

# Content Server DQL Reference Manual



Version 5.3 SP1  
September 2005

Copyright © 1994-2005 EMC Corporation

# Table of Contents

---

<b>Preface</b> .....	11
<b>Chapter 1 DQL Language Elements</b> .....	13
Literals .....	13
Integer literals .....	13
Floating point literals .....	14
Character string literals .....	15
ID literals .....	15
Date literals .....	15
Default formats .....	16
Short date formats .....	16
ANSI format .....	17
Other character string formats .....	17
Date literal keywords .....	18
Date output formats .....	18
Special keywords .....	19
Functions .....	20
Scalar functions .....	20
UPPER .....	21
LOWER .....	21
SUBSTR .....	21
Aggregate functions .....	22
COUNT .....	22
MIN .....	23
MAX .....	23
AVG .....	24
SUM .....	24
Date functions .....	24
DATEDIFF .....	25
DATEADD .....	26
DATEFLOOR .....	26
DATETOSTRING .....	27
The ID function .....	27
The MFILE_URL function .....	28
Examples .....	29
Predicates .....	29
Arithmetic operators .....	29
Comparison operators .....	30
Column and attribute predicates .....	30
Predicates for columns and single-valued attributes .....	31
Predicates for repeating attributes .....	32
Pattern matching with LIKE .....	34
The percent sign .....	34
The underbar .....	34
Matching cases .....	35
The ESCAPE character .....	35
SysObject predicates .....	35

	The TYPE predicate .....	36
	The FOLDER predicate .....	36
	The CABINET predicate.....	37
	Logical operators .....	38
	AND and OR .....	38
	NOT .....	38
	Order of precedence .....	39
	DQL reserved words.....	40
<b>Chapter 2</b>	<b>DQL Statements</b> .....	41
<b>Chapter 3</b>	<b>Administration Methods</b> .....	167
	Invoking administration methods .....	167
	Using Documentum Administrator .....	168
	Using the EXECUTE statement .....	168
	Using the Apply method.....	168
	Scope of the administration methods .....	168
	Administration method operations .....	169
<b>Chapter 4</b>	<b>Using DQL</b> .....	337
	Introducing DQL.....	337
	Quoting object type and attribute names .....	340
	NULLs, default values, and DQL.....	340
	Testing for default and NULL values .....	341
	Default values and aggregate functions.....	342
	Sorting and nulls .....	342
	Repeating attributes in queries .....	343
	Modifying repeating attributes .....	343
	Adding new values.....	343
	Inserting values .....	344
	Appending values .....	344
	Updating values .....	345
	Deleting values.....	345
	Forcing index correspondence in query results .....	346
	Performance note for Sybase or MS SQL Server users.....	347
	Querying virtual documents .....	347
	Full-text searching and virtual documents.....	348
	Querying registered tables .....	349
	Referencing registered tables in queries.....	349
	Security controls.....	349
	Default object-level permissions and table permits.....	350
	Caching queries.....	350
	Using execquery in an application .....	351
	Privileges, permissions, and queries.....	354
<b>Appendix A</b>	<b>Using DQL Hints</b> .....	355
	General guideline for all.....	355
	SQL_DEF_RESULT_SET N.....	356
	FORCE_ORDER .....	357
	RETURN_TOP N.....	358
	Database-specific implementations .....	358
	SQL Server.....	358

Subqueries and the hint .....	358
DB2 .....	358
Oracle and Sybase .....	359
Effects of a SEARCH clause .....	359
Recommended use .....	359
OPTIMIZE_TOP N .....	360
FETCH_ALL_RESULTS N.....	360
SQL Server, the hint, and subqueries .....	361
OPTIMIZATION_LEVEL level_1 level_2.....	361
UNCOMMITTED_READ .....	361
IN and EXISTS .....	362
ROW_BASED.....	363
Effects on returned results.....	363
Effects on query syntax rules .....	364
FTDQL and NOFTDQL.....	364
Including multiple hints limiting rows returned .....	365
Passthrough hints .....	365
Syntax .....	365
Error handling and debugging .....	366
<b>Appendix B Implementing Java Evaluation of Docbasic Expressions .....</b>	<b>367</b>
Overview of Docbasic expression handling .....	367
Migrating Docbasic expressions to Java.....	369
Disabling or re-enabling Java evaluation .....	371
Docbasic expression components support .....	372
Operators.....	372
Supported functions for Java evaluation.....	373
Unsupported functions for Java evaluation.....	378
Supported constants .....	379
Unsupported constants .....	380
Implicit objects .....	381
<b>Appendix C DQL Quick Reference .....</b>	<b>383</b>
The DQL statements .....	383
DQL reserved words .....	390
<b>Appendix D Document Query Language examples .....</b>	<b>393</b>
Basic examples .....	393
The simplest format .....	393
Using the WHERE clause .....	394
Searching repeating attributes in a WHERE Clause .....	394
Using aggregate functions.....	395
Using the GROUP BY clause .....	395
Using the HAVING clause.....	396
The ORDER BY clause .....	396
Using the asterisk (*) in queries .....	396
Searching cabinets and folders .....	397
Querying registered tables .....	398
Querying virtual documents .....	398
Determining the components .....	399

## List of Figures

Figure 2-1.	Sample virtual document .....	127
Figure 4-1.	Query processing .....	353
Figure B-1.	Architectural implementationn of Docbasic expressions for object types.....	368
Figure D-1.	Virtual document model .....	399

