTION clause evaluates to true, if the search primary is found in one of the specified sections.

Section names are not valid XPath expressions that are evaluated during query execution. If no model file is used, the default section names are phrased in XPath notation. The absolute path expression to the element (such as `/father/child/grandchild`) is used as the name for identifying the section. Full XPath expressions are not supported as section names.

### context-argument IN SAME context-unit AS context-argument AND context-argument ...

This condition lets you search for a combination of text-literals occurring in the same paragraph or same sentence. Context arguments are always equivalent to text-literal-lists, and thesaurus expansion may be used to expand a text-literal to such a list.

The condition evaluates to true, if there is a context-unit (paragraph or sentence, respectively) in the document, which contains at least one of the text-literals of each expanded context-argument. This can be seen in the following example:

```
("a","b") IN SAME PARAGRAPH AS ("c","d")
AND THESAURUS "t1" EXPAND SYNONYM TERM OF "e".
```

Assuming e1, e2 are synonyms of e, the following paragraphs would match:

```
".. a c e .." , ".. a c e1.." , "a c e2..",
".. a d e .." , ".. a d e1.." , "a d e2..",
".. b c e .." , ".. b c e1.." , "b c e2..",
".. b d e .." , ".. b d e1.." , "b d e2..".
```

### PRECISE FORM OF

A keyword that causes the word (or each word in the phrase) following PRECISE FORM OF to be searched for exactly as typed. This form of search is case-sensitive; that is, the use of upper- and lowercase letters is significant. For example, if you search for mice, you do not find "Mouse".

This parameter requires that the index configuration parameter Respect case is set to yes. This configuration setting cannot be changed after the index has been built.

## STEMMED FORM OF

A keyword that causes the word (or each word in the phrase) following STEMMED FORM OF to be reduced to its word stem before the search is carried out. This form of search is not case-sensitive. For example, if you search for mouse, you find "Mouse".

The way in which words are reduced to their stem form is language-dependent. Currently, only English stemming is supported and the word must follow regular inflection endings.

## FUZZY FORM OF

A keyword for making a "fuzzy" search, which is a search for terms that have a similar spelling to the search term. This is particularly useful when searching in documents that were created by an Optical Character Recognition (OCR) program. Such documents often include misspelled words. For example, the word economy could be recognized by an OCR program as economy. Note that successful matches are only returned for words in a document where the first three characters match. In the above example, ecanomy is not a match. Fuzzy search cannot be used if a word in the search atom contains a masking character.

### *match-level*

An integer between 1 and 100 specifying the degree of similarity, where 100 is more similar than 1. 100 specifies an "exact match", and 60 is already considered a very "fuzzy value". The fuzzier the match level is, the longer the elapsed search time, since more documents qualify for the search. The default match level is 70.

### **WEIGHT** *number*

Associates a text-literal with a weight value to change the default score. The allowed weight values are integers between 0 (the lowest score weighting) and 1000 (the highest); the default value is 100.

### *word-or-phrase*

A word or phrase to be searched for. The characters that can be used within a word are language-dependent. It is also language-dependent whether words need to be separated by separator characters. For English and most other languages, each word in a phrase must be separated by a blank character.

To search for a character string that contains double quotation marks, type the double quotation marks twice. For example, to search for the text "wildcard" character, use:

```
""wildcard"" character"
```

Note that in the example, it is only possible to search for one set of quotation marks. You cannot search for two quotation marks in a sequence. There is also a maximum length of 128 bytes for each word or phrase.

## **Masking characters**

A word can contain the following masking characters:

### **\_ (underscore)**

Represents any single character.

### **% (percent)**

Represents any number of arbitrary characters. If a word consists of a single %, then it represents an optional word of any length. A word cannot be composed exclusively of masking characters, except when a single % is used to represent an optional word. If

you use a masking character, you cannot use the THESAURUS keyword. Masking characters cannot be used inside thesaurus query parts. If they are used in combination, search results are unpredictable. Masking characters cannot follow a non-alphanumeric character. Masking characters cannot be used inside a fuzzy search as masking always expands into a single word.

**ESCAPE** *escape-character*

A character that identifies the next character as one to be searched for and not as one to be used as a masking character. For example, if an escape-character is \$, then \$%, \$\_, and \$\$ represent %, \_, and \$ respectively. Any % and \_ characters not preceded by \$ represent masking characters.

During search, you are only allowed to use single-byte escape characters. No double-byte characters are allowed.

**THESAURUS** *thesaurus-name*

A keyword used to specify the name of the thesaurus to be used to expand a text-literal. The thesaurus name is the file name (without its extension) of a thesaurus that has been compiled using the thesaurus compiler. It must be located in <os-dependent>/sql11b/db2ext/thes. Alternatively, the full path can be specified preceding the file name.

**EXPAND** *relation*

Specifies which relation is used to expand the text-literal using the thesaurus. The thesaurus has predefined relations described in the DB2EXTTH command. These are referred to using the following keywords:

- SYNONYM, a symmetrical relationship expressing equivalence.
- RELATED, a symmetrical relationship expressing association.
- BROADER, a directed hierarchical relationship that can be followed by specified depth levels.
- NARROWER, a directed hierarchical relationship that can be followed by specified depth levels.

For user-defined relations, use RELATION(number), that corresponds to the relation definition in DB2TEXTTH.

**TERM OF** *text-literal*

The text-literal, to which other search terms are to be added from the thesaurus.

*count* **LEVELS**

A keyword used to specify the number of levels (the depth) of terms in the thesaurus that are to be used to expand the search term for the given relation. If you do not specify this keyword, a count of 1 is assumed. The value of depth must be a positive integer value.

**ATTRIBUTE** *attribute-name*

Searches for documents that have attributes matching the specified condition. The attribute-name refers to the name of an attribute expression in the CREATE INDEX command, or to an attribute definition in the document model file.

The attribute-factor is allowed for attributes of type double only. The precision of the value is guaranteed for 15 digits. Numbers that consist of

16 digits and above are rounded. Usage of masking characters is not allowed in attribute-name, valueFrom and, valueTo. For an explanation, see the following:

**BETWEEN** *valueFrom* **AND** *valueTo*

A BETWEEN attribute factor evaluates to true if the value of the attribute is greater than (not equal to) *valueFrom* and smaller than (not equal to) *valueTo*.

**>***valueFrom*

A ">" attribute factor evaluates to true if the value of the attribute is greater than (not equal to) valueFrom.

**<***valueTo*

A "<" attribute factor evaluates to true if the value of the attribute is lower than (not equal to) valueTo.

If the attribute name in the CREATE INDEX command is specified with quotes, or is defined in a model file, the specified attribute name must match exactly. Whereas, if no quotes are specified in the CREATE INDEX command, the attribute name must be in uppercase.

**IS ABOUT** *language word-or-phrase*

An option that lets you specify a free-text search argument. Using IS ABOUT, you can search for any (but not necessarily all) of the words that you specify in word-or-phrase in any order in a document. The closer together the terms used in word-or-phrase are and the more terms that are included in a document, the higher the returned score for the document.

The parameter language is optional and must be set only for Thai (TH\_TH) where it is required for tokenization purposes, and for Turkish (TR\_TR), where it is required for proper case mapping.

Note that IS ABOUT is useful only if document score values are requested and the search results are ordered by score values.

---

## Chapter 52. SQL scalar search function and the SQL table-valued function

Net Search Extender provides SQL scalar search functions and an SQL table-valued function for searching text documents stored in DB2.

This section describes the following SQL search functions.

Search function	Purpose
"CONTAINS scalar function"	Searches for text in a particular document.
"NUMBEROFMATCHES scalar function"	Searches and returns the number of matches found.
"SCORE scalar function" on page 240	Searches and returns the score value of a found text document.
"DB2EXT.TEXTSEARCH command" on page 240	The SQL table-valued function returns a table of found primary keys, a number of matches, and/or score values.
"DB2EXT.HIGHLIGHT" on page 243	To get information about why a document qualified as a search result.

---

### CONTAINS scalar function

#### Purpose

The CONTAINS scalar function searches for text in a text document indexed by Net Search Extender. It returns the INTEGER value 1 if the document contains the text, or any relation specified in the search argument. Otherwise, it returns 0.

#### Function syntax

►►—CONTAINS—(—*column-name*—,—*search-argument*—)—————►►

#### Function parameters

##### column name

The name of a table column. The column must have an associated text index. You can create text indexes by using the administration command DB2TEXT CREATE INDEX.

##### search-argument

A string of type VARCHAR containing the terms to be searched.

**Note:** You cannot use the CONTAINS query on a text index created on a view.

---

### NUMBEROFMATCHES scalar function

#### Purpose

The NUMBEROFMATCHES scalar function searches in text documents and returns an INTEGER value indicating how many matches resulted per document.

## Function syntax

►►—NUMBEROFMATCHES—(—*column-name*—,—*search-argument*—)—►►

## Function parameters

### column name

The name of a table column. The column must have an associated text index. You can create text indexes by using the administration command `DB2TEXT CREATE INDEX`.

### search-argument

A string of type VARCHAR containing the terms to be searched.

**Note:** You cannot use the NUMBEROFMATCHES query on a text index created on a view.

---

## SCORE scalar function

### Purpose

The SCORE scalar function searches in text documents and returns a score value for each document found, indicating how well the found document is described by the search argument.

SCORE returns a DOUBLE value. As the search term appears more frequently in the document, the score of the document increases.

## Function syntax

►►—SCORE—(—*column-name*—,—*search-argument*—)—►►

## Function parameters

### column name

The name of a column. The column must have an associated text index. You can create text indexes by using the administration command `DB2TEXT CREATE INDEX`.

### search-argument

A string of type VARCHAR containing the terms to be searched.

**Note:** You cannot use the SCORE query on a text index created on a view. The values returned by score are only meaningful if they are compared to other values retrieved from the same index.

---

## DB2EXT.TEXTSEARCH command

### Purpose

In addition to the stored procedure search and the SQL scalar search functions, Net Search Extender provides two SQL table-valued functions which look very similar to the stored procedure.



Both table-valued functions are called `db2ext.textsearch`. The only difference between them is that one supports the `HIGHLIGHT` function and has two additional parameters, `numberOfHits` and `hitInformation`.

Both table-valued functions return the results from the user tables sorted according to the sort criteria defined in the parameter called `INITIAL SEARCH RESULT ORDER` in the `CREATE INDEX` command. If the SQL query statement containing the table-valued function has a join at the end of the statement (for example, something like `where T.primkey = S.key`), the order of the result rows depends on the join method and not on the order defined in the `CREATE INDEX` command.

- You cannot use the table-valued function on tables with a compound primary key.
- The table-valued function may be used in a partitioned database environment only if the user table is stored in a single-partition tablespace. You must also ensure that you connect to the correct node using the `DB2NODE` environment variable.

#### 1. `db2ext.textsearch` without highlight support

```
db2ext.textSearch
(
  query          VARCHAR(4096),
  indexSchema    VARCHAR(128),
  indexName      VARCHAR(128),
  resultFirstRow INTEGER,
  resultNumberRows INTEGER,
  primKeyBinding <supported types>, // same type as primary key
)

return table
(
  primKey        <supported types>, // same type as primary key
  numberOfMatches INTEGER,
  score          DOUBLE,
  totalNbResults INTEGER
)
```

#### 2. `db2ext.textsearch` with highlight support

```
db2ext.textSearch
(
  query          VARCHAR(4096),
  indexSchema    VARCHAR(128),
  indexName      VARCHAR(128),
  resultFirstRow INTEGER,
  resultNumberRows INTEGER,
  primKeyBinding <supported types>, // same type as primary key
  numberOfHits  INTEGER
)

return table
(
  primKey        <supported types>, // same type as primary key
  numberOfMatches INTEGER,
  score          DOUBLE,
  totalNbResults INTEGER,
  hitInformation BLOB(20K)
)
```

## Function parameters

The following are input parameters.

**query** See Chapter 51, "Syntax of search arguments," on page 231 for additional information.

**indexSchema, indexName**  
Identifies the index to search.

**resultFirstRow**  
The result list of the query is returned in parts. This parameter describes which row of the query result list is the first one to be entered into the result table of the table-valued function. The value must be  $\geq 0$ .

Note that the number 0 identifies the first row in the query result list.

**resultNumberRows**  
This parameter describes how many rows of the query result list are entered into the result table of the table-valued function, and where 0 means that all the results need to be returned.

Note that this is different from the result limit query parameter that determines the maximum size of the query result list.

**primaryKeyBinding**  
The type of this parameter determines the type of the primaryKey Output parameter. If the text index has been created for a base table with a primary key of type `<type1>`, then primaryKeyBinding must also be of type `<type1>`.

Additionally, the parameter determines the scope of the text search. If primaryKeyBinding is set to NULL ("CAST(NULL as <type1>)", the scope of the search will be all the documents stored in the index. Alternatively, you can restrict the search to documents primaryKeyBinding is bound to.

For example, if primaryKeyBinding is set to CAST(5 as BIGINT), you restrict the search to the single document with the BIGINT primary key value of "5".

Note that only single column primary keys of the following types are supported: SMALLINT, INTEGER, BIGINT, REAL, DOUBLE, VARCHAR FOR BIT DATA, DATE, TIME, and TIMESTAMP.

**numberOfhits**  
This option specifies the number of terms that are highlighted using the highlighting function called `db2ext.highlight`. If 0 is specified, all of the hits are highlighted up to a maximum of 1100 hits. This process may be time-consuming.

## Function parameters

The following return values are stored in a temporary table which needs to be joined to your user table if further results are requested. Note that the NUMBEROFMATCHES, SCORE, TOTALNUMBEROFRESULTS, and HITINFORMATION are only calculated if they are requested in your select statement.

**primKey**  
The primary key of the found document.

**numberofmatches**

NUMBEROFMATCHES is an INTEGER value indicating how many matches resulted for each document.

**score** Score returns a DOUBLE value. As the search term increases in frequency in the document, the document score increases.

**totalNumberOfResults**

The query result list denotes how many results were found. Note that each row has the same value.

Also note that when you use the STOP SEARCH AFTER, or the RESULT LIMIT together with the SCORE syntax in a query, this number is no longer reliable.

**hitInformation**

The hit information returned by `db2ext.textsearch` is necessary for highlight processing. Currently, hit information for approximately 1100 hits can be contained in this output parameter. If the number of hits exceeds this threshold, hit information for these further hits is ignored.

Note that this value is only returned if you specify `numberOfHits`.

**Usage**

With the SQL table-valued function you are able to search on views in the same way you do with the stored procedure search. Only with the SQL table-valued function no shared memory is needed, so the index does not need to have a cache that must be activated.

This function is primarily for those users who have used an SQL query within the stored procedure search. However, the restriction is that only a single column primary key on base tables is supported.

The following example shows how you can work on a multi-column primary key table:

```
select s.id from
db2ext.sample s, table (db2ext.textSearch(
    "characteristics",
    'DB2EXT',
    'COMMANDS',
    1,
    20,
    cast(NULL as INTEGER))) t
where s.id = t.primkey
```

In this example, you must first create a view on this table with a single unique key and then create the index on this view.

For an example of using the SQL table-valued function with the `db2ext.highlight` function, see "DB2EXT.HIGHLIGHT."

---

**DB2EXT.HIGHLIGHT****Purpose**

Use the `db2ext.highlight` function to get information that can be used to display why a particular document qualifies as a search result. More specifically, it can be used to:

- get hits
- get hits and surrounding text
- get the document with user-defined highlight tags surrounding the hits.

Note that the `db2ext.highlight` function can only be used with the `db2ext.textsearch` table-valued function. The table-valued function searches the index providing the results for the HIGHLIGHT function to use.

## Function syntax

```
▶▶—db2ext.highlight—————▶▶
▶—(—document-content—,—hit-information—,—hit-processing-information—)————▶▶
```

## Function parameters

The following are input parameters:

### document content CLOB(100K)

Only UTF8 documents of TEXT or serialized XML format are supported. For highlighting natively stored XML documents, the XML data must be serialized to CLOB by using the XMLSERIALIZE SQL/XML function.

To increase the CLOB value, use the “DB2EXTTH command” on page 220.

### hit information BLOB(20K)

A string containing hit information. This is returned by the `db2ext.textsearch` function, if the `numberOfHits` parameter is specified.

### hit processing information VARCHAR(1024)

This parameter is a list of option value pairs separated by a comma ‘,’ character with each string character surrounded by “ ” characters. It specifies how highlighting should be processed for the specified document. If none of the options is specified, the original document content is returned unmodified.

#### TAGS = (“STRING”, “STRING”)

This option enables the user to specify the tags to be inserted before and after a hit in the document. If this option is omitted, no tags are added before and after a hit in the document.

#### WINDOW\_NUMBER = INTEGER

This option specifies how many parts (or windows) of the document should be returned by the highlight function. Each window contains one or more hits and the first hit in each window determines the part of the document returned to the user. These hits may or may not have text surrounding the hit.

If this option is omitted, 0 is taken as the default and the entire document containing start and end tags (if specified) is returned. In this case, the WINDOW\_SIZE option is ignored.

#### WINDOW\_SIZE = INTEGER

This option specifies the recommended size of the window in bytes. This actual size may vary, depending on the number of hits, length of hits and the start and end tag sizes. If the option is omitted, 0 is the default and only hits without surrounding text will be returned.

**WINDOW\_SEPARATOR = "STRING"**

This option specifies the tag used to separate one window from the next window. If the option is omitted, "..." is the default value.

**FORMAT = "STRING"**

This option specifies the format of the document. Valid values are XML or TEXT. If this option is omitted, then TEXT is taken as the default value. Ensure that the format value is the same as that specified during indexing.

**MODEL\_NAME = "STRING"**

This option specifies the model name related to the specified XML document. Note that if the FORMAT is TEXT, this option results in an error condition.

**SECTIONS = ("section-name1", ..., "section-nameN")**

For XML documents, highlighting can be restricted to relevant sections. For example, they can be defined in the model file. To specify these sections, separate the one or more section names with a comma. If this option is omitted, highlighting is performed on the whole XML document. Note that if the FORMAT is TEXT, this option is ignored.

Section specification ("section-name1",..., "section-nameN") used in DB2EXT.HIGHLIGHT must be the same as is for the DB2EXT.TEXTSEARCH function.

## Function parameters

The following are return parameters.

**CLOB(200K)**

The HIGHLIGHT function returns a CLOB value containing the document parts modified by the HIGHLIGHT function.

## Usage

The following example shows how you can use the HIGHLIGHT function:

```
select p.id,
       p.title,
       db2ext.highlight(p.content,
                       t.hitinformation,
                       'TAGS = ("<bf>", "</bf>"),
                       WINDOW_NUMBER = 5,
                       WINDOW_SIZE = 200,
                       WINDOW_SEPARATOR = "...",
                       FORMAT = "XML",
                       SECTIONS = ("section1-name", "section2-name"))
FROM patent p, table (db2ext.textsearch(
  "relational database systems",
  'DB2EXT',
  'TI_FOR_CONTENT',
  0,
  20,
  CAST(NULL as BIGINT),
  15)) t
WHERE p.id = t.primkey
```

Using documents larger than 100 KB will cause the SQL query to terminate and produce an SQL error (SQL 1476N and sql error -433). To avoid this, use the

db2exthl command to increase the permitted document content size.

<b>Note:</b> Special characters, such as "newline" will be returned as is.
--

Highlighting natively stored XML documents requires a serialization of these XML documents to CLOB before they can be passed to the HIGHLIGHT table-valued function. The following example shows you how you can use the HIGHLIGHT function on natively stored XML documents using the XMLSERIALIZE SQL/XML function. The patent content of the sample below is stored as native XML. Notice that FORMAT="XML" is also specified:

```
select p.id,
       p.title,
       db2ext.highlight(XMLSERIALIZE(p.content AS CLOB(100K)),
                        t.hitinformation,
                        'TAGS = ("<bf>","</bf>"),
                        FORMAT = "XML",
                        SECTIONS = ("section1-name", "section2-name")')
FROM patent p, table (db2ext.textsearch(
    'xml database systems',
    'DB2EXT', 'TI_FOR_XML',
    0,
    20,
    CAST(NULL as BIGINT),
    15)) t
WHERE p.id = t.primkey
```

## Restrictions

- Only XML and flat text documents are supported.
- Only UTF8 databases are supported. For binary documents, you need to ensure that the documents are encode in UTF8.
- Thai documents are not supported.
- If there is a mismatch between the document format used during indexing and query time the HIGHLIGHT function will return unpredictable results. This is particularly true for cases where a transformation function is used for obtaining the text during indexing. Any change to the transformation function between the time of the indexing, and the time of submitting the search request that affects the position of text tokens in the output will render the results of the highlight function invalid.
- Only hits found in the text parts of a document will be highlighted.
- The highlight table-valued function can only be used with the db2ext.textsearch function.
- String values cannot contain the " character.

---

## Chapter 53. Stored procedure search function

Net Search Extender provides a stored procedure search for returning predefined result tables. The result table is specified in the cache table section during create index. Use the stored procedure search when you need to return a small number of results in a specific order.

An example would be a Web application where the first 20 rows with the best score are returned, but the rest of the results can also be returned in increments of 20 rows.

**Note:** The stored procedure function may be used in a partitioned database environment only if the user table is stored in a single-partition tablespace.

You must also ensure that you connect to the correct partition using the DB2NODE environment variable.

---

### DB2EXT.TEXTSEARCH for stored procedure search

#### Function syntax

db2ext.TextSearch(

IN	query	VARCHAR(4096),
IN	indexSchema	VARCHAR(128),
IN	indexName	VARCHAR(128),
IN	resultFirstRow	INTEGER,
IN	resultNumberRows	INTEGER,
IN	scoringFlag	INTEGER,
IN	searchTermCountsFlag	INTEGER,
OUT	searchTermCounts	VARCHAR(4096),
OUT	totalNumberOfResults	INTEGER )

#### Function parameters

The following are input parameters.

**Query** See Chapter 51, "Syntax of search arguments," on page 231 for further information.

**indexSchema, indexName**

To identify the index to search.

**resultFirstrow**

The query result list is returned in parts. The parameter describes which row of the query result list is the first one to be put into the result set of the stored procedure. The first row in the query result list is identified by the number 0.

**resultNumberRows**

This parameter describes how many rows of the query result list are put into the result set of the stored procedure.

This is not to be confused with the "result limit" expression in the query, which determines the maximum size of the query result list.

The value should be  $\geq 0$ . Where 0 means that all the results need to be returned.

**Note:** If a larger result set is requested, ensure that a temporary user table space is available. If there is none available, then create a table space. The following example creates a table space on a UNIX platform:

```
db2 "create user temporary tablespace tempts managed by system
    using ('/work/temptts.ts')"
```

### **scoringFlag**

0 means that there is no scoring and 1 means that there is scoring. If scoring is requested, an additional column with the score values is returned with the highest value first.

### **searchTermCountsFlag**

This controls the searchTermCounts processing. If searchTermCountsFlag is 0, searchTermCounts is not calculated.

## **Function parameters**

The following are output parameters.

### **searchTermCounts**

The number of occurrences of each search term query in the index. These counts are returned as a blank separated list in the order of search terms in the query.

See the **searchTermCountsFlag** for information.

### **totalNumberOfResults**

The total number of results found in the query result list.

Also note that when you use the STOP SEARCH AFTER, or the RESULT LIMIT together with the scoringFlag syntax in a query, this number is no longer reliable.

## **Usage**

The columns in the result set returned by the stored procedure are given by the CACHE TABLE option of the DB2TEXT CREATE INDEX command. If scoringFlag=1, then a column of type double is added. This column contains the SCORE value.

Use the following options to increase the performance of a second query with the same string as the first query. Note that this must be in a different cursor window with no totalNumberOfResults required:

- If you do not require scoring, add the following syntax: STOP SEARCH AFTER x DOCUMENTS, where x is the resultFirstRow + resultNumberRows.
- If you require scoring, add the following syntax: STOP SEARCH AFTER y DOCUMENTS, where y is equal to the totalNumberOfResults in the first query.

To ensure that you connect to the correct node for searching, it may be necessary to set the DB2NODE environment variable.

For UNIX, use the following command:

```
export DB2NODE=<no>
```

Note that it is important that all physical nodes have a synchronized time.

For Windows, use:

```
set DB2NODE= <no>
```



**Note:** A fenced user ID that is different from the instance owner ID does not work with partitioned databases.



---

## Chapter 54. Net Search Extender Messages

Note that the SQL states returned from the search function are 38600 plus the CTE error number.

---

### CTE0000 - CTE0499

---

**CTE0015W No dependent text-maintained table found requiring reset.**

**Explanation:** The RESET PENDING command executes a set integrity statement for dependent tables used to manage full-text search. No dependent table was found that matches these criteria for the given table.

**CTE0016W Failed to set integrity for *schema-name.table-name*.**

**Explanation:** The RESET PENDING command executes a set integrity statement for dependent tables used to manage full-text search. The command failed to complete successfully.

For more details see the topic called "RESET PENDING" in the DB2 Information Center.

**User response:** Run a "Set Integrity for <schema>.<table> IMMEDIATE CHECKED" for the failed table.

**CTE0017W Failed to set integrity for dependent tables.**

**Explanation:** The RESET PENDING command executes a set integrity statement for dependent tables used to manage full-text search. The command failed to complete successfully.

For more details, see the topic called "RESET PENDING" in the DB2 Information Center.

**User response:** Run a "Set integrity for <dependent-table> immediate checked" for any dependent tables for the base table that are still in pending mode.

**CTE0100E A DB2 operation failed. DB2 information: %2 %4.**

**Explanation:** A DB2 error occurred that does not allow further processing.

**User response:** For more detailed information on this DB2 error, use the following command: db2 ? SQLxxx.

**CTE0101E A search engine operation failed. Reason code: %2, %3, %4, %5, %6.**

**Explanation:** A Search Engine error occurred that does not allow further processing.

**User response:** For more detailed information, see the Search Engine reason code descriptions.

**CTE0102E A general system function failed. Error: %2.**

**Explanation:** A system error occurred that does not allow further processing.

**User response:** Additional information can be found on UNIX in the errno.h header file.

**CTE0103E An internal error occurred. Location: %1, %2.**

**Explanation:** An internal processing error that does not allow further processing. Try to start and stop the update and locking services, as well as DB2.

**User response:** If the error persists, start a trace and also check the db2diag.log.

**CTE0104E Memory allocation error (search engine).**

**Explanation:** The system has run out of memory.

**User response:** Increase the available memory size for the instance owner, or stop other processes running parallel.

**CTE0105E Memory allocation error.**

**Explanation:** The system has run out of memory.

**User response:** Increase the available memory size for the user, or stop other processes running parallel.

**CTE0106E Table %1.%2 has no primary key.**

**Explanation:** You tried to create an index on a table that does not have a primary key.

**User response:** Call the db2 alter table to ensure the existence of a primary key. Then try to create the index again.

**CTE0107E Directory %1 does not exist.**

**Explanation:** You specified a directory which does not exist.

**User response:** Create the directory, ensure accessibility to the instance owner. Then try to specify the directory again. Note that in a distributed DB2 environment, this directory has to exist on every physical node.

**CTE0108E The internal size %4 of the key columns on object %1.%2 is larger than maximum allowed size of %3.**

**Explanation:** The internal representation of the key columns exceeds the maximum size.

**User response:** Change the layout of the table before creating the index again. Use smaller key columns, which also benefit performance.

**CTE0109E The number of key columns %3 on object %1.%2 is larger than the allowed maximum of %4.**

**Explanation:** A maximum number of 14 key columns is supported.

**User response:** Change the layout of the table before creating the index again.

**CTE0110E The primary key on object %1.%2 is larger than allowed.**

**Explanation:** The primary key length exceeds the maximum size. The primary key length limit is based on the page size of the table space used by the table:

Max Key Length	Page size
1007	4K
2031	8K
4079	16K
4094	32K

Note that if the primary key consists of more than one column, the mentioned limits must be reduced by 2 bytes for each additional column.

**User response:** Change the layout of the table before creating the index again. Use smaller key columns, which also benefit performance.

**CTE0111E The file %1 is not readable.**

**Explanation:** The file specified cannot be read.

**User response:** Check the access rights for the file. Take into account that the Stored Procedure runs as a fenced user ID, which may also require rights to work on this file.

**CTE0112E The file %1 cannot be opened.**

**Explanation:** The file specified could not be opened.

**User response:** Verify that the file is correctly specified.

**CTE0113E Error converting model file %1 to UTF-8 encoding.**

**Explanation:** The specified CCSID or the default database CCSID does not match the model file CCSID.

**User response:** Ensure correct specification of the model file CCSID.

**CTE0114E Unable to register document model %1 in file %2.**

**Explanation:** The model file could not be used.

**User response:** Check that the model file syntax is correctly specified.

**CTE0115E A locking problem occurred. Lock Manager information: %1 %2.**

**Explanation:** An internal locking problem occurred.

**User response:** Check the current locks using the db2text control command. Using the same command, clean up the pending locks. If this does not help, stop and restart the locking and update services.

**CTE0116E Operation conflicts with existing lock.**

**Explanation:** You have tried to use a command that is currently not allowed when other commands are running on the index.

**User response:** Check the locks held on this index to see which commands are currently running. Wait until the other commands have finished. If the operation is no longer running but the lock is still active, clean up the locks for the index and try again.

**CTE0117E All available lock space for databases is used. Change the configuration.**

**Explanation:** You tried to work on more databases than are configured in your lock file.

**User response:** Change the number of databases you want to work in parallel with in your lock configuration db2extlm.cfg. Restart the update and locking services using the db2text stop and db2text start commands.

---

**CTE0118E** All available lock space for indexes on a databases is used. Change the configuration.

**Explanation:** You tried to work on more indexes for one database than are configured in your lock file.

**User response:** Change the number of indexes you want to work in parallel with in your lock configuration file db2extlm.cfg. Restart the update and locking services using the db2text stop and db2text start commands.

---

**CTE0119E** All available space for locks on an index is used.

**Explanation:** The operations you are running require more locks for one index than are configured in your lock configuration file.

**User response:** Change the number of locks you want to work in parallel with in your lock configuration db2extlm.cfg. Restart the update and locking services using the db2text stop and db2text start commands.

---

**CTE0120E** Update and locking services configuration file error.

**Explanation:** The configuration file db2extlm.cfg is in error.

**User response:** Check the db2extlm.cfg file and correct the error. Restart the update and locking services using the command db2text start.

---

**CTE0121E** The update and locking services configuration file cannot be opened.

**Explanation:** The file db2extlm.cfg could not be opened.

**User response:** Check if the file exists and that it can be accessed. If the file cannot be accessed, try to update your db2 instance using db2iupdt.

---

**CTE0122E** A syntax error was found in the update and locking services configuration file.

**Explanation:** A syntax error was found in the update and locking services configuration file.

**User response:** Check the update and locking services configuration file for errors.

---

**CTE0126E** The update and locking service input file %1 is corrupted.

**Explanation:** A required file for update and locking services is corrupted.

**User response:** Check if the file exists and if it can be accessed. If you can access the file, rename the file and restart the update and locking services. The file should

be created again. However, this action removes all of the specified frequency updates for create index.

---

**CTE0127E** An update and locking service error has occurred. Reason code: %1.

**Explanation:** An internal error has occurred in the update and locking service area.

**User response:** Stop DB2 and Net Search Extender and then clean up your shared resources. Try to start both again. If this does not work, report the problem to your IBM representative.

---

**CTE0129E** NULL values are not allowed to be passed as parameters.

**Explanation:** DB2 has passed a NULL value to an internal user-defined function.

**User response:** First make sure the specified base table has a primary key. Change your select statement to avoid this problem. Switch on the trace function and pass the returned information on to IBM Services.

---

**CTE0130E** The specified search argument exceeds the maximum length. The current search argument length is %1 and the maximum supported length is %2.

**Explanation:** The length of the specified search argument is %1. The maximum length must not exceed %2.

**User response:** Reduce the length of your search argument to %2.

---

**CTE0131E** The user-defined function %1.%2 does not exist.

**Explanation:** The specified user-defined function does not exist in this database.

**User response:** Check the name specified for this user-defined function, or register the user-defined function in the database you are using.

---

**CTE0132E** The text index %1.%2 does not exist.

**Explanation:** The specified text index does not exist in this database.

**User response:** Check the name specified and the database you are using. Use the db2ext.text columns view to see the existing text indexes.

---

**CTE0133E** The text index %1.%2 already exists.

**Explanation:** The text index that you specified already exists in this database.

**User response:** Check the name specified and the database you are using. Use the db2ext.text columns

view to see the existing text indexes.

---

**CTE0135E** The object %1.%2 does not exist.

**Explanation:** The specified object name does not exist in this database.

**User response:** Check the object name specified and the database you are using.

---

**CTE0136E** The column %1 does not exist in %2.%3.

**Explanation:** The specified column does not exist.

What to do

Check the column name that you specified. Check the table, view, or database you are using.

---

**CTE0137E** The table space %1 does not exist.

**Explanation:** The specified table space does not exist in this database.

**User response:** Check the name specified and the database you are using.

---

**CTE0138E** The table space %1 is not regular.

**Explanation:** The specified table space is not regular. The event table can only be created in a regular table space.

**User response:** Use this command again with a regular table space.

---

**CTE0139E** The environment variable %1 is not set.

**Explanation:** A required environment variable is not set.

**User response:** Check your environment, specify the required variable, and use the command again.

---

**CTE0140E** The database %1 is already enabled for text.

**Explanation:** The database you specified is already enabled for text.

**User response:** Check the name that you specified. Also check the DB2DBDFT variable that implies an implicit connection.

---

**CTE0141E** The database %1 is not enabled for text.

**Explanation:** The database you specified is not enabled for text.

**User response:** Check the database name you specified and the DB2DBDFT variable. If the database name is correct, use the command db2text enable database for text.

---

**CTE0142E** The command requires control authority on %1.%2 granted to user %3.

**Explanation:** You do not have the authority to use this command.

**User response:** Only the owner of this table can use this command or provide you with the required authorization.

---

**CTE0143E** The command requires database administration authority for user %1.

**Explanation:** You do not have the required authority to use this command.

**User response:** Only the owner of the database can use this command or provide you with the required authorization.

---

**CTE0144E** There is at least one text index active in database %1.

**Explanation:** You cannot disable your database until all text indexes are dropped.

**User response:** See the db2ext.text columns view for the existing indexes. Drop the existing indexes using the DROP INDEX command or specify the FORCE option with the DISABLE DATABASE command.

---

**CTE0145E** The CCSID %1 is not supported.

**Explanation:** The CCSID that you specified is not supported.

**User response:** Specify a valid CCSID.

---

**CTE0146E** The language %1 is not supported.

**Explanation:** The specified language is not supported.

**User response:** Specify a valid language.

---

**CTE0147E** The format %1 is not supported.

**Explanation:** The specified format is not supported.

**User response:** Specify a valid format.

---

**CTE0148E** The specified format %1 does not accept a model file.

**Explanation:** The format %1 does not support model files.

**User response:** Use a format that accepts a model file, or remove the model file from your command.

---

**CTE0149E** Too many terms (beginning with %1) are specified for the index update frequency.

**Explanation:** The syntax for the update frequency is not correct.

**User response:** Ensure that the DAY, HOUR, and MINUTE parameters are only specified once.

---

**CTE0150E** Unexpected end of command. Check the command syntax.

**Explanation:** The command syntax is not correct, or you are creating an index on a range partitioned table and you omitted the ADMINISTRATION TABLES IN clause.

**User response:** Check the command syntax. Verify that you specified the required parameters. If you are creating an index on a range partitioned table, specify the ADMINISTRATION TABLES IN clause.

---

**CTE0151E** Token %1 is unexpected. Check the command syntax.

**Explanation:** The syntax of the command is not correct.

**User response:** Check the command syntax and verify that the token you are using is allowed in the specific command.

---

**CTE0152E** Token %1 is too long.

**Explanation:** The token is too long.

**User response:** Check the command syntax and verify that the token is reduced to the maximum size allowed.

---

**CTE0153E** Token %1 occurs twice in the update frequency.

**Explanation:** You specified an incorrect syntax for the update frequency.

**User response:** Ensure that the DAY, HOUR, and MINUTE parameters are only specified once.

---

**CTE0154E** The value %1 for %2 is out of range. The valid range is %3 - %4.

**Explanation:** You specified an incorrect value. The value should be in the allowed range.

**User response:** Update your command. Change the value to match those in the allowed range.

---



---

**CTE0155E** The search string is empty.

**Explanation:** You specified an empty search string.

**User response:** Check that the search string includes valid alphanumeric characters.

---

**CTE0157E** Syntax error near %1.

**Explanation:** You specified an incorrect search syntax.

**User response:** Check the syntax near %1. Correct and try again.

---

**CTE0158E** The freetext search string is missing.

**Explanation:** Specify a freetext string.

**User response:** Check that the search string after "is about" includes valid alphanumeric characters.

---

**CTE0159E** Search string exceeds the allowed length of %1.

**Explanation:** The search string is too long.

**User response:** Reduce the size of the search string and try again.

---

**CTE0160E** No section name has been specified in the search string.

**Explanation:** You need to specify a valid section name.

**User response:** Add a valid section name and try again.

---

**CTE0162E** The escape command could not be processed.

**Explanation:** Your search string includes too many special characters that can be used as masking characters.

**User response:** Reduce the number of special characters in your search term, or avoid the escape command. The following special characters can be used: ! \* + , \_ . : ; { } ~ | ? [ ] ` = \

---

**CTE0163E** No thesaurus name specified in thesaurus clause.

**Explanation:** A thesaurus search is requested without a thesaurus name.

**User response:** Specify a thesaurus name in your search argument.

---

**CTE0164E** Syntax error in thesaurus relation %1.

**Explanation:** The specified syntax for the thesaurus relation is not correct.

**User response:** Update the thesaurus relation

---



according to the syntax specification.

---

**CTE0166E**    **Freetext must be the last statement in search query.**

**Explanation:** It is not allowed to have further operators after the "is about" token.

**User response:** Rewrite the query string. The last operator must be "is about".

---

**CTE0167E**    **Syntax error in free text query %1.**

**Explanation:** The syntax for the free text string is not correct.

**User response:** Update the free text string according to the syntax specification.

---

**CTE0168E**    **A left parenthesis in a section statement is missing.**

**Explanation:** The syntax for the section statement is not correct.

**User response:** Update the section statement according to the syntax specification.

---

**CTE0169E**    **A comma or right parenthesis is missing in a section statement.**

**Explanation:** The syntax for the section statement is not correct.

**User response:** Update the section statement according to the syntax specification.

---

**CTE0170E**    **A closing double quote is missing.**

**Explanation:** The specified syntax for the search term is not correct.

**User response:** Update the search term according to the syntax specification.

---

**CTE0171E**    **An open double quote for a section name is missing.**

**Explanation:** The syntax for the section statement is not correct.

**User response:** Update the section statement according to the syntax specification.

---

**CTE0172E**    **The closing double quote for the section name is missing.**

**Explanation:** The syntax for the section statement is not correct.

**User response:** Update the section statement according to the syntax specification.

---

**CTE0173E**    **One escape character must be defined in an escape clause.**

**Explanation:** There can be no more than one character in an escape clause.

**User response:** Remove the additional characters in the escape clause.

---

**CTE0174E**    **A blank character is not allowed as an escape character.**

**Explanation:** It is not allowed to have a blank character in an escape clause.

**User response:** Change the escape clause to a clause with a valid character.

---

**CTE0175E**    **An escape clause is defined but no mask character is found in the search phrase.**

**Explanation:** An escape clause is specified without using a mask character.

**User response:** Remove the escape clause.

---

**CTE0176E**    **The succeeding character of an escape character in the phrase is neither the same character nor a mask character.**

**Explanation:** The character after the escape character must be either a masking character or the escape character itself.

**User response:** Change the search string to correctly use the escape character.

---

**CTE0177E**    **The number value %1 is invalid.**

**Explanation:** The specified number in the search argument is not valid.

**User response:** Check the documentation about the valid range. Update the value in the search argument.

---

**CTE0178E**    **Mask characters in fuzzy phrase must be preceded by an escape character.**

**Explanation:** Masking together with fuzzy search is not allowed.

**User response:** Update the search string with an escape character.

---

**CTE0179E**    **Thesaurus name %1 exceeds allowed length of %2.**

**Explanation:** Primary keys longer than 60 bytes are not supported.

**User response:** Change the layout of the table before creating the index again.



---

**CTE0180E** Thesaurus %1 can not be found.

**Explanation:** The thesaurus specified cannot be found.

**User response:** Check that the thesaurus files are located in the thesaurus directory or fully qualified.

---

**CTE0181E** Library %1 cannot be loaded.

**Explanation:** A library cannot be found.

**User response:** Check that the library is located in the library path and available. Start and stop DB2 to ensure that the current settings are used.

---

**CTE0182E** Function %1 cannot be loaded from library %2.

**Explanation:** A library entry point cannot be loaded.

**User response:** The library accessed seems to be invalid. Check that the library is specified only once.

---

**CTE0183E** Error occurred using shared system resources.

**Explanation:** A request to share system resources like shared memory or semaphores cannot be fulfilled.

**User response:** Check the current system status and configuration. On UNIX, use the `ipcs` command to check the resources. Stop all applications, such as DB2 and Net Search Extender. If further resources are listed, clean them up using `ipcrm`.

---

**CTE0184E** No db2text start command was issued.

**Explanation:** A command was called which requires the locking and update services.

**User response:** Start the update and locking services with `db2text start`.

---

**CTE0184N** No DB2TEXT START command was issued.

**Explanation:** A Net Search Extender (NSE) command was called which requires the NSE instance services to be started. You might get this messages if the service is not started.

In a multiple partitioned database environment:

- NSE instance services must be started on all database partitions.
- For more information about starting NSE instance services in a multiple partitioned database environment, see the topic called "Starting Net Search Extender (NSE) instance services in a multiple partitioned database environment".

**User response:** Use the `DB2TEXT START` command to start NSE instance services.

---

**CTE0185E** The Net Search Extender (NSE) instance services are already active.

**Explanation:** A `db2text start` is issued but the update and locking services are already running.

**User response:** No further action.

---

**CTE0185N** The update and locking services are already active.

**Explanation:** This message can be returned when the `db2text start` command is issued when the update and locking services are already running.

**User response:** No further action is required.

---

**CTE0186E** Update and locking service error occurred, check the db2diag log file for details.

**Explanation:** An update and locking service error occurred.

What to do

Check the `db2diag` log file for further information, or clean up your shared resources. See also CTE0183E.

---

**CTE0187E** Update and locking services are still active, use FORCE option to stop the services.

**Explanation:** The `db2text stop` command has not stopped the locking services, there are still processes running.

**User response:** Check with `db2text control` which processes are running and wait for those to finish. If you need to stop them, use the `FORCE` option.

---

**CTE0188E** There is a temporary problem using update and locking services. Please try again.

**Explanation:** The `db2text stop` command has not stopped the locking services. Programs are still running or an inconsistent situation is found.

**User response:** Check with `db2text control` which processes are running and wait for those to finish. To stop them, use the `FORCE` option.

---

**CTE0189E** The executable program %1 cannot be found.

**Explanation:** The program file cannot be located or accessed.

**User response:** Check if the program file is located in the `bin` or `adm` directory of the DB2 server. The installation is corrupt if the file cannot be found.

---

---

**CTE0190E** The executable program %1 cannot be started.

**Explanation:** The program cannot be started.

**User response:** Check if the program is located in the bin or adm directory of the DB2 server and that the appropriate libraries are installed. For further information, call the program manually on the server.

---

**CTE0191E** The drop index operation is incomplete. Check the db2diag log files for details.

**Explanation:** The drop index operation is incomplete, possibly caused by the FORCE option.

**User response:** Using the FORCE option drops everything regardless of any errors. Check the index directory for pending files and remove these manually.

---

**CTE0192E** Errors occurred in an update index operation. Check event table %1. %2 and db2diag log files for details.

**Explanation:** During the index update process, any document errors are written to the event table.

**User response:** Check the event table for more information about the document errors. Clean up the event log after the problems have been fixed.

---

**CTE0194E** The type %1 of column %2 is not supported.

**Explanation:** You used a column that is not in the list of the supported ones.

**User response:** Check CREATE INDEX for a list of valid columns for keys and indexing. Make the appropriate changes to the command and try again.

---

**CTE0195E** %1 is not an absolute path.

**Explanation:** An absolute path on the server is required.

**User response:** Check the path and write an absolute path in the command.

---

**CTE0198E** No corresponding text index.

**Explanation:** There is no text index on the column.

**User response:** Check if the text index still exists.

---

**CTE0199E** There is no text index corresponding to column *column-number* of table *table-name*.

**Explanation:** You tried to search on a column without a text index.

**User response:** Check the column you are searching on, or create a text index on the column.

---

**CTE0199N** No text index corresponding to column *column-number* of table *.table-name*.

**Explanation:** A text search query on a table column failed because it has no valid and active text search index.

**User response:** Check the column you are searching on, or create a text index on the column. If the table column has an invalid text search index, drop the text index and create a new one.

For example, a text search index gets invalidated if the underlying base table is dropped and recreated.

---

**CTE0200E** At least one command option must be specified.

**Explanation:** The ALTER INDEX command changes the characteristics of an index, such as the update and storage options. None of the characteristics to be changed was specified.

**User response:** Specify at least one command option. Refer to the command syntax for all possible options.

---

**CTE0201E** There is a conflict with an existing text index on the same column.

**Explanation:** A text index defined on the same column was created with different parameters from this create index command.

**User response:** Correct the parameter values in the create index command. Make sure that following parameters have the same value for the existing index and the index to be created: ccsid, language, format, document model, index configuration, column function, and attributes.

---

**CTE0202E** The object %1.%2 must be a view when key columns are specified.

**Explanation:** The specified object is not a view. The KEY COLUMNS FOR INDEX ON VIEW clause is only allowed when indexing a column of a view.

**User response:** Remove the KEY COLUMNS FOR INDEX ON VIEW(SQL-columnname-list) clause.

---

**CTE0203E** The text index %1.%2 was not created with the CACHE TABLE option. This is required for command execution.

**Explanation:** This command can only be executed if the specified index was created with the CACHE TABLE option.

**User response:** Create an index with CACHE TABLE option. Refer to the documentation for the command syntax.

---

**CTE0204E** An attribute name is missing. Add "AS <attribute name>" to the attribute expression.

**Explanation:** Whenever a column expression is used in the attribute expression, an attribute name must be supplied. For example: (C1+C2 AS myname).

**User response:** Add "AS <attribute name>" to the attribute expression.

---

**CTE0205E** CACHE TABLE expressions are not valid.

**Explanation:** The column list in the cache table expression is not valid.

**User response:** Correct the cache table column list in the create index command. Make sure the columns exist in the specified table. If a function is applied on a column, verify that it is used correctly.

---

**CTE0206E** ATTRIBUTE expressions are not valid.

**Explanation:** The column list in the attribute expression is not valid.

**User response:** Correct the attribute column list in the create index command. Make sure the columns exist in the specified table. If a function is applied on a column, verify that it is used correctly.

---

**CTE0207E** KEY COLUMNS FOR INDEX ON VIEW not specified for index on view %1.%2.

**Explanation:** If indexes on views are created, the KEY COLUMNS FOR INDEX ON VIEW(SQL-columnname-list) clause must be specified. The list of column names specifies the columns that UNIQUELY identify a row in the view.

**User response:** Include the KEY COLUMNS FOR INDEX ON VIEW(SQL-columnname-list) clause in the create index command.

---

**CTE0208E** INITIAL SEARCH RESULT ORDER columns are not valid.

**Explanation:** The column list in the INITIAL SEARCH RESULT ORDER(SQL-order-by list) expression is not valid.

**User response:** Correct the order by column list in the create index command. Check if the syntax is correct and the columns exist in the specified table. If a function is applied on a column, verify that it is used correctly.

---

**CTE0209E** The type %1 of attribute column %2 is not supported, type DOUBLE is required.

**Explanation:** For attribute columns, the only supported data type is DOUBLE.

**User response:** Make sure the attribute columns of the table with the text column to be indexed are of type DOUBLE. It may be possible to use cast operators in attribute column expressions. Refer to the SQL Reference for data types which can be cast to double.

---

**CTE0210E** The value %1 for index configuration parameter %2 is not valid. A valid value is %3.

**Explanation:** The specified value for the configuration parameter is incorrect. For valid values of the parameters refer to the command syntax.

**User response:** Correct the index configuration parameter value in the create index command.

---

**CTE0211E** %1 is not a valid index configuration parameter.

**Explanation:** The index configuration option is not known.

**User response:** Check the create index command syntax. Valid index configuration options are TreatNumbersAsWords and IndexStopWords. These have to be comma separated: index configuration(treatnumberaswords 1, indexstopwords 1).

---

**CTE0212E** Internal index configuration file %1 could not be saved.

**Explanation:** The internal configuration file for the index could not be saved.

**User response:** Make sure the instance owner has write permissions to the directory the file should be saved in. If a file with the same name already exists, make sure that it is writable for the instance owner.

---

**CTE0213E** Internal index configuration file template %1 could not be loaded.

**Explanation:** The internal index configuration file template could not be read.

**User response:** Make sure the file exists in the correct location and is readable.

---

**CTE0214E** Internal error when setting new entry [%1],%2=%3 for index configuration file.

**Explanation:** Internal error while writing an internal configuration file for the index.

**User response:** If the file exists, check if it is readable and writable for the instance owner. Check that there is enough space on the device where the file is located.

**CTE0215E** Index creation on alias %1.%2 is not supported. Use base table %3.%4 instead.

**Explanation:** The index cannot be created on the alias.

**User response:** Type in the create index command with the base table.

**CTE0217E** The schedule service is already active.

**Explanation:** The service is already active, you do not need to start it.

**User response:** No action required.

**CTE0218E** The function *function-name* failed with error code *error-code*.

**Explanation:** A Windows function failed with the specified error code which does not allow further processing.

**User response:** Use the specified Windows system error code to get detailed error information.

**CTE0218N** Function *function-name* failed with error code *error-code*.

**Explanation:** A Windows operating system function failed with the specified error code which does not allow further processing.

**User response:** Use the specified Windows system error code to get detailed error information.

**CTE0219E** The service named *service-name* could not be opened. Error code *error-code*.

**Explanation:** The specified service cannot be found on the Windows system.

**User response:** Check if the specified service is installed on the Windows system. Use the specified Windows system error code to get detailed error information.

**CTE0219N** The service *service-name* could not be opened. Error code *error-code*.

**Explanation:** The specified service cannot be found on the Windows operating system.

This message can be returned when the database manager attempted to start the Net Search Extender (NSE) services in response to the DB2TEXT START command, but could not start the services on one or more of the database partitions.

For example, if Net Search Extender (NSE) is not

installed on one or more computers where the database partitions are located, then the database manager will not be able to successfully start Net Search Extender (NSE) instance services.

**User response:** Use the specified Windows system error code to get detailed error information.

**CTE0220E** The DB2 instance profile path could not be found.

**Explanation:** Internal DB2 function to obtain the DB2 instance profile path failed.

**User response:** Create a DB2 instance without specifying the instance profile path information and retry the command.

**CTE0221E** UpdateFrequency %1 is incorrectly specified.

**Explanation:** The syntax for the update frequency statement is not correct.

**User response:** Correct the update frequency statement according to the syntax specification.

**CTE0222E** The schedule service input file %1 is corrupted.

**Explanation:** The scheduler file containing index update information is corrupted.

**User response:** Use your system editor and try to correct the problem. Maybe an entry has been truncated, or the ending line character has been deleted. If this does not restore the file content, try the following: Call command db2text stop to stop the scheduler. Delete the scheduler service file. Call command db2text start to start the scheduler. Use command db2text alter index ... to recreate the update frequency entries for all concerned indexes.

**CTE0223E** File %1 could not be closed.

**Explanation:** The file specified cannot be closed.

**User response:** Verify that the file is correctly specified.

**CTE0224E** File %1 could not be copied to %2.

**Explanation:** The first file cannot be copied to the second file.

**User response:** Verify that the files are correctly specified. Check if the second file already exists and is read only. Also check if there is enough free space on the system.

---

**CTE0225E** File %1 could not be removed.

**Explanation:** The file specified cannot be removed from the system.

**User response:** Verify that the file is specified correctly and check the file access rights.

---

**CTE0227E** A write operation on file %1 failed.

**Explanation:** The file specified is not writable.

**User response:** Verify that the file is correctly specified and check the file access rights. Also check if there is enough free space on the system.

---

**CTE0228E** The user has insufficient access rights at the operating system level.

**Explanation:** The command requires administrator rights at the operating system level.

**User response:** Ensure that you have operating system administrator rights. Check if you are a member of the administrator group.

---

**CTE0231E** %1 is not defined in same nodegroup (%4) as the table space of %2.%3.

**Explanation:** The table space of the administration tables is required to be distributed over different nodes in exactly the same way as the table containing the text column to be indexed. To enforce this, it is checked whether the specified tablespace is defined in the same nodegroup.

**User response:** Specify a table space that is defined in the same nodegroup as the table containing the text column to be indexed.

---

**CTE0232E** The specified or default table space %1 is not single-noded. This is necessary for an index on a view, or when the CACHE TABLE option is specified.

**Explanation:** An index on a view or with the CACHE TABLE option enabled is only supported for tables on a single node.

**User response:** Put the table in a single-noded table space if the default table space caused this error. Alternatively, specify another single-noded table space, if you specified a multi-noded table space.

---

**CTE0233E** There is a conflicting administration command running. Please retry this command later.

**Explanation:** Another administration command is still running or terminated abnormally without releasing the command lock.

**User response:** Check with CONTROL LIST which

locks are still active. If there is an active lock but no command running, clear the lock manually using the CONTROL CLEAR command. Be aware that someone else may be running the administration command holding the lock.

---

**CTE0234E** There is a conflicting administration command running on a text index. Please retry this command later, or specify the FORCE option of a DISABLE DATABASE command.

**Explanation:** Another administration command is still running or terminated abnormally without releasing the command lock.

**User response:** Check with CONTROL LIST which locks are still active. If there is an active lock but no command running, clear the lock manually using the CONTROL CLEAR command. Be aware that someone else may be running the administration command holding the lock. For a DISABLE DATABASE command you may specify the FORCE option which stops all other commands on that database.

---

**CTE0235E** No valid license found for Net Search Extender.

**Explanation:** There was no valid license found for Net Search Extender.

**User response:** Check whether the license was correctly installed with db2lic. Make sure existing instances are updated after the product install.

---

**CTE0236E** Only Node0 is supported on MPP instances.

**Explanation:** Text Indexes can only be created on MPP instances, if the table with the text column to be indexed resides on Node0.

**User response:** Check the node group of the tablespace in which the table is defined.

---

**CTE0237E** Internal error: log table %1.%2 contains an invalid operation %3.

**Explanation:** The log table keeps track of operations executed on the table containing the indexed text column. This table might be corrupted, as it contains an entry not written by Net Search Extender.

**User response:** Check the log table and delete the corrupted entry.

---

**CTE0238E** Internal error: table %1.%2 contains an incorrect syntax expression in column %3.

**Explanation:** There is an error in the expression list in the specified text column.



**User response:** Check the delimiter Begin and End pairs.

**CTE0239E Internal error: total length of index properties %1 exceeds maximum %2.**

**Explanation:** The maximum size of the index properties (1016 bytes) is exceeded. The properties contain the instance, index, and work directory as well as other information.

**User response:** Make sure these path names are not too long.

**CTE0240E Internal error: setting environment variable %1 failed.**

**Explanation:** Setting the specified environment variable failed. There may be a problem with the environment setup.

**User response:** Check your OS specific guidelines.

**CTE0242E Value %1 for parameter %2 is invalid.**

**Explanation:** The search stored procedure or the table valued function DB2EXT.TEXTSEARCH was called with invalid parameters.

**User response:** Correct the parameter values of the search stored procedure or table valued function. For valid parameters refer to the documentation.

**CTE0243E The cache for text index %1.%2 has not been activated.**

**Explanation:** A Net Search Extender operation requires an activated cache. The cache is currently not activated. These are the possible reasons: The cache has never been activated after the last DB2TEXT START command. The cache has been explicitly deactivated with the DB2TEXT DEACTIVATE CACHE command.

**User response:** Perform a DB2TEXT ACTIVATE CACHE command for the index and rerun the Net Search Extender operation.

**CTE0244E Internal error: call to %1 returns rc=%2, SQLCODE=%3.**

**Explanation:** An internal processing error occurred when calling an internal function.

**User response:** If the error persists, start a trace and check the db2diag.log. Report the error.

**CTE0245E The requested cache size exceeds the available cache size. Increase the maximum cache size to a value > %1 or decrease the pctfree value.**

**Explanation:** The cache size necessary to load all data exceeds the MAXIMUM CACHE SIZE value for an

index. This can be detected during activation of the cache (the DB2TEXT ACTIVATE command), or by an index update operation while the cache is activated.

**User response:** If the error was reported in a DB2EXT ACTIVATE command, recalculate the maximum cache size using the DB2EXT.MAXIMUM\_CACHE\_SIZE function and alter the MAXIMUM\_CACHE\_SIZE setting for the index. Eventually decrease the PCTFREE value. If the maximum number of documents is exceeded during incremental update, rebuild the cache with the commands db2 deactivate cache and db2text activate cache recreate.

**CTE0246E File %1 is empty.**

**Explanation:** A DB2TEXT CREATE INDEX command failed because the document model file specified in the command is empty.

**User response:** Specify a valid document model file in the command.

**CTE0247E A Net Search Extender stored procedure could not be created.**

**Explanation:** A DB2TEXT ENABLE DATABASE command failed to create the internal stored procedure DB2EXT.CTESRVSP.

**User response:** Check the additional DB2 error message associated with a CREATE PROCEDURE statement for details. If the error cannot be corrected by removing an existing stored procedure with an identical name, start a trace and report the error.

**CTE0248E The generated search string is too long. Reduce the complexity of search query.**

**Explanation:** A Net Search Extender query is too long or too complex to be processed by the base search engine. The complexity is affected by thesaurus expansions, FUZZY FORM OF expressions, and masking characters.

**User response:** Reduce complexity or length of the query.

**CTE0249E The executable program named *program-name* terminated abnormally.**

**Explanation:** When executing a Net Search Extender command, the executable *program-name* was called, but terminated abnormally.

**User response:** Verify that the executable was not terminated explicitly by user interaction, for example, a signal. If not, start a trace, rerun the command, and report the error.

---

**CTE0249N Executable program *program-name* terminated abnormally.**

**Explanation:** While executing a Net Search Extender (NSE) command, the database manager ran the named executable program. That program terminated abnormally.

**User response:**

1. Verify, that the executable was not terminated explicitly by user interaction.
2. If the program was not terminated by a user, rerun the command that failed.
3. If the problem persists:
  - Turn on the DB2 trace utility to collect diagnostic information.
  - Rerun the command that failed
  - If the command fails again, contact DB2 support with the collected diagnostic information.

---

**CTE0250E The return type %1 of column type transformation function %2.%3 is not supported.**

**Explanation:** In a DB2TEXT CREATE INDEX command a column type transformation was specified that returns an unsupported datatype. Supported datatypes are: CHARACTER, VARCHAR, LONG VARCHAR (deprecated), CLOB, GRAPHIC, VARGRAPHIC, LONG VARGRAPHIC (deprecated), DBCLOB, BLOB, and XML.

**User response:** Choose a different column type transformation function.

---

**CTE0251E Internal error: the column type %1 is not supported.**

**Explanation:** A column type is used that is not in the list of supported types.

**User response:** Check create index for a list of valid columns for Keys and Indexing. Make the appropriate changes to the command and try again. If the error persists, start a trace and also check the db2diag.log. Report the error to IBM Services.

---

**CTE0252E The parameter %1 is missing.**

**Explanation:** Internal error - when executing a Net Search Extender command, an administration executable program was called with a missing parameter "%1".

**User response:** Try to change Net Search Extender parameter commands to avoid the problem. If the error persists, switch on the trace function and report the error to IBM Services.

---

**CTE0253E The document listed in the log view was not found.**

**Explanation:** The contents of a text document that is listed in the log view has changed and could not be accessed.

**User response:** Check that the document exists and the read/access permissions of the text documents to be included in the index.

---

**CTE0254E The cache for index %1 is already activated.**

**Explanation:** The index has already been activated with the ACTIVATE CACHE command.

**User response:** Check the specified index name and the database that you are using.

---

**CTE0255E A column name for a cache result column expression is missing. Add "AS <cache column name>" to the expression.**

**Explanation:** A cache result column expression must be named. For example: 'C1+C2 AS myresult'.

**User response:** Add "AS <cache column name>" to the expression.

---

**CTE0256E The query necessary to select data for indexing failed. Reduce the complexity of the attribute, cache table, or the initial search result order expressions.**

**Explanation:** Net Search Extender creates a query from the expressions in your command to select data for indexing from the database. The query failed because it was too complex.

**User response:** Reduce the complexity of attribute, cache table, or initial search result order expressions.

---

**CTE0257E : Error creating shared memory.**

**Explanation:** The shared memory resource could not be created due to a previous error or permission problem.

**User response:** Check db2diag.log for further information, or clean up your shared resources. See also error CTE0183E.

---

**CTE0258E Shared memory version error.**

**Explanation:** The shared memory resource could not be accessed because it is corrupted or there is a version conflict.

**User response:** Check db2diag.log for further information. Disable and re-enable the database and then try again.

---

**CTE0259E** Cannot insert entry in global shared memory. Entry already exists.

**Explanation:** An entry to be inserted in global shared memory already exists because of a previous error.

**User response:** Check db2diag.log for further information. Restart the update and locking services using the commands db2text stop and db2text start.

---

**CTE0260E** Cannot access entry in global shared memory. Entry not found.

**Explanation:** An entry to be removed from global shared memory does not exist because of a previous error.

**User response:** Check db2diag.log for further information. Try to restart the update and locking services using the commands db2text stop and db2text start.

---

**CTE0261E** There is at least one cache activated for a text index in this instance. Deactivate the cache for any activated index using the DEACTIVATE CACHE command, or use the FORCE option to stop.

**Explanation:** The db2text stop command can only be used if you run a DEACTIVATE CACHE command for all text indexes that have been activated with the ACTIVATE CACHE command.

**User response:** Deactivate the cache for any activated index using the DEACTIVATE CACHE command or use the FORCE option to stop.

---

**CTE0262E** The value for parameter %1 is too long.

**Explanation:** The value exceeds the maximum allowable size.

**User response:** Check the maximum size.

---

**CTE0263E** The text index %1.%2 was created with the RECREATE INDEX ON UPDATE option. In this context, the UPDATE MINIMUM or COMMITCOUNT FOR UPDATE may not be specified.

**Explanation:** Update minimum and commitcount for update are only effective if the index is updated incrementally.

**User response:** If you want to recreate the index each time an update is performed, remove the UPDATE MINIMUM and COMMITCOUNT FOR UPDATE settings. If you want to use UPDATE MINIMUM and COMMITCOUNT FOR UPDATE, do not specify RECREATE INDEX ON UPDATE.

---

**CTE0264E** Errors occurred in an activate index operation. Check event view %1.%2 and the db2diag.log for details.

**Explanation:** During the index activate process, errors are written to the event table and the db2diag.log file.

**User response:** Check the event table for more information about the document errors. Clean up the event log after the problems have been fixed.

---

**CTE0265E** The table space of a user table or administration table space (%1) is not only defined on node 0.

**Explanation:** If text indexes are created on MPP instances, the table space of the user table must only reside on Node0.

**User response:** Use a table where the table space resides on Node0.

---

**CTE0266E** ValueFrom %1 must be smaller than ValueTo %2.

**Explanation:** The values specified in the attribute search are not valid. If the search syntax is 'BETWEEN ValueFrom AND ValueTo', the lower boundary (ValueFrom) must be smaller than upper boundary(ValueTo).

**User response:** Change the boundaries in the 'BETWEEN ValueFrom AND ValueTo' clause.

---

**CTE0267E** The Net Search Extender database objects in the database %1 are in an inconsistent state.

**Explanation:** At least one Net Search Extender object is missing or corrupted. Either the database has not been migrated after installation of a new Net Search Extender product version, or a database user has changed or dropped Net Search Extender internal object(s). In this case, all text indexes are lost and the database has to be disabled for text.

**User response:** For a database migration to the current version please follow the migration description found in the Net Search Extender documentation. Alternatively, issue a DB2TEXT DISABLE DATABASE command using the FORCE option. You can then enable the database for text again by using the DB2TEXT ENABLE DATABASE command.

---

**CTE0270E** Logtable %1.%2 could not be modified after incremental update. Entries are to be processed during the next UPDATE.

**Explanation:** When starting an incremental index update, a timestamp is created. This serves as a threshold for change records to be processed. Changes occurring concurrently to the incremental update are then processed later, during the next update. In certain



situations, there can be changes in transactions that are uncommitted at the time the update starts, but are committed while the index update is being performed. This may potentially lead to inconsistencies.

To avoid such an inconsistent situation, the change records prior to the threshold timestamp are not deleted from the logtable, although they have been partially processed. On the next incremental update the changes will be re-applied to the index.

**User response:** On the next index update the changes are re-applied to the index. In case of delete operations, this can lead to the following error: CTE0101E: ItlEnReasonCode\_Docmap\_ docid\_not\_found.

Note that this error can be ignored, as the document was already deleted. If CTE0270E errors frequently occur, consider dropping and re-creating the index with a modified timestamp threshold for incremental index update. For example: db2text "CREATE INDEX ... INDEX CONFIGURATION(UPDATEDDELAY 30)"

This means that processing during an incremental update run only changes records older than 30 seconds and avoids interference with concurrent change transactions of less than 30 seconds.

---

**CTE0271E Cache not usable, DEACTIVATE and ACTIVATE RECREATE required.**

**Explanation:** The cache is in an inconsistent state because the maximum cache size has been reached.

**User response:** Check that the maximum cache size is still sufficient. Then call the following db2text commands: DEACTIVATE CACHE and ACTIVATE CACHE RECREATE.

---

**CTE0272E Insufficient cache size. Increase the PCTFREE value or use DEACTIVATE and ACTIVATE [RECREATE] to recreate the cache.**

**Explanation:** All reserved memory for the cache has been used.

**User response:** Rebuild the cache by using the following sequence of db2text commands: DEACTIVATE CACHE, ALTER INDEX MAXIMUM CACHE SIZE, and ACTIVATE CACHE RECREATE.

---

**CTE0273E The cache for index %1, %2 is already activated.**

**Explanation:** The index has already been activated with the ACTIVATE CACHE command.

**User response:** Check the specified index name and the database that you are using.

---

**CTE0274E The target database system %1 for the connection is not supported.**

**Explanation:** You tried to execute a DB2TEXT command with a connection to a database system that is not supported by Net Search Extender.

---

**CTE0275E The type and version information for server %2 could not be found.**

**Explanation:** The type and version information for the server could not be found in the DB2 catalog view 'SERVERS'.

**User response:** Make sure that the DB2 federated environment is set up correctly.

---

**CTE0277E A cache memory segment could not be attached.**

**Explanation:** The system cannot allocate enough memory to load a large cache segment, or the cache segment cannot be opened because it has been previously deleted.

**User response:** Check your system settings and increase the amount of paging space and free memory. For large cache sizes you may need to prepare your system. Refer to the Net Search Extender documentation. Use the DEACTIVATE and ACTIVATE [RECREATE] commands to recreate the cache. If the problem persists, check db2diag.log for additional information.

---

**CTE0278E On an AIX 32-bit system, change the MAXDATA setting before activating a large cache.**

**Explanation:** When you use the search stored procedure on an AIX 32-bit system, you may need to change the MAXDATA setting for the db2fmp executable.

**User response:** Refer to the Net Search Extender documentation for details about changing the MAXDATA setting.

---

**CTE0279E The size of the cached data has reached a system limit.**

**Explanation:** By decreasing the PCTFREE value, you can increase the maximum data size during cache activation. This enables the system to reserve less freespace in the cache.

**User response:** Use a lower PCTFREE value or reduce your amount of text data to be cached. Use the DEACTIVATE and ACTIVATE [RECREATE] commands to recreate the cache.

---

**CTE0280E** There is not enough disk space to write persistent cache files.

**Explanation:** The system can not write a large enough file for persistent cache in the cache directory.

**User response:** Change the persistent cache directory to an empty file system by using the ALTER INDEX command. Alternatively, reduce the cache size by decreasing the PCTFREE or MAXIMUM CACHE SIZE values or by using a temporary cache.

---

**CTE0281E** Deletion of persistent cache file %1 has failed.

**Explanation:** The file does not exist or cannot be accessed.

**User response:** Check if this file still exists and delete it manually.

---

**CTE0282E** The number of documents in the cache has reached a system limit.

**Explanation:** By decreasing the PCTFREE value, you can increase the maximum number of document entries to be cached during cache activation. This enables the system to reserve less freespace in the cache.

**User response:** Use a lower PCTFREE value or reduce the amount of document entries in the cache. Use the DEACTIVATE and ACTIVATE [RECREATE] commands to recreate the cache.

---

**CTE0283E** A cache memory segment could not be created.

**Explanation:** The system cannot allocate enough memory for loading a large cache segment into memory. By decreasing the PCTFREE value, you achieve a smaller cache segment size.

**User response:** Check your system settings and increase the amount of paging space and free memory. You can also decrease the cache size by using a lower PCTREE value. For large cache sizes, you may need to prepare your system. Refer to the Net Search Extender documentation. Use the DEACTIVATE and ACTIVATE [RECREATE] commands to recreate the cache. If the problem persists, check db2diag.log for additional information.

---

**CTE0284E** The text index is located on node %1, but the search function was called on node %2.

**Explanation:** The search stored procedure or table valued function DB2EXT.TEXTSEARCH was not called on the node where the index is located. The search function will not automatically be distributed to the correct node.

**User response:** Set the DB2NODE environment

variable to the node where the index is connected before connecting to the database.

---

**CTE0285E** Search function is not allowed for a text index which is distributed to multiple nodes.

**Explanation:** The table valued function DB2EXT.TEXTSEARCH must not be called with indexes that are distributed to multiple nodes, since it will not be automatically distributed to the correct nodes, but executed on the coordinator node.

**User response:** Use the CONTAINS, SCORE or NUMBEROFMATCHES function in a multiple node environment.

---

**CTE0286E** No row found in %1."IBMSNAP\_REGISTER" for source table %2.%3 and capture change table %4.%5.

**Explanation:** No valid entry was found in the IBMSNAP\_REGISTER table for the replication capture table characteristics specified in the DB2TEXT CREATE INDEX command. A valid entry must contain the specified source table for the index incolumns SOURCE\_OWNER and SOURCE\_NAME, with SOURCE\_VIEW\_QUAL=0 and the specified replication capture table in columns PHYS\_CHANGE\_OWNER and PHYS\_CHANGE\_TABLE.

The specified source table was not registered as a replication source for the replication capture table.

**User response:** Register the source table correctly for DB2 Replication, or specify a correct replication capture table for the source table.

---

**CTE0287E** Invalid value %1 for %2 in "%3"."IBMSNAP\_REGISTER" for source table %4.%5 and capture change table %6.%7.

**Explanation:** A replication setting found in the IBMSNAP\_REGISTER table is not allowed. Possible causes: 1.The column CHG\_UPD\_TO\_DEL\_INS does not contain the value 'Y'. 2.The column CCD\_CONDENSED contains the value 'Y'.

**User response:** When registering the source table for DB2 Replication, ensure that update operations are transformed into pairs of delete and insert operations. In addition, ensure that no condensed replication capture tables are used.

---

**CTE0288E** Source table %1.%2 and capture change table %3.%4 are on different servers (%5 and %6).

**Explanation:** The specified source table and replication capture table must reside on the same server.

---

**CTE0289E The wrapper %1 is not supported.**

**Explanation:** The wrapper is not supported. Refer to the Net Search Extender documentation for a list of supported wrappers.

---

**CTE0290E The alias %1.%2 is not allowed in the replication clause.**

**Explanation:** You are not allowed to specify an alias for a nickname in a replication clause.

**User response:** Specify the nickname instead of the alias, or create a new nickname for the remote table.

---

**CTE0291E The specified format is not allowed for a column of type XML.**

**Explanation:** For columns of type XML only the 'XML' format is allowed.

**User response:** Specify format 'XML' or none.

---

**CTE0292E Windows Exception %1 was caught, address=%2, flags=%3.**

**Explanation:** A windows exception occurred. Exception name, address and flags are provided.

---

**CTE0293E Windows exception %1 was caught.**

**Explanation:** A Windows exception occurred.

---

**CTE0294E Search argument processing problem.**

**Explanation:** A search argument processing error occurred due to an incorrect environment setup.

**User response:** Check that the locale charmap value matches the DB2 code page and is available on the system.

---

**CTE0295E Invalid CCSID %1 specified for non-binary text column.**

**Explanation:** For non-binary text column data types, DB2 always stores the data in the database CCSID. Only the database CCSID is valid for non-binary text columns.

**User response:** Omit the CCSID-clause or specify a valid CCSID.

---

**CTE0296E The library %1 cannot be found on %2. Please check the Net Search Extender installation.**

**Explanation:** You tried to execute a Net Search Extender command by using the DB2 Control Center. Net Search Extender is not installed correctly on the target system.

**User response:** Check if Net Search Extender is

installed correctly on the target system.

---

**CTE0297E The database %1 is associated with DB2 Text Search**

**Explanation:** You are not allowed to associate a database with more than one text search component.

**User response:** If you want to continue using DB2 Text Search with the database, then no action is required. If you would rather use DB2 Net Search Extender, disable the database from DB2 Text Search and try this command again.

---

**CTE0298E The database %1 is already associated with DB2 Text Search. This command cannot be executed.**

**Explanation:** You are not allowed to associate a database with more than one text search component.

**User response:** If you do not want to use DB2 Net Search Extender, disable the database using the DB2TEXT DISABLE DATABASE command. If you would rather use DB2 Net Search Extender, try this command again after disabling the database with DB2 Text Search.

---

**CTE0300N User *user-id* does not have authority to create text index on the table *schema-name.table-name*.**

**Explanation:** "CREATE INDEX .. FOR TEXT" requires one of the following:

- DBADM authority
- CONTROL privilege on the table
- INDEX privilege on the table with either IMPLICIT\_SCHEMA privilege on the database or CREATEIN privilege on the index schema.

For more information about DB2 Net Search Extender create index, see the topic called "Net Search Extender CREATE INDEX command" in the DB2 Information Center.

**User response:** Ensure that the user has authority to perform the create index operation.

---

**CTE0301E The instance owner does not have the authority to perform the Net Search Extender command.**

**Explanation:** The user attempted to run a text index command without the instance owner id holding the appropriate authority

Different DB2 Text Search index commands require varying levels of database authority. For a list of required authorities for DB2 Net Search Extender commands, see the topic called "DB2 Net Search Extender command authority requirements" in the DB2 Information Center.

**User response:** Ensure that the instance owner has DBADM with DATAACCESS authority in order for the user to perform Net Search Extender operations.

---

**CTE0302E** User *user-id* does not have authority to update the DB2 Net Search Extender index named *index-schema-name.index-name*.

**Explanation:** Update text index requires one of the following:

- DATAACCESS authority
- CONTROL privilege on the table.

For more information about DB2 Net Search Extender update, see the topic called "Net Search Extender UPDATE INDEX command" in the DB2 Information Center.

**User response:** Ensure that the user has authorization to perform the update index operation.

---

**CTE00303N** Table column *column-name.table-name* has no valid and active text index.

**Explanation:** A text search query on a table column failed because it has no valid and active text index.

**User response:** Ensure that the table column has a valid and active text index before performing a text search query.

If the table column has an invalid text index, drop the text index and create a new one.

For example, a text index gets invalidated if the underlying base table is dropped and recreated.

---

**CTE0304N** Index *schema-name.index-name* is not valid and cannot be used.

**Explanation:** The text index is invalid and cannot be updated or used.

**User response:** Drop and recreate the invalid text index.

---

**CTE0305E** Failed to retrieve the text-maintained dependent tables for *schema-name.table-name*.

**Explanation:** The command executes a query to retrieve information from the text index administration tables. Querying the text index information failed.

**User response:** Ensure that the text index administration tables are accessible. For more information see the db2diag.log.

---

**CTE0306N** The specification of a COMMITCOUNT value is not allowed for a text index created with the REPLICATION CAPTURE TABLE option.

**Explanation:** If the REPLICATION CAPTURE TABLE option is specified the COMMITCOUNT value cannot be specified for the index.

**User response:** Retry the command without the COMMITCOUNT option for the index since REPLICATION CAPTURE TABLE option is being used.

---

**CTE0307N** The table specified in the replication clause must not be a view.

**Explanation:** The viewname has been specified in the replication clause instead of the table name.

**User response:** Retry command with the view name replaced with the table name for the replication clause.

---

**CTE0308N** A text index on a view cannot be created using the REPLICATION CAPTURE TABLE option.

**Explanation:** The viewname has been specified in the replication clause instead of the table name.

**User response:** Retry command with the view name replaced with the table name for the replication clause.

---

**CTE0309W** Search result truncated by mask limit.

**Explanation:** The result set for the search was truncated due to the mask resolution limit.

**User response:** To get the complete set of results for the search query, increase the mask resolution limit and re execute the search query.

---

**CTE0310N** Index update is not allowed due to insufficient disk space.

**Explanation:** You are not allowed to update index since the available space for index update is less than the estimated disk space required.

**User response:** Check the available free space for the filesystem that holds index and work directory for the index. Increase the available space on the filesystem and try update index again.

---

**CTE0360E** A specific error message displays

**Explanation:** This message number is reserved for specific Net Search Extender errors.

**User response:** Respond to the specific error message.

---

---

**CTE0451E** The specified document format %1 is not supported by the highlighting UDF.

**Explanation:** The document format %1 does not support highlighting.

**User response:** Use a document format that is supported by the highlighting UDF.

---

**CTE0452E** Syntax error near option %1 in the highlighting UDF.

**Explanation:** You specified an incorrect syntax near the specified option.

**User response:** Check the syntax near option %1. Correct and try again.

---

**CTE0453E** The return size of the highlighting UDF is too small.

**Explanation:** The requested parts of the highlighted document does not fit into the return parameter of the highlighting UDF.

**User response:** Decrease the window number, the window size and/or the number of sections from which hits should be displayed. This will reduce the document parts returned to the user.

---

**CTE0454E** Error converting the parameters of the highlighting UDF from codepage %1 to codepage UTF8.

**Explanation:** The parameters of the highlighting UDF in the specified CCSID (which may be the default database CCSID), cannot be converted to UTF8.

**User response:** Ensure correct specification of the CCSID.

---

**CTE0455E** The database codepage %1 is not supported in the highlighting UDF.

**Explanation:** The database has a codepage which is not supported by the highlighting UDF.

---

**CTE0456E** The highlighting UDF only supports documents in codepage UTF8.

**Explanation:** Only documents in codepage UTF8 support the highlighting UDF.

---

**CTE0457E** The value %1 for parameter %2 is not valid in the highlighting UDF.

**Explanation:** A value for a highlighting parameter is not valid.

**User response:** Check the parameter value and ensure that the value is allowed in the data range.

---

---

**CTE0458E** Usage: db2exthl <new size in kilo bytes>.

**Explanation:** The parameter for the db2exthl utility is not correct.

**User response:** Provide a value between 1 and 1048576.



---

## CTE0500 - CTE0999

---

### CTE0841E Missing command option %1.

**Explanation:** A required command option was not specified.

**User response:** Check the specified parameters and add the missing parameter.

### CTE0842E No value is specified for the command option %1.

**Explanation:** A required value for a command option was not specified.

**User response:** Check the specified parameters and add the missing option.

### CTE0843E No numeric value is specified for the command option %1.

**Explanation:** A string instead of a number has been specified.

**User response:** Check the specified parameters and change the string to the correct number.

### CTE0844E The definition file path %1 is too long.

**Explanation:** The specified path is too long and could not be processed.

**User response:** Use a shorter path and try again.

### CTE0845E No definition file is specified.

**Explanation:** The definition file needs to be specified.

**User response:** Add a valid definition file and try the call again.

### CTE0846E The definition file name %1 is too long.

**Explanation:** The specified definition file name is too long.

**User response:** Reduce the length of the definition file name to the size allowed.

### CTE0847E The definition file %1 does not exist.

**Explanation:** The specified definition file could not be found.

**User response:** Check that the definition file is in the correct path and can be accessed by the current user.

### CTE0849E The dictionary file %1 could not be locked.

**Explanation:** The process was not able to lock the dictionary file. Either you do not have write access, or

another process has opened the file for writing.

**User response:** Check the running processes to ensure that no process is locking the dictionary file and check your access rights.

### CTE0850E Output file %1 already exists.

**Explanation:** The specified output file could not be overwritten.

**User response:** Check that you are able to create the thesaurus in the specified directory.

### CTE0851E The integrity of the dictionary file %1 is lost.

**Explanation:** The thesaurus dictionary files are corrupted.

**User response:** Clean up the directory and compile your definition file again.

### CTE0852E Dictionary file %1 version error.

**Explanation:** Your dictionary file was generated with an older version of the thesaurus compiler.

**User response:** Compile your definition file again with the current version of the thesaurus compiler.

### CTE0853E The existing dictionary %1 cannot be overwritten.

**Explanation:** An existing dictionary cannot be overwritten.

**User response:** Check your write access right on the dictionary file, its directory location and subdirectory location.

### CTE0855E A thesaurus term is incorrectly specified.

**Explanation:** There is a syntax error in your definition file.

**User response:** Check your Net Search Extender documentation for information on creating a thesaurus definition file and thesaurus support.

### CTE0856E The definition file %1 is empty.

**Explanation:** An empty definition file is not allowed.

**User response:** Check your Net Search Extender documentation for information on creating a thesaurus definition file and thesaurus support.

---

---

**CTE0857E** No block starting line found in file *file-name* at line *line-number*.

**Explanation:** There is a syntax error in your thesaurus definition file.

**User response:** A block has to start with ':WORDS'. Check your Net Search Extender documentation for information on thesaurus concepts.

---

**CTE0858E** An invalid relationship is specified in file *%1* at line *%2*.

**Explanation:** There is a syntax error in your definition error.

**User response:** You have to examine your 'associated-term-definition'. Check your Net Search Extender documentation for information on creating a thesaurus definition file.

---

**CTE0859E** The relationship number is out of range in file *%1* at line *%2*.

**Explanation:** The user-defined relations are all based on the associative type. They are identified by unique numbers between 1 and 128.

**User response:** Verify your relationship numbers.

---

**CTE0860E** No terms are defined in file *%1* at line *%2*.

**Explanation:** Required terms are not specified.

**User response:** Check your Net Search Extender documentation for information on creating a thesaurus definition file.

---

**CTE0861E** The thesaurus term in file *%1* at line *%2* is too long.

**Explanation:** The length of the thesaurus term is restricted to 64 bytes.

**User response:** Alter the size of your thesaurus term and try again.

---

**CTE0862E** Strength is incorrectly specified in file *%1* at line *%2*.

**Explanation:** There is a syntax error in your definition file.

**User response:** Check your Net Search Extender documentation for information on creating a thesaurus definition file and thesaurus support.

---

**CTE0863E** Strength is out of range in file *%1* at line *%2*.

**Explanation:** The strength value should be specified between 1 and 100.

**User response:** Change the strength value so that it is a numerical value from 1 to 100.

---

**CTE0864E** Internal error: Thesaurus compiler failed with reason code *%1*.

**Explanation:** An internal processing error occurred that does not allow further processing. Try to start and stop the update and locking services, as well as DB2.

**User response:** If the error persists, start a trace and also check the db2diag.log.

---

**CTE0865E** The directory *%1* could not be created.

**Explanation:** The specified directory could not be created.

**User response:** Check if the directory already exists and the permissions on the directory.

---

**CTE0866E** The directory *%1* could not be removed.

**Explanation:** The directory could not be removed.

**User response:** Check that you have write permissions on the specified directory.





---

## Chapter 55. Windows system errors

The following is a list of Windows system errors:

### System errors

- 1 Incorrect function.
- 2 The system cannot find the file specified.
- 3 The system cannot find the path specified.
- 4 The system cannot open the file.
- 5 Access is denied.
- 6 The handle is invalid.
- 8 Not enough storage is available to process this command.
- 14 Not enough storage is available to complete this operation.
- 15 The system cannot find the drive specified.
- 29 The system cannot write to the specified device.
- 30 The system cannot read from the specified device.
- 32 The process cannot access the file because it is being used by another process.
- 36 Too many files opened for sharing.
- 38 Reached the end of the file.
- 39 The disk is full.
- 80 The file exists.
- 82 The directory or file cannot be created.
- 100 Cannot create another system semaphore.
- 101 The exclusive semaphore is owned by another process.
- 102 The semaphore is set and cannot be closed.
- 103 The semaphore cannot be set again.
- 104 Cannot request exclusive semaphores at interrupt time.
- 105 The previous ownership of this semaphore has ended.
- 110 The system cannot open the device or file specified.
- 111 The file name is too long.
- 112 There is not enough space on the disk.
- 121 The semaphore timeout period has expired.
- 126 The specified module could not be found.
- 127 The specified procedure could not be found.
- 147 Not enough resources are available to process this command.
- 155 Cannot create another thread.

- 161 The specified path is invalid.
- 164 No more threads can be created in the system.
- 170 The requested resource is in use.
- 183 Cannot create a file when that file already exists.
- 187 The specified system semaphore name was not found.
- 206 The filename or extension is too long.
- 267 The directory name is invalid.
- 288 Attempt to release mutex not owned by caller.
- 298 Too many posts were made to a semaphore.
- 998 Invalid access to memory location.
- 1051 A stop control has been sent to a service that other running services are dependent on.
- 1052 The requested control is not valid for this service.
- 1053 The service did not respond to the start or control request in a timely fashion.
- 1054 A thread could not be created for the service.
- 1055 The service database is locked.
- 1056 An instance of the service is already running.
- 1057 The account name is invalid or does not exist, or the password is invalid for the account name specified.
- 1058 The service cannot be started, either because it is disabled or because it has no enabled devices associated with it.
- 1059 Circular service dependency was specified.
- 1060 The specified service does not exist as an installed service.
- 1061 The service cannot accept control messages at this time.
- 1062 The service has not been started.
- 1063 The service process could not connect to the service controller.
- 1064 An exception occurred in the service when handling the control request.
- 1066 The service has returned a service-specific error code.
- 1067 The process terminated unexpectedly.
- 1068 The dependency service or group failed to start.
- 1069 The service did not start due to a logon failure.
- 1070 After starting, the service hung in a start-pending state.
- 1071 The specified service database lock is invalid.
- 1072 The specified service has been marked for deletion.
- 1073 The specified service already exists.
- 1078 The name is already in use as either a service name or a service display name.

- 1079 The account specified for this service is different from the account specified for other services running in the same process.
- 1082 No recovery program has been configured for this service.
- 1154 One of the library files needed to run this application is damaged.
- 1219 The credentials supplied conflict with an existing set of credentials.
- 1242 The service is already registered.
- 1243 The specified service does not exist.
- 1244 The operation being requested was not performed because the user has not been authenticated.
- 1245 The operation being requested was not performed because the user has not logged on to the network. The specified service does not exist.
- 1392 The file or directory is corrupted and unreadable.
- 1455 The paging file is too small for this operation to complete.
- 1793 The user's account has expired.



---

## Chapter 56. Net Search Extender information catalogs

Net Search Extender stores important information about defaults, configurations, text indexes, and formats in catalog tables. To view this information, you can query some views on tables.

The following views and tables reflect the current configuration of your system:

- Database level information views:
  - db2ext.dbdefaults
- Index level information views:
  - db2ext.textindexes
  - db2ext.textindexformats
  - db2ext.indexconfiguration
- Table views for a text index:
  - Event view
  - Log table view
  - Staging table view

---

### Views for database-level information

The view db2ext.dbdefaults displays all the default values for the Net Search Extender database.

The defaults on the database level can be changed and are available as attribute-value pairs in this view:

db2ext.dbdefaults

```
db2 select DEFAULTNAME, DEFAULTVALUE from DB2EXT.DBDEFAULTS
```

*Table 8. db2ext.dbdefaults view*

Attribute	Default Value	Notes
CCSID	CCSID of database	Default CCSID for documents. This is applied if no CCSID is specified in the CREATE INDEX command.
FORMAT	TEXT	Document default format. This is applied if no format is specified in the CREATE INDEX command.

Table 8. db2ext.dbdefaults view (continued)

Attribute	Default Value	Notes
INDEXDIRECTORY	See the path name under Notes	Directory for full-text index files. This is applied if no index directory is specified in the CREATE INDEX command.  For Linux and UNIX operating systems:  INSTHOME/sqllib/db2ext/indexes where INSTHOME is the instance owner home directory.  For Windows operating systems:  INSTPROFDIR\instance-name\db2ext\ indexes where INSTPROFDIR is the instance profile directory
LANGUAGE	EN_US	The document language.
MODELCCSID	CCSID of database	CCSID of document model files.
UPDATECOMMITCOUNT	0	The number of changes processed in one transaction during an update.
CLEARCOMMITCOUNT	0	The number of changes processed in one transaction during an CLEAR INDEX command.
UPDATEFREQUENCY	NONE	When to check for updates in new indexes.
UPDATEMINIMUM	1	The minimum number of changes before update is executed.
WORKDIRECTORY	See the path name under Notes	Directory for index temporary files.  For Linux and UNIX operating systems:  INSTHOME/sqllib/db2ext/indexes where INSTHOME is the instance owner home directory.  For Windows operating systems:  INSTPROFDIR\instance-name\db2ext\ indexes where INSTPROFDIR is the instance profile directory
CACHEDIRECTORY	See the path name under Notes	Default directory for PERSISTENT CACHE option of the CREATE INDEX command.  For Linux and UNIX operating systems:  INSTHOME/sqllib/db2ext/indexes where INSTHOME is the instance owner home directory.  For Windows operating systems:  INSTPROFDIR\instance-name\db2ext\ indexes where INSTPROFDIR is the instance profile directory
PCTFREE	50	The percentage of the cache left free for future inserts.

Table 8. db2ext.dbdefaults view (continued)

Attribute	Default Value	Notes
USERPERSISTENTCACHE	1	Use the persistent cache.
AUTOMATICREORG	1	The REORGANIZE option in the CREATE INDEX command. This means automatic reorganization.
TREATNUMBERSASWORDS	0	Do not interpret sequences of characters and numbers as separate words, even if they are adjacent characters. For example, the 0 default means that tea42at5 is considered as one word.
INDEXSTOPWORDS	1	Index all text including stop words.
VERSION		NSE V9.7 Current version number of Net Search Extender.
UPDATEDELAY	0	Specifies the duration in seconds for incremental update without capture tables. Only entries older than this duration will be taken from the log table. This is to avoid lost updates. For example, document changes that are not reflected in the index in transaction scenarios where user transactions interfere with update commands. Therefore, the UpdateDelay parameter should be set to the maximum duration of a user write transaction on the table the index was created on.
AUXLOGNORM	OFF	Do not enable the extended text-maintained staging infrastructure for non-partitioned tables by default. The staging infrastructure can be enabled for a text index with explicit index configuration AUXLOG ON.
AUXLOGPART	ON	By default, enable the extended text-maintained staging infrastructure for range-partitioned tables. The staging infrastructure can be disabled for a text index with explicit index configuration AUXLOG OFF.
LOCKSCHEULERFILE	0	Set to '1' to block concurrent write-access to the scheduler file ctedem.dat. In some scenarios with a high degree of parallelism to create, drop, or alter update frequencies, update processes may otherwise be started unnecessarily.

**Note:** On Windows operating systems the default index directory has changed. In the DB2 Net Search Extender Version 9.5, the value was DB2-installation-path\db2ext\indexes. See Directory structure for your installed DB2 database product (Windows) for the DB2 installation path

---

## Views for index-level information

You can query information at an index-level using the following Net Search Extender views:

- `db2ext.textindexes`
- `db2ext.textindexformats`
- `db2ext.indexconfiguration`
- `<index eventview name schema>.<index eventview name>`

For backward compatibility reasons, DB2 Text Information Extender views `db2ext.textcolumns`, `db2ext.formats`, and `db2ext.models` are still supported, but deprecated.

Note that in the `db2ext.textcolumns` view the `OPERATION`, `OPERATIONBEGIN`, and `OPERATIONEND` columns are no longer supported.

### **db2ext.textindexes view**

Each database enabled for Net Search Extender contains a `db2ext.textindexes` view. This contains information on settings, statistics, and defaults for the created text indexes in this database.

When you create a text index, new entries are created in `db2ext.textindexes`. When you drop the text indexes, these entries are deleted.

You can query the view to obtain information about the indexes. This is an example using the index schema:

```
db2 "select COLNAME from DB2EXT.TEXTINDEXES where INDSHEMA='myschema'
    and INDNAME='myindex'"
```

Note, however, that you cannot modify the view using normal SQL data manipulation commands, or explicitly create or drop the catalog view. Additional contents of the view are found in the following table.

Also note that the replication parameters are not included in this view.

*Table 9. db2ext.textindexes view*

Attribute	Type	Notes
INDSCHEMA	VARCHAR(128)	Schema name of the text index.
INDNAME	VARCHAR(128)	Name of the text index.
TABSCHEMA	VARCHAR(128)	The table name of the schema for base tables, nicknames, and views.
TABNAME	VARCHAR(128)	Alias name the index was created on.
COLNAME	VARCHAR(128)	Column the index was created on.
CCSID	INTEGER	Document CCSID for this index.
LANGUAGE	VARCHAR(5)	Document language for this index.
FUNCTIONSCHEMA	VARCHAR(128)	Schema of the column mapping function.
FUNCTIONNAME	VARCHAR(18)	Name of the column mapping function.
INDEXDIRECTORY	VARCHAR(256)	Directory for full-text index files.
WORKDIRECTORY	VARCHAR(256)	Directory for index temporary files.



Table 9. db2ext.textindexes view (continued)

Attribute	Type	Notes
CACHEDIRECTORY	VARCHAR(256)	Directory for persistent cache (if persistentcache=1).
UPDATEFREQUENCY	VARCHAR(300)	Trigger criterion for applying automatic updates to this index.
UPDATEMINIMUM	INTEGER	Minimum number of documents that must be changed before an update is performed.
EVENTVIEWSHEMA	VARCHAR(128)	Schema of the event view created for this index.
EVENTVIEWNAME	VARCHAR(128)	Name of the event view created for this index.
LOGVIEWSHEMA	VARCHAR(128)	Schema of the log view created for an index.
LOGVIEWNAME	VARCHAR(128)	Name of the log view created for an index (important for incremental update on views).
COMMITCOUNT	INTEGER	Default for commitcount update.
NUMBER_DOCS	INTEGER	Total number of documents currently in the index. Note that during an index update, this value is only updated if the commitcount is set.
REORG_SUGGESTED	INTEGER	Indicates if performance can be improved by running UPDATE INDEX REORGANIZE. This parameter is only true (1) if at least one of the nodes has an index reorganization suggested.
REORGAUTOMATIC	INTEGER	1, if the index gets automatically reorganized during the update operation.
RECREATEONUPDATE	INTEGER	1, if the index gets automatically reorganized during the update operation.
CREATIONTIME	TIMESTAMP	Time of index creation.
UPDATETIME	TIMESTAMP	Time of last update. If UPDATE TIME is equal to CREATION TIME, then no update has been processed.
PERSISTENTCACHE	INTEGER	1, if persistent cache is used.
MAXIMUMCACHESIZE	INTEGER	Maximum size of cache.
PCTFREE	INTEGER	Percentage of cache left free for future inserts.
CACHETABLE	VARCHAR(32000)	Column expression list for the CACHE TABLE.
RESULTORDER	VARCHAR(32000)	SQL-order-by for INITIAL RESULT ORDER.
ATTRIBUTES	VARCHAR(32000)	Column expression list for ATTRIBUTES.
VIEWKEYCOLUMNS	VARCHAR(32000)	Key columns for index on view.

Table 9. db2ext.textindexes view (continued)

Attribute	Type	Notes
AUXSTAGINGSHEMA	VARCHAR(16)	Schema for text-maintained staging tables; set to SYSIBMTS
AUXSTAGINGNAME	VARCHAR(48)	System-generated name of the text-maintained staging table (only when configured).

## db2ext.indexconfiguration view

Index configuration parameters are available in the db2ext.indexconfiguration view. The view is available through normal SQL query facilities. This is an example using the index name:

```
db2 "select VALUE from DB2EXT.INDEXCONFIGURATION where INDSHEMA='myschema'
and INDNAME='myindex' and PARAMETER ='INDEXSTOPWORDS'"
```

Additional contents of the view are found in the following tables.

Table 10. db2ext.indexconfiguration view

Attribute	Type	Notes
INDSHEMA	VARCHAR(128)	Schema name of the index.
INDNAME	VARCHAR(128)	Name of the index.
PARAMETER	VARCHAR(30)	Type of parameter.
VALUE	VARCHAR(512)	Value of the parameter.

For the PARAMETER and VALUE attributes, there are several values available.

Table 11. db2ext.indexconfiguration view

Attribute and values	Attribute and values
PARAMETER	VALUE
- TREATNUMBERASWORDS	- 0 or 1
- INDEXSTOPWORDS	- 0 or 1
- UPDATEDELAY	- seconds >= 0
AUXLOGPART	- ON or OFF
AUXLOGNORM	- ON or OFF

For further information, see the CONFIGURATION option of the CREATE INDEX command.

## db2ext.textindexformats view

Format and model information for indexes is available in the db2ext.textindexformats view. This is an example using the index name:

```
db2 "select FORMAT from DB2EXT.TEXTINDEXFORMATS where INDSHEMA='myschema'
and INDNAME='myindex'"
```

Additional contents of the view are found in the following table.

Table 12. *db2ext.textindexformats* view

Attribute	Type	Notes
INDSCHEMA	VARCHAR(128)	Schema name for the index (used as prefix for indexname and schemaname in the log table).
INDNAME	VARCHAR(128)	Index name specified in CREATE INDEX command.
FORMAT	VARCHAR(30)	The model is bound to this format.
MODELNAME	VARCHAR(30)	The name of a document model.
MODELFILE	VARCHAR(256)	File containing the model definition.
MODELCCSID	INTEGER	CCSID of MODELFILE.
DEFAULT	INTEGER	Currently 1, as multiple formats in an index are not currently supported.

## Table views for a text index

You can query information at an index level using these DB2 Net Search Extender views:

- Event view
- Log table view

### Event view

This view allows you to get information on indexing status and error events, and when problems occur during indexing, for example, a document not being found. These index update events are then written to the index's event table.

This view allows you to get information on indexing status and error events, and when problems occur during indexing, for example, a document not being found. These index update events are then written to the index's event table.

The schema and name are stored in the `db2ext.textindexes` view. To get the name of the event view, use the following example:

```
db2 "select EVENTVIEWSHEMA, EVENTVIEWNAME from DB2EXT.TEXTINDEXES
     where INDSCHEMA = 'myschema' and INDNAME = 'myindex'
```

The event view for an index consists of the following columns.

Table 13. *The event view*

Attribute	Type	Notes <sup>®</sup>
OPERATION	INTEGER	The operation on the user table is be reflected in full-text index (insert = 0/ update = 1/ delete = 2).  When using a replication capture table, update operations are split into a delete and an insert operation. In this case, an insert operation in the event table can be either from an insert, or an update operation on the source table the index was created on.

Table 13. The event view (continued)

Attribute	Type	Notes®
TIME	TIMESTAMP	Timestamp of event entry creation.
REASON	INTEGER	Reason code. For a list of reason codes, see Chapter 57, "Text Search Engine reason codes," on page 287.
SEVERITY	INTEGER	The severity of the table entry. For example, 1 is for information purposes, 4 indicates a warning, and 8 means a table entry error.
MESSAGE	VARCHAR(1024)	Additional text information.
KEY1, ... KEY14	Depends on the user table	First primary key column of user table to the last primary key column (a maximum of 14).
PARTITION	INTEGER	The database partition number on which this error occurs. In a non-partitioned database environment, this is 0.

The events can be cleared by the DB2TEXT CLEAR EVENTS command.

**Note:** Informational events, such as starting, committing, and finishing update processing are also available in this view.

In this case, Key1, ... Key14 and OPERATION all have NULL values.

In the case of indexes on views, the PK01, ..., PK14 columns relate to the columns specified in the KEY COLUMNS clause of the CREATE INDEX command.

## Log tables, views, and nicknames

The purpose of the log table is to store the change operations on the user table or view that require synchronization with the external full-text index.

For indexes created on regular tables or nickname tables, there are triggers created on the user table to feed change information into the log table. However, if replication capture tables are used, no log table is created and the replication capture table is used instead.

For log tables, the update command reads the entries and deletes them after successful synchronization.

However, in the case of indexes on views, triggers cannot fill the log table. As you can update the view, the user is responsible for this task.

Table 14. The log table view

Attribute	Type	Notes
OPERATION	INTEGER	The type of change on the user table that requires index synchronization: (0 = insert, 1 = update, 2 = delete).
TIME	TIMESTAMP	The timestamp for the creation of a row in this table.

Table 14. The log table view (continued)

Attribute	Type	Notes
PK01 ... PKnm	Same as user table	In case of errors, the column where the problem occurred. They are a copy of the primary key columns of the user table or the equivalent key columns in case of an index on a view.

The user who creates the table is able to select, update, insert, and delete this view.

Table 15. The text-maintained staging table

Attribute	Type	Notes
PK 1..n	Same as user table	A copy of the primary key definition(s) of the user table.
Globaltransid	CHAR(8)	Internal transaction id
Globaltranstime	CHAR(13)	Timestamp
Operationtype	Integer	Insert 1
		Delete -1

The content of the table is maintained by the text search, however an administrator may delete entries in the table.

If you specify a replication capture table in the create index command, no log table is created and the replication capture table is used instead. The replication capture table must contain the following columns:

Table 16. The replication capture table

Attribute	Type	Notes
IBMSNAP_OPERATION	INTEGER	The type of change on the CD or CCD table that requires index synchronization: (I = insert, U= update, D= delete).  When using a replication capture table, update operations are split into a delete and an insert operation. In this case, an insert operation in the event table can be either from an insert, or an update operation on the source table the index was created on.
IBMSNAP_COMMITSEQ	CHAR	Maps to the corresponding column of the CD or CCD table.
IBMSNAP_INTENTSEQ	CHAR	Maps to the corresponding column of the CD or CCD table.
PK01 ... PKnm	Same as user table	In case of errors, the column where the problem occurred. They are the primary key columns of the user table.

The user who defines the table is able to select, update, insert, and delete with the grant option.



---

## Chapter 57. Text Search Engine reason codes

- 0 Operation performed successfully - no error occurred.
- 1 An invalid handle was passed to a function.
- 2 Function could not allocate enough memory.
- 3 Function could not perform due to access limitations or security restrictions.
- 4 The operation is not supported for this version of the Text Search Engine run time.
- 5 The operation is currently not enabled.
- 6 The application violated the Text Search Engine protocol by calling Text Search Engine functions in illegal order.
- 7 An unexpected error occurred. Please report this to your service representative.
- 8 An invalid language was specified.
- 9 The specified language is valid but not supported by the Text Search Engine run time.
- 10 An invalid CCSID was specified.
- 11 The specified CCSID is valid but not supported by the Text Search Engine run time.
- 12 An invalid document ID was specified.
- 13 The specified document format is valid but not supported by the Text Search Engine run time.
- 14 An invalid document format was specified.
- 15 The operation could not succeed due to access limitation during file input/output.
- 16 The operation could not succeed due to read errors during file input/output.
- 17 The operation could not succeed due to read errors during file input.
- 18 The operation could not succeed due to write errors during file output.
- 19 The operation could not succeed due to seek errors during file input/output.
- 20 The operation could not succeed due to tell errors during file input/output.
- 21 The operation could not succeed due to close errors during file input/output.
- 22 The operation could not succeed due to errors during rename operations.
- 23 The operation could not succeed due to errors during remove operations.
- 24 The operation could not succeed due to errors during mkdir operations.
- 25 One or more function arguments did have an invalid value (for example, an unexpected null pointer or an invalid enumeration type value).

- 26 The specified directory does not exist.
- 27 An unexpected Text Search Engine error occurred. Please examine the Text Search Engine error code in the error info object for further details.
- 28 An unexpected COS error occurred. Please report this error.
- 29 An attempt was made to update an empty document.
- 30 The specified argument is not supported for this operation.
- 31 The date attribute parser found an invalid value when trying to parse a date attribute.
- 32 The number attribute parser found an invalid value when trying to parse a number attribute.
- 33 Attribute name invalid, probably too long.
- 35 Reserve number for future use.
- 36 The input document contains an attribute (DATE, NUMBER, or STRING) that exceeds the length limit for attributes. The attribute text has been truncated to that limit.
- 38 The warning threshold as set by the user has been exceeded. As a consequence, this error has been generated.
- 39 The input document could not be indexed. It contains too many nested fields.
- 40 The limit of different attributes for one of the attribute types has been exceeded for this index.
- 46 The iterator is not (or no longer) valid, because its list is empty or deleted.
- 47 The function is not supported for the passed kind of handle. This error occurs, for example, when trying to use `itlQueryResultEntryObtainData` on a list iterator that does not represent a query result iterator.
- 48 This warning is issued if a stop word file cannot be found for the specified language and resource path.
- 49 This warning is issued if a stop word file does not contain any stop words.
- 50 This warning is issued if a stop word file does contain invalid data.
- 100 The index could not be opened because it does not exist with the specified name and/or directory.
- 101 The specified index name is not a valid index name.
- 102 The specified index directory is not a valid directory name.
- 103 The operation cannot be performed because the Text Search Engine detected a corruption of the index structure and/or index file sets.
- 104 The specified index cannot be created because it already exists with the given name and directory.
- 109 Before any other operation can be performed on this index, a rollback operation must be performed.
- 110 The index configuration file does not contain the mandatory section as specified in the error context.
- 111 The index configuration file does not contain the mandatory option as specified in the error context.



- 112 The index configuration file contains invalid data in the option as specified in the error context.
- 113 The index configuration file does not match the Text Search Engine version.
- 200 The specified document model name is not a valid model name.
- 201 The specified document model field name is not a valid field name.
- 202 The specified document model is not known.
- 203 The specified document model already exists and cannot be redefined.
- 204 Too many or too large document models have been added to the index.
- 205 The document model contains too many elements.
- 206 The document model element contains a parameter (XML attribute) that is not allowed for this type of element.
- 207 The document model element contains a parameter value that is not allowed for this type of parameter (XML attribute).
- 208 The document model element does not contain a required parameter (XML attribute), like "name".
- 209 The document model does not seem to be XML, or starts with an unexpected XML element.
- 210 The given XPath (locator value) contains an unexpected token.
- 211 The given XPath (locator value) contains an unexpected axis specifier (name followed by two colons).
- 212 The given XPath (locator value) contains an unexpected node test.
- 213 The document model directory file (extension .mdx) is corrupted.
- 214 The document model index file (extension .mox) is corrupted.
- 215 The document contains an XML element which is mapped to a document attribute and which contains another document attribute. The inner attribute is ignored.
- 216 The given parameter value is too long as a GPP or HTML tag.
- 217 The document model contains a duplicate field definition.
- 218 The document model contains a duplicate attribute definition.
- 300 The operation cannot be performed because the Text Search Engine detected a corruption in the index files used for document name mapping.
- 301 The operation cannot be performed because the Text Search Engine detected an invalid document number.
- 302 The operation cannot be performed because the Text Search Engine detected an invalid document identifier.
- 303 The operation cannot be performed because the Text Search Engine found no index entry for the document identifier.
- 304 The operation cannot be performed because the Text Search Engine found no index entry for the document number.
- 305 The operation cannot be performed because the Text Search Engine detected an overflow in used document numbers.

- 306 The document identifier that the application tried to index has appeared already in the list of documents. The Text Search Engine does not support duplicate document identifiers appearing in one indexing sequence, that is, before the update has been committed.
- 340 The term strength is not valid.
- 341 The relation number is not valid, must be in.
- 342 The relation type is invalid, use one of the defines described in API.
- 343 The phrase (term) is too long.
- 344 Unexpected end of file encountered while reading.
- 345 Version conflict detected when reading index/thesaurus files.
- 346 Overflow in thesaurus buffers.
- 347 Invalid name, probably too long a name, for file or directory.
- 348 Lookup did not find term (phrase) in dictionary or entry in definition file does not contain mandatory term.
- 349 Definition file is empty.
- 350 Thesaurus dictionary or definition file as specified via input parameter does not exist.
- 351 Syntax errors in definition file.
- 352 The Relationship was specified incorrectly.
- 352 The Relationship number was out of range.
- 360 An invalid single character masking was used.
- 361 An invalid multiple character masking was used.
- 362 Operator arity is smaller than number of operands given in query.
- 363 Operator value out of range defined by ItlEnOperator enumeration.
- 364 Value for rank formula out of enumeration range.
- 365 Number identifying proximity segment is out of range.
- 366 Query is under construction and cannot be redefined or reset.
- 367 Scope given as previous search result denotes empty result.
- 368 Invalid call requesting to add field names before setting the first one.
- 369 Invalid search flag requesting an invalid comparison with index content is ignored. If, for example, a case-sensitive comparison was requested for an index that was build in a case insensitive manner, this reason code is shown in the error information.
- 370 Masking of strings is not supported for Thai or DBCS languages.
- 371 No valid query input. For example, the search terms is available.
- 372 Invalid comparison operations requested.
- 373 Invalid comparison operations requested.
- 374 Search index handle was requested for an empty index.
- 375 The combination of operator and requested operator mode is not supported.

- 380 Search result is incomplete, search was discontinued due to threshold.
- 381 Index lookup revealed that query contained stopwords.
- 401 The operation cannot be performed because the Text Search Engine detected a corruption in the index files used for field/attribute name mapping.
- 402 The operation cannot be performed because the Text Search Engine detected an invalid field or attribute name.
- 403 The operation cannot be performed because the given field or attribute name is unknown.
- 404 The limit of different attributes for one of the attribute types or of different fields has been exceeded for this index.
- 500 The document/data contains an invalid character sequence (in a UTF8, UTF16, or DBCS source).
- 501 The code page converter was in error.
- 502 The document/data contains an incomplete character sequence (in a UTF8, UTF16, or DBCS source).
- 503 The code page converter has an invalid descriptor.
- 600 The XML document contains an asynchronous entity. For example, an unquoted XML attribute value.
- 602 Invalid character reference, (for example, or).
- 603 Invalid binary entity reference.
- 604 XML Parser Expat could not be created.
- 605 An attribute name in tag must be unique.
- 607 XML Parser found an invalid external entity reference.
- 608 Documents includes an incorrect token, like missing a < or >.
- 609 XML documents must have an enclosing tag, and after this enclosing end tag no text is allowed.
- 610 A processing instruction is not allowed at its position. For example, the first processing instruction is not the prolog <?xml .. ?>.
- 611 An element is a sequence of start tag, content, and end tag. This error occurred, for example, from the sequence "<s> text /s>", because the end tag is not correct.
- 612 Memory allocation failed in XML parser.
- 614 Invalid parameter entity reference.
- 615 A non-complete character, maybe only the first byte of a 2-byte UTF8 character.
- 616 Recursive entity reference.
- 617 XML Syntax error; for example text outside the enclosing start and end tag.
- 618 Every start tag needs a matching end tag.
- 619 Unclosed cdata section.
- 620 Unclosed token; for example text after the last token in a document.
- 621 There is an entity in the document that could not be resolved.

- 622 Unexpected error.
- 631 Could not parse field or attribute information in meta-tag. Tag must have the format `<meta name="abc" content="xyz">`; maybe attributes name or content of the meta-tag not correct.
- 632 The entity could not be transformed to a character.
- 650 Different field definitions begin with the same start tag.
- 651 One start tag contains another, so the tags are ambiguous.
- 652 If a field and an attribute use the same start tag, then they must use the same end tag or both no end tag.
- 653 A field not yet closed if the document ends.
- 654 No document model is specified for the structured format. The document will be parsed as plain text document without field or attribute informations.
- 670 The operation could not be performed, because it requires the "Outside In" (TM) libraries, which could not be found.
- 671 The operation could not be performed, because a required procedure from the "Outside In" (TM) libraries could not be loaded. Probably the libraries are outdated or corrupted.
- 672 An error occurred while the document was processed with "Outside In".

---

## Part 10. Troubleshooting



---

## Chapter 58. Tracing faults

If you need to report an error to an IBM<sup>®</sup> representative, you may be asked to switch on tracing so that information can be written to a file that can be used for locating the error.

As system performance is affected when tracing is switched on, only use the trace facility when directed by an IBM Support Center representative, or by your technical support representative.

To turn tracing on, use the DB2 facility:

```
db2trc on
```

See the DB2 documentation for further information.

To receive information specific to Net Search Extender, a mask with component in 96 can be used:

```
db2trc on -m *.*.96.*.*
```

In the case of severe errors, it may also help to look in db2diag log file.





---

## Chapter 59. Dropping DB2 objects without using the correct Net Search Extender commands

---

### Dropping a table

Before you drop a table with one or more text indexes, you must issue the following command for each text index:

```
db2text drop index <index_name> for text
```

If you accidentally drop a table before you drop the indexes, parts of the indexes will still exist, for example, the administration tables and text index files.

To remove these files, drop the indexes using the `db2text drop index` command, even though the table no longer exists.

---

### Dropping a database

Before you drop a database with one or more text indexes, issue the following command for each text index:

```
db2text drop index <index_name> for text
```

If you do not use this command, you must delete all of the index files manually in `index_directory` and `index_work_directory`.

Note that if the indexes belonging to the dropped database were created during an automatic update, you will need to edit the scheduler file `ctedem.dat`.

To do this, enter the following:

For UNIX:

```
db2text stop force  
cd ~/sql11ib/db2ext
```

or for Windows:

```
db2text stop force  
cd <db2_install_path>\sql11ib\<db2_instance_name>\db2ext
```

Open the file `ctedem.dat` in the directory and remove all entries referring to the dropped database.



---

## Chapter 60. Installation return codes on Windows

### **setup.exe return codes on Windows**

The setup.exe return codes in the setup.log are the following:

- 0 Success
- -1 General error
- -2 Invalid mode
- -3 Required data not found in the setup.iss file
- -4 Not enough memory available
- -5 File does not exist
- -6 cannot write to the response file
- -7 Unable to write to the log file
- -8 Invalid path to the install shield silent response (.iss) file
- -9 Not a valid list type
- -10 Data type invalid
- -11 Unknown error during setup
- -12 Dialog boxes are out of order
- -51 Cannot create the specified folder
- -52 Cannot access the specified file or folder
- -53 Invalid option selected



---

## Chapter 61. Hints and tips

### Authorization

When you issue the DB2TEXT START command on Windows, ensure that you are a member of the Administrators group. Otherwise, the DB2TEXT START command fails and returns the following message: CTE0218 Function "OpenSCManager()" failed with error code "5".

### Authorization

On Windows, ensure that the Net Search Extender instance service DB2EXT-<DB2\_instance\_name> runs under a user account rather than the system account. If you run under the system account, you can not enable your database.

### Language

As Net Search Extender event log messages are always displayed in the DB2 server language, event log messages for commands issued from the DB2 Control Center might be displayed in a different language to language set for the DB2 Control Center.

### Modifying the db2cli.ini file

If you modified the db2cli.ini file and have problems using Net Search Extender, restore the original version of db2cli.ini.

### Client server interoperability

You must install the same Net Search Extender fix pack level on both the server and client.

Interoperability between client to server, and vice versa, is only possible for the supported platforms.

### Log size

If the indexing process is not completed because the error and warning messages require more space than the available DB2 log size, DB2 will rollback the whole transaction and not commit the log table entries. This means that you will not be able to see the entries.

Check the DB2 documentation for information on how to increase your transaction log size to avoid this situation.

### DBCS object names

If you use DBCS object names in db2text administration commands, it is necessary to enclose those names in double quotes to avoid uppercase transformation.

### Incremental index update on nicknames

If initial updates on two or more indexes are started at the same time, the update command might return a SQL0803N error. In this case, try the update command again.

### Single masking and character normalization

Words like 'über' are normalized and are stored in the index in their normalized form ('ueber'). Therefore, if you issue a query containing single character masking, for example '\_ber', you will not find 'über'.

### Using duplicate cache column names

If you use duplicate cache column names, you will not get an error during

text index creation or index update, but you will not be able to search. When you try to search, you will get an SQL error message stating that duplicate columns were used.

#### **Incorrect shared memory size**

If the specified maximum cache size is too small in a db2text activate cache command, the required cache size that is displayed in the resulting error message is incorrect.

Check for the correct cache size by using the DB2EXT.MAXIMUM\_CACHE\_SIZE and DB2EXT.PCTFREE functions. Correct the maximum cache size by using the db2text alter index command and activate the cache again.

#### **Unicode tables in a non Unicode database**

You are not allowed to create a text index on a Unicode table if your database does not support Unicode.

#### **Incorrect codepage for LANG variable in a query on Linux**

If you use a 7-bit ASCII codepage for the LANG variable setting in a query, the following error is displayed: SearchString parse: check LANG & locale charmap values..

To avoid this happening, change your LANG variable to a 8-bit LANG value, restart DB2, and try searching again.

#### **File access problems**

If you cannot access files, for example, a document model, a text index, or a thesaurus, ensure that you used the correct password and have the correct authorization for running the Net Search Extender instance services. This applies especially to shared resources on mapped network drives.

#### **Cache cannot be used**

During search or activate cache, the following error message might be displayed: CTE0271 Cache not usable, DEACTIVATE and ACTIVATE RECREATE required. To solve this problem, check your system settings and try to increase the amount of paging space and free memory.

#### **Instance services not dropped after an uninstall**

If the instance services are not dropped after you uninstall Net Search Extender, use the following tool to drop the services manually: ctereg <instancename> unregister. For example, ctereg db2-0 unregister.

#### **ctepcrx terminates abnormally on UNIX**

Check that the used instance owner does not have a separate fenced user ID. To check this, open the file .fenced in <instance\_home\_dir>/sqllib/adm and check if the instance owner is also the fenced user.

#### **Index update fails with SQL0668N**

If the index update fails with SQL0668N (rc=1), the access to the text-maintained staging table was blocked due to a database operation on the base table that requires integrity processing for dependent tables.

Use the RESET PENDING command to unblock the staging table, and re-run the index update command

#### **When SQL scalar search functions cannot be used on tables spanning multiple partitions**

In situations where the text search function does not directly reference the partitioned table or if it is a member of a subselect that contains an OUTER JOIN clause, the search will return an error (SQL0270N rc=109). If this occurs, attempt to manually rewrite the query to a different form, and reissue it.

**CTE0249 error when verifying partitioned database environment setup**

On Windows, if you verify the partitioned database environment setup and encounter the error *CTE0249 Executable program "cteprisc" terminated abnormally*, perform the following actions:

1. For each of the computers in the partitioned database environment configuration, check that the environment has been prepared for a partitioned DB2 server, as instructed in . Specifically, ensure that the "Trust computer for delegation" checkbox on the General tab of each computer's account Properties dialog box in the Active Directory Users and Computers console is checked.
2. Ensure that the Windows service "DB2 Remote Command Server" has been started on all the participating computers.

**CTE0150E error when creating a text index on a range partitioned table without the ADMINISTRATION TABLES IN clause**

If you try to create a text index on a range partitioned table please make sure to specify the ADMINISTRATION TABLES IN clause without fail. See CTE0150E for more information.





---

## Part 11. Appendixes



---

## Appendix A. Overview of the DB2 technical information

DB2 technical information is available through the following tools and methods:

- DB2 Information Center
  - Topics (Task, concept and reference topics)
  - Help for DB2 tools
  - Sample programs
  - Tutorials
- DB2 books
  - PDF files (downloadable)
  - PDF files (from the DB2 PDF DVD)
  - printed books
- Command line help
  - Command help
  - Message help

**Note:** The DB2 Information Center topics are updated more frequently than either the PDF or the hardcopy books. To get the most current information, install the documentation updates as they become available, or refer to the DB2 Information Center at [ibm.com](http://ibm.com).

You can access additional DB2 technical information such as technotes, white papers, and IBM Redbooks® publications online at [ibm.com](http://ibm.com). Access the DB2 Information Management software library site at <http://www.ibm.com/software/data/sw-library/>.

### Documentation feedback

We value your feedback on the DB2 documentation. If you have suggestions for how to improve the DB2 documentation, send an e-mail to [db2docs@ca.ibm.com](mailto:db2docs@ca.ibm.com). The DB2 documentation team reads all of your feedback, but cannot respond to you directly. Provide specific examples wherever possible so that we can better understand your concerns. If you are providing feedback on a specific topic or help file, include the topic title and URL.

Do not use this e-mail address to contact DB2 Customer Support. If you have a DB2 technical issue that the documentation does not resolve, contact your local IBM service center for assistance.

---

## DB2 technical library in hardcopy or PDF format

The following tables describe the DB2 library available from the IBM Publications Center at [www.ibm.com/shop/publications/order](http://www.ibm.com/shop/publications/order). English and translated DB2 Version 9.7 manuals in PDF format can be downloaded from [www.ibm.com/support/docview.wss?rs=71&uid=swg2700947](http://www.ibm.com/support/docview.wss?rs=71&uid=swg2700947).

Although the tables identify books available in print, the books might not be available in your country or region.

The form number increases each time a manual is updated. Ensure that you are reading the most recent version of the manuals, as listed below.

**Note:** The *DB2 Information Center* is updated more frequently than either the PDF or the hard-copy books.

*Table 17. DB2 technical information*

<b>Name</b>	<b>Form Number</b>	<b>Available in print</b>	<b>Last updated</b>
<i>Administrative API Reference</i>	SC27-2435-01	Yes	November, 2009
<i>Administrative Routines and Views</i>	SC27-2436-01	No	November, 2009
<i>Call Level Interface Guide and Reference, Volume 1</i>	SC27-2437-01	Yes	November, 2009
<i>Call Level Interface Guide and Reference, Volume 2</i>	SC27-2438-01	Yes	November, 2009
<i>Command Reference</i>	SC27-2439-01	Yes	November, 2009
<i>Data Movement Utilities Guide and Reference</i>	SC27-2440-00	Yes	August, 2009
<i>Data Recovery and High Availability Guide and Reference</i>	SC27-2441-01	Yes	November, 2009
<i>Database Administration Concepts and Configuration Reference</i>	SC27-2442-01	Yes	November, 2009
<i>Database Monitoring Guide and Reference</i>	SC27-2458-00	Yes	August, 2009
<i>Database Security Guide</i>	SC27-2443-01	Yes	November, 2009
<i>DB2 Text Search Guide</i>	SC27-2459-01	Yes	November, 2009
<i>Developing ADO.NET and OLE DB Applications</i>	SC27-2444-00	Yes	August, 2009
<i>Developing Embedded SQL Applications</i>	SC27-2445-01	Yes	November, 2009
<i>Developing Java Applications</i>	SC27-2446-01	Yes	November, 2009
<i>Developing Perl, PHP, Python, and Ruby on Rails Applications</i>	SC27-2447-00	No	August, 2009
<i>Developing User-defined Routines (SQL and External)</i>	SC27-2448-01	Yes	November, 2009
<i>Getting Started with Database Application Development</i>	GI11-9410-01	Yes	November, 2009
<i>Getting Started with DB2 Installation and Administration on Linux and Windows</i>	GI11-9411-00	Yes	August, 2009

*Table 17. DB2 technical information (continued)*

<b>Name</b>	<b>Form Number</b>	<b>Available in print</b>	<b>Last updated</b>
<i>Globalization Guide</i>	SC27-2449-00	Yes	August, 2009
<i>Installing DB2 Servers</i>	GC27-2455-01	Yes	November, 2009
<i>Installing IBM Data Server Clients</i>	GC27-2454-00	No	August, 2009
<i>Message Reference Volume 1</i>	SC27-2450-01	No	November, 2009
<i>Message Reference Volume 2</i>	SC27-2451-01	No	November, 2009
<i>Net Search Extender Administration and User's Guide</i>	SC27-2469-01	No	November, 2009
<i>Partitioning and Clustering Guide</i>	SC27-2453-01	Yes	November, 2009
<i>pureXML Guide</i>	SC27-2465-01	Yes	November, 2009
<i>Query Patroller Administration and User's Guide</i>	SC27-2467-00	No	August, 2009
<i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i>	SC27-2468-00	No	August, 2009
<i>SQL Procedural Languages: Application Enablement and Support</i>	SC27-2470-00	Yes	August, 2009
<i>SQL Reference, Volume 1</i>	SC27-2456-01	Yes	November, 2009
<i>SQL Reference, Volume 2</i>	SC27-2457-01	Yes	November, 2009
<i>Troubleshooting and Tuning Database Performance</i>	SC27-2461-01	Yes	November, 2009
<i>Upgrading to DB2 Version 9.7</i>	SC27-2452-01	Yes	November, 2009
<i>Visual Explain Tutorial</i>	SC27-2462-00	No	August, 2009
<i>What's New for DB2 Version 9.7</i>	SC27-2463-01	Yes	November, 2009
<i>Workload Manager Guide and Reference</i>	SC27-2464-00	Yes	August, 2009
<i>XQuery Reference</i>	SC27-2466-01	No	November, 2009

*Table 18. DB2 Connect-specific technical information*

<b>Name</b>	<b>Form Number</b>	<b>Available in print</b>	<b>Last updated</b>
<i>Installing and Configuring DB2 Connect Personal Edition</i>	SC27-2432-01	Yes	November, 2009
<i>Installing and Configuring DB2 Connect Servers</i>	SC27-2433-01	Yes	November, 2009

Table 18. DB2 Connect-specific technical information (continued)

Name	Form Number	Available in print	Last updated
DB2 Connect User's Guide	SC27-2434-01	Yes	November, 2009

Table 19. Information Integration technical information

Name	Form Number	Available in print	Last updated
Information Integration: Administration Guide for Federated Systems	SC19-1020-02	Yes	August, 2009
Information Integration: ASNCLP Program Reference for Replication and Event Publishing	SC19-1018-04	Yes	August, 2009
Information Integration: Configuration Guide for Federated Data Sources	SC19-1034-02	No	August, 2009
Information Integration: SQL Replication Guide and Reference	SC19-1030-02	Yes	August, 2009
Information Integration: Introduction to Replication and Event Publishing	GC19-1028-02	Yes	August, 2009

## Ordering printed DB2 books

If you require printed DB2 books, you can buy them online in many but not all countries or regions. You can always order printed DB2 books from your local IBM representative. Keep in mind that some softcopy books on the *DB2 PDF Documentation DVD* are unavailable in print. For example, neither volume of the *DB2 Message Reference* is available as a printed book.

Printed versions of many of the DB2 books available on the *DB2 PDF Documentation DVD* can be ordered for a fee from IBM. Depending on where you are placing your order from, you may be able to order books online, from the IBM Publications Center. If online ordering is not available in your country or region, you can always order printed DB2 books from your local IBM representative. Note that not all books on the *DB2 PDF Documentation DVD* are available in print.

**Note:** The most up-to-date and complete DB2 documentation is maintained in the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7>.

To order printed DB2 books:

- To find out whether you can order printed DB2 books online in your country or region, check the IBM Publications Center at <http://www.ibm.com/shop/publications/order>. You must select a country, region, or language to access publication ordering information and then follow the ordering instructions for your location.
- To order printed DB2 books from your local IBM representative:

1. Locate the contact information for your local representative from one of the following Web sites:
  - The IBM directory of world wide contacts at [www.ibm.com/planetwide](http://www.ibm.com/planetwide)
  - The IBM Publications Web site at <http://www.ibm.com/shop/publications/order>. You will need to select your country, region, or language to the access appropriate publications home page for your location. From this page, follow the "About this site" link.
2. When you call, specify that you want to order a DB2 publication.
3. Provide your representative with the titles and form numbers of the books that you want to order. For titles and form numbers, see "DB2 technical library in hardcopy or PDF format" on page 307.

---

## Displaying SQL state help from the command line processor

DB2 products return an SQLSTATE value for conditions that can be the result of an SQL statement. SQLSTATE help explains the meanings of SQL states and SQL state class codes.

To start SQL state help, open the command line processor and enter:

```
? sqlstate or ? class code
```

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, ? 08003 displays help for the 08003 SQL state, and ? 08 displays help for the 08 class code.

---

## Accessing different versions of the DB2 Information Center

For DB2 Version 9.7 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/>.

For DB2 Version 9.5 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/>.

For DB2 Version 9.1 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>.

For DB2 Version 8 topics, go to the *DB2 Information Center* URL at: <http://publib.boulder.ibm.com/infocenter/db2luw/v8/>.

---

## Displaying topics in your preferred language in the DB2 Information Center

The DB2 Information Center attempts to display topics in the language specified in your browser preferences. If a topic has not been translated into your preferred language, the DB2 Information Center displays the topic in English.

- To display topics in your preferred language in the Internet Explorer browser:
  1. In Internet Explorer, click the **Tools** —> **Internet Options** —> **Languages...** button. The Language Preferences window opens.
  2. Ensure your preferred language is specified as the first entry in the list of languages.
    - To add a new language to the list, click the **Add...** button.

**Note:** Adding a language does not guarantee that the computer has the fonts required to display the topics in the preferred language.

- To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
- 3. Refresh the page to display the DB2 Information Center in your preferred language.
- To display topics in your preferred language in a Firefox or Mozilla browser:
  1. Select the button in the **Languages** section of the **Tools** —> **Options** —> **Advanced** dialog. The Languages panel is displayed in the Preferences window.
  2. Ensure your preferred language is specified as the first entry in the list of languages.
    - To add a new language to the list, click the **Add...** button to select a language from the Add Languages window.
    - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
  3. Refresh the page to display the DB2 Information Center in your preferred language.

On some browser and operating system combinations, you must also change the regional settings of your operating system to the locale and language of your choice.

---

## Updating the DB2 Information Center installed on your computer or intranet server

A locally installed DB2 Information Center must be updated periodically.

### Before you begin

A DB2 Version 9.7 Information Center must already be installed. For details, see the “Installing the DB2 Information Center using the DB2 Setup wizard” topic in *Installing DB2 Servers*. All prerequisites and restrictions that applied to installing the Information Center also apply to updating the Information Center.

### About this task

An existing DB2 Information Center can be updated automatically or manually:

- Automatic updates - updates existing Information Center features and languages. An additional benefit of automatic updates is that the Information Center is unavailable for a minimal period of time during the update. In addition, automatic updates can be set to run as part of other batch jobs that run periodically.
- Manual updates - should be used when you want to add features or languages during the update process. For example, a local Information Center was originally installed with both English and French languages, and now you want to also install the German language; a manual update will install German, as well as, update the existing Information Center features and languages. However, a manual update requires you to manually stop, update, and restart the Information Center. The Information Center is unavailable during the entire update process.

### Procedure



This topic details the process for automatic updates. For manual update instructions, see the “Manually updating the DB2 Information Center installed on your computer or intranet server” topic.

To automatically update the DB2 Information Center installed on your computer or intranet server:

1. On Linux operating systems,
  - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `/opt/ibm/db2ic/V9.7` directory.
  - b. Navigate from the installation directory to the `doc/bin` directory.
  - c. Run the `ic-update` script:  
`ic-update`
2. On Windows operating systems,
  - a. Open a command window.
  - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `<Program Files>\IBM\DB2 Information Center\Version 9.7` directory, where `<Program Files>` represents the location of the Program Files directory.
  - c. Navigate from the installation directory to the `doc\bin` directory.
  - d. Run the `ic-update.bat` file:  
`ic-update.bat`

### Results

The DB2 Information Center restarts automatically. If updates were available, the Information Center displays the new and updated topics. If Information Center updates were not available, a message is added to the log. The log file is located in `doc\eclipse\configuration` directory. The log file name is a randomly generated number. For example, `1239053440785.log`.

---

## Manually updating the DB2 Information Center installed on your computer or intranet server

If you have installed the DB2 Information Center locally, you can obtain and install documentation updates from IBM.

### About this task

Updating your locally-installed *DB2 Information Center* manually requires that you:

1. Stop the *DB2 Information Center* on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to apply updates. The Workstation version of the DB2 Information Center always runs in stand-alone mode. .
2. Use the Update feature to see what updates are available. If there are updates that you must install, you can use the Update feature to obtain and install them

**Note:** If your environment requires installing the *DB2 Information Center* updates on a machine that is not connected to the internet, mirror the update site to a local file system using a machine that is connected to the internet and has the *DB2 Information Center* installed. If many users on your network will be

installing the documentation updates, you can reduce the time required for individuals to perform the updates by also mirroring the update site locally and creating a proxy for the update site.

If update packages are available, use the Update feature to get the packages. However, the Update feature is only available in stand-alone mode.

3. Stop the stand-alone Information Center, and restart the *DB2 Information Center* on your computer.

**Note:** On Windows 2008, Windows Vista (and higher), the commands listed later in this section must be run as an administrator. To open a command prompt or graphical tool with full administrator privileges, right-click the shortcut and then select **Run as administrator**.

### Procedure

To update the *DB2 Information Center* installed on your computer or intranet server:


1. Stop the *DB2 Information Center*.
  - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click **DB2 Information Center** service and select **Stop**.
  - On Linux, enter the following command:  

```
/etc/init.d/db2icdv97 stop
```
2. Start the Information Center in stand-alone mode.
  - On Windows:
    - a. Open a command window.
    - b. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the *Program\_Files\IBM\DB2 Information Center\Version 9.7* directory, where *Program\_Files* represents the location of the Program Files directory.
    - c. Navigate from the installation directory to the *doc\bin* directory.
    - d. Run the *help\_start.bat* file:  

```
help_start.bat
```
  - On Linux:
    - a. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the */opt/ibm/db2ic/V9.7* directory.
    - b. Navigate from the installation directory to the *doc/bin* directory.
    - c. Run the *help\_start* script:  

```
help_start
```

The systems default Web browser opens to display the stand-alone Information Center.

3. Click the **Update** button (). (JavaScript™ must be enabled in your browser.) On the right panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check the selections you want to install, then click **Install Updates**.
5. After the installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center:
  - On Windows, navigate to the installation directory's *doc\bin* directory, and run the *help\_end.bat* file:

help\_end.bat

**Note:** The help\_end batch file contains the commands required to safely stop the processes that were started with the help\_start batch file. Do not use Ctrl-C or any other method to stop help\_start.bat.

- On Linux, navigate to the installation directory's doc/bin directory, and run the help\_end script:

```
help_end
```

**Note:** The help\_end script contains the commands required to safely stop the processes that were started with the help\_start script. Do not use any other method to stop the help\_start script.

7. Restart the *DB2 Information Center*.

- On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click **DB2 Information Center** service and select **Start**.
- On Linux, enter the following command:  

```
/etc/init.d/db2icdv97 start
```

## Results

The updated *DB2 Information Center* displays the new and updated topics.

---

## DB2 tutorials

The DB2 tutorials help you learn about various aspects of DB2 products. Lessons provide step-by-step instructions.

### Before you begin

You can view the XHTML version of the tutorial from the Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

Some lessons use sample data or code. See the tutorial for a description of any prerequisites for its specific tasks.

### DB2 tutorials

To view the tutorial, click the title.

#### **“pureXML®” in *pureXML Guide***

Set up a DB2 database to store XML data and to perform basic operations with the native XML data store.

#### **“Visual Explain” in *Visual Explain Tutorial***

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

---

## DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2 database products.

### DB2 documentation

Troubleshooting information can be found in the *DB2 Troubleshooting Guide* or the Database fundamentals section of the *DB2 Information Center*. There you will find information about how to isolate and identify problems using

DB2 diagnostic tools and utilities, solutions to some of the most common problems, and other advice on how to solve problems you might encounter with your DB2 database products.

#### **DB2 Technical Support Web site**

Refer to the DB2 Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes, Authorized Program Analysis Reports (APARs or bug fixes), fix packs, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the DB2 Technical Support Web site at [http://www.ibm.com/software/data/db2/support/db2\\_9/](http://www.ibm.com/software/data/db2/support/db2_9/)

---

## **Terms and Conditions**

Permissions for the use of these publications is granted subject to the following terms and conditions.

**Personal use:** You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

**Commercial use:** You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

---

## Appendix B. Notices

This information was developed for products and services offered in the U.S.A. Information about non-IBM products is based on information available at the time of first publication of this document and is subject to change.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan, Ltd.  
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

**The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web

sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information that has been exchanged, should contact:

IBM Canada Limited  
Office of the Lab Director  
8200 Warden Avenue  
Markham, Ontario  
L6G 1C7  
CANADA

Such information may be available, subject to appropriate terms and conditions, including, in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (*your company name*) (*year*). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. *\_enter the year or years\_*. All rights reserved.

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)<sup>®</sup> are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

The following terms are trademarks or registered trademarks of other companies

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java<sup>™</sup> and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Intel<sup>®</sup>, Intel logo, Intel Inside<sup>®</sup>, Intel Inside logo, Intel<sup>®</sup> Centrino<sup>®</sup>, Intel Centrino logo, Celeron<sup>®</sup>, Intel<sup>®</sup> Xeon<sup>®</sup>, Intel SpeedStep<sup>®</sup>, Itanium<sup>®</sup>, and Pentium<sup>®</sup> are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- Microsoft, Windows, Windows NT<sup>®</sup>, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.





---

# Index

## A

- ACTIVATE CACHE command 197
- administration commands
  - database administrators
    - DB2EXTHL 194
    - DISABLE DATABASE 193
    - ENABLE DATABASE 191
    - summary 191
  - instance owners
    - CONTROL 185
    - START 187
    - STOP 188
    - summary 185
  - text table owners
    - ACTIVATE CACHE 197
    - ALTER INDEX 199
    - CLEAR EVENTS 203
    - COPYRIGHT 226
    - CREATE INDEX 204
    - DB2EXTTH 220
    - DEACTIVATE CACHE 217
    - DROP INDEX 219
    - HELP 225
    - RESET PENDING 221
    - summary 197
    - UPDATE INDEX 222
- ALTER INDEX command 199
- attribute values
  - number attributes 165

## B

- books
  - ordering 310

## C

- character normalization 155
- CLEAR EVENTS command 203
- client/server environment
  - overview 19
- column transformation function 8
- commands
  - CONTROL 185
  - START 187
  - STOP 188
- COMMITCOUNT
  - keyword 212
- compiling
  - thesaurus 145
- CONTAINS function
  - Net Search Extender 239
  - parameters 239
- CONTROL command 185
- COPYRIGHT command
  - details 226
  - parameters 226
- CREATE INDEX command 204

## D

- database partitions
  - Net Search Extender 45
- DB2 Information Center
  - languages 311
  - updating 312, 313
  - versions 311
- DB2EXT.HIGHLIGHT function 243
  - input parameters 244
  - returned parameters 245
- DB2EXT.TEXTSEARCH command
  - details 247
  - input parameters 247
  - returned parameters 248
- DB2EXT.TEXTSEARCH function
  - details 240
  - input parameters 242
  - returned parameters 242
- DB2EXTHL command 194
- DB2EXTTH command 220
  - parameters 221
- db2nse\_deinstall command 228
- db2nse\_install command 227
- db2nsels command 228
- db2text commands 15
- DEACTIVATE CACHE command 217
  - parameters 218
- DISABLE DATABASE command 193
- document formats
  - Outside In filtering software 53
- document models
  - default 166
  - document attributes
    - limitations 180
    - overview 165
  - document type definition 177
  - Outside In filter technology 174
  - overview 164
  - plain-text
    - defining 167
    - example 163
  - reference information 177
  - text fields
    - limitations 180
    - overview 165
  - XML 170
- documentation
  - overview 307
  - PDF files 307
  - printed 307
  - terms and conditions of use 316
- DROP INDEX command
  - parameters 219
- DROP INDEX Net Search Extender command 219
- dropping
  - text indexes 112

## E

- ENABLE DATABASE command 191

externally stored data  
overview 9

## F

fault finding 295  
functions  
Net Search Extender search 239

## G

GPP document model  
details 167  
document type definition 177  
element parameters 167  
indexing GPP documents 167

## H

help  
configuring language 311  
SQL statements 311  
HELP command  
DB2 Net Search Extender 225  
HTML document model DTD 177

## I

incremental updates based on integrity processing 11  
indexes  
clearing events 111  
nicknames 10  
information catalogs  
overview 277  
views  
database-level information 277  
index-level information 280  
text index tables 283  
installation  
Net Search Extender  
db2nse\_install command 227  
db2nsels command 228  
directory names 29  
file names 29  
Outside In filtering software 31  
partitioned DB2 servers 23  
system requirements 21  
UNIX 25  
verification (UNIX) 33  
verification (Windows) 33  
Windows 27  
instance services  
locking services 65  
overview 8  
starting 63  
stopping 63  
update services 69  
L  
locking services  
details 65  
viewing snapshot 67

## M

memory  
requirements  
AIX 47  
Linux 48  
Solaris 47  
Windows 47  
stored procedures 47  
migration  
Net Search Extender 37  
text indexes 37

## N

Net Search Extender (NSE)  
administration tables  
overview 9  
views 9  
character normalization 155  
code pages 51  
column names 49  
Control Center 71  
creating text indexes  
cache table panel 80, 83  
name panel 74  
summary panel 83  
target panel 75  
text properties panel 77  
update characteristics panel 79  
databases  
disabling 93  
enabling 72, 91  
directory location 43  
document formats 51  
extended text-maintained staging infrastructure 57  
incremental updates 57  
index events  
clearing 87  
displaying 87  
index names 49  
index status 88  
index storage 43  
instance services 8, 61  
migration 37  
outer joins 135  
overview 3, 13  
planning 41  
reason codes 287  
resources 45  
specific characters as parts of words 155  
starting 72  
stop words 154  
stopping 72  
table names 49  
text index caches  
activating 87  
deactivating 88  
text indexes  
administering 73  
altering 84  
creating 74  
dropping 86  
maintaining 84  
updating 87  
Text Search Engine 153  
tips 301

- Net Search Extender (NSE) *(continued)*
  - tokenization 153
  - XML support 10
- notices 317
- NUMBEROFMATCHES function 239
  - parameters 240

## O

- ordering DB2 books 310
- outer joins
  - text search 135
- Outside In filters
  - attribute values 180
  - details 174
  - indexing 175

## P

- partitioned database environments
  - Net Search Extender 9
- partitioned tables
  - overview 11
- performance
  - indexes
    - degradation causes 107
    - recommendations 107
- problem determination
  - information available 315
  - tutorials 315

## R

- reorganization
  - text indexes 109
- replication
  - capture tables 204
- replication capture table 211
- RESET PENDING command 221
- return codes
  - installation 299

## S

- scenarios
  - Net Search Extender
    - overview 139
    - SQL scalar search 139
    - SQL table-valued function 141
    - stored procedure search 140
- SCORE function
  - details 240
  - parameters 240
- search argument keywords
  - THESAURUS 237
- search arguments
  - overview 231
  - syntax 231
- search functions
  - CONTAINS 239
  - DB2EXT.HIGHLIGHT 243
  - DB2EXT.TEXTSEARCH 240
  - details 115
  - NUMBEROFMATCHES 239
  - overview 239

- search functions *(continued)*
  - performance considerations 137
  - samples 125
  - SCORE 240
  - SQL scalar 117
  - SQL table-valued 129
    - stored procedures 127, 247
- search parameters 234
- servers
  - tracing faults 295
- SQL
  - table-valued function
    - highlighting 129
- SQL scalar search functions
  - CONTAINS 117
  - NUMBEROFMATCHES 117
  - overview 5
  - SCORE 118
  - searching on more than one column 133
  - SQL search arguments 119
- SQL search arguments
  - Boolean AND operator 119
  - Boolean NOT operator 120
  - Boolean OR operator 119
  - escape character in masking 121
  - free-text search 123
  - fuzzy search 120
  - numeric attribute search 122
  - terms in any sequence 119
  - terms in fixed sequence 121
  - terms in same paragraph 121
  - terms in same sentence 121
  - terms in sections 122
  - thesaurus search 122
  - wildcard character masking 120
- SQL statements
  - help
    - displaying 311
- START Net Search Extender command 187
- STOP command 188
- stop words 154
- stored procedure search
  - DB2EXT.TEXTSEARCH command 247
  - function 247
  - overview 6
- structured documents
  - details 163
  - document models 163
  - Net Search Extender 157
  - XML document processing 159

## T

- table-valued function
  - overview 7
- terms and conditions
  - publications 316
- text indexes
  - altering 111
  - backing up 113
  - clearing index events 111
  - creating
    - binary data types 97
    - disabling database 93
    - enabling database 91
    - nicknames 98
    - overview 89, 95

- text indexes (*continued*)
  - creating (*continued*)
    - range-partitioned tables 105
    - stored procedure searches 99
    - unsupported data types 97
  - details 109
  - dropping 112
  - dropping database 297
  - dropping table 297
  - maintaining 89
  - migrating 37
  - removing files from /tmp directory 113
  - reorganizing 109
  - restoring 113
  - updating 109
  - viewing status 112
  - views 103
- Text Search Engine
  - character normalization 155
  - using specific characters as parts of words 155
- thesaurus search
  - syntax 237
  - THESAURUS keyword 237
- thesauruses
  - compiling 145
  - compiling definition file 146
  - creating 145
  - creating definition file 145
  - defining relations 145
  - error messages 149
  - overview 143
  - predefined relations 144
  - search syntax 234
  - structure 143
  - support 146
  - supported coded character set identifiers (CCSIDs) 148
  - THESAURUS keyword 234
- tokenization 153
- tracing faults 295
- troubleshooting
  - online information 315
  - tutorials 315
- tutorials
  - list 315
  - problem determination 315
  - troubleshooting 315
  - Visual Explain 315

## U

- uninstalling
  - Net Search Extender
    - db2nse\_deinstall command 228
    - overview 35
    - UNIX 35
    - Windows 35
- UPDATE INDEX Net Search Extender command 222
- updates
  - DB2 Information Center 312, 313
- user roles
  - database administrators 55
  - DB2 instance owner 55
  - Net Search Extender 55
  - text table owners 55
- using specific characters as parts of words 155

## V

- views
  - db2ext.textindexes 112
- views for Net Search Extender
  - database-level information
    - db2ext.dbdefaults 277
  - event information 283
  - index-level information
    - db2ext.indexconfiguration 282
    - db2ext.textconfiguration 280
    - db2ext.textindexes 280
    - db2ext.textindexformats 282
    - db2ext.textinformatms 280
  - log table 284

## W

- Windows
  - system errors 273

## X

- XML decomposition
  - Net Search Extender 166
- XML document model
  - details 170
  - document type definition 177
  - XPath expression semantics 178
- XML document processing
  - customized XML document model 160
  - default XML document model 159
  - XQuery support 161





Printed in USA

SC27-2469-01



Spine information:

IBM DB2 9.7 for Linux, UNIX, and Windows **Version 9 Release 7**

**Net Search Extender Administration and User's Guide**