# List of Tables

Table 1-1.	Valid ranges for floating point literals .....	14
Table 1-2.	Default character string formats for dates.....	16
Table 1-3.	Arithmetic operators.....	30
Table 1-4.	Valid comparison operators for DQL statements.....	30
Table 1-5.	Predicates for columns and single-valued attributes.....	31
Table 1-6.	Predicates for repeating attributes .....	32
Table 2-1.	ALTER GROUP syntax description.....	43
Table 2-2.	ALTER TYPE syntax description.....	45
Table 2-3.	Alter Type operations .....	47
Table 2-4.	CHANGE...OBJECT syntax description.....	61
Table 2-5.	CREATE GROUP syntax description.....	65
Table 2-6.	CREATE...OBJECT syntax description.....	67
Table 2-7.	CREATE TYPE syntax description .....	71
Table 2-8.	DB2 tablespace page sizes and associated maximum row lengths .....	73
Table 2-9.	DELETE syntax description.....	87
Table 2-10.	DELETE...OBJECT syntax description .....	89
Table 2-11.	DROP GROUP syntax description .....	93
Table 2-12.	DROP TYPE syntax description.....	94
Table 2-13.	EXECUTE syntax description .....	96
Table 2-14.	Administration methods by category for the EXECUTE statement .....	97
Table 2-15.	GRANT syntax description .....	103
Table 2-16.	INSERT syntax description.....	105
Table 2-17.	REGISTER syntax description.....	108
Table 2-18.	REVOKE syntax description.....	113
Table 2-19.	SELECT syntax description .....	115
Table 2-20.	DQL standard hints .....	149
Table 2-21.	UNREGISTER syntax description .....	153
Table 2-22.	UPDATE syntax description.....	155
Table 2-23.	UPDATE...OBJECT syntax description .....	158
Table 3-1.	Administration methods by category .....	169
Table 3-2.	BATCH_PROMOTE arguments.....	175
Table 3-3.	CHECK_CACHE_CONFIG arguments .....	179
Table 3-4.	Attributes in the CHECK_CACHE_CONFIG result object .....	180
Table 3-5.	CHECK_RETENTION_EXPIRED arguments.....	182
Table 3-6.	CHECK_SECURITY arguments.....	186
Table 3-7.	DB_STATS arguments.....	191
Table 3-8.	Query result object attributes for DB_STATS administration method.....	191

Table 3–9.	DELETE_REPLICA arguments .....	193
Table 3–10.	DO_METHOD arguments.....	197
Table 3–11.	Attributes of the query result object returned by DO_METHOD .....	200
Table 3–12.	DROP_INDEX arguments .....	206
Table 3–13.	ESTIMATE_SEARCH arguments.....	208
Table 3–14.	EXEC_SQL arguments .....	211
Table 3–15.	EXPORT_TICKET_KEY arguments .....	213
Table 3–16.	GET_FILE_URL arguments .....	217
Table 3–17.	GET_INBOX arguments.....	219
Table 3–18.	SysObject-related attribute names for GET_INBOX query result objects .....	221
Table 3–19.	GET_PATH arguments.....	225
Table 3–20.	HTTP_POST arguments.....	228
Table 3–21.	Query result object attributes for HTTP_POST.....	230
Table 3–22.	IMPORT_REPLICA arguments.....	234
Table 3–23.	IMPORT_TICKET_KEY arguments.....	236
Table 3–24.	LIST_RESOURCES arguments.....	240
Table 3–25.	Collection attributes for LIST_RESOURCES .....	240
Table 3–26.	Collection attributes for LIST_RESOURCES .....	242
Table 3–27.	LIST_SESSIONS arguments.....	244
Table 3–28.	Complete information returned by LIST_SESSIONS (BRIEF_INFO=F) .....	244
Table 3–29.	Brief information returned by LIST_SESSIONS (BRIEF_INFO=T).....	246
Table 3–30.	Query result object attributes for LIST_TARGETS.....	248
Table 3–31.	LOG_OFF arguments.....	251
Table 3–32.	MAKE_INDEX arguments .....	253
Table 3–33.	MARK_FOR_RETRY arguments.....	258
Table 3–34.	MIGRATE_CONTENT arguments .....	261
Table 3–35.	MODIFY_TRACE arguments .....	269
Table 3–36.	MOVE_INDEX arguments .....	271
Table 3–37.	PING arguments .....	273
Table 3–38.	PURGE_AUDIT arguments.....	275
Table 3–39.	PUSH_CONTENT_ATTRS arguments .....	285
Table 3–40.	RECOVER_AUTO_TASKS argument.....	288
Table 3–41.	REGISTER_ASSET arguments .....	290
Table 3–42.	Queue item attribute values set by REGISTER_ASSET .....	291
Table 3–43.	REORGANIZE_TABLE arguments .....	292
Table 3–44.	REPLICATE arguments.....	295
Table 3–45.	RESTORE_CONTENT arguments.....	300
Table 3–46.	ROLES_FOR_USER arguments.....	302
Table 3–47.	SET_APIDEADLOCK arguments .....	304
Table 3–48.	Valid operation names for SET_APIDEADLOCK.....	305
Table 3–49.	SET_CONTENT_ATTRS arguments.....	307
Table 3–50.	Example settings for content metadata attributes in content objects.....	310
Table 3–51.	SET_OPTIONS arguments .....	312



Table 3–52.	Trace options for SET_OPTIONS .....	313
Table 3–53.	SET_STORAGE_STATE arguments .....	316
Table 3–54.	SYNC_REPLICA_RECORDS arguments .....	321
Table 3–55.	TRANSCODE_CONTENT arguments.....	325
Table 3–56.	Queue item attributes set by TRANSCODE_CONTENT .....	327
Table 3–57.	UPDATE_STATISTICS arguments .....	329
Table 3–58.	WEBCACHE_PUBLISH arguments .....	332
Table 3–59.	Valid arguments for ARGUMENTS .....	332
Table 4–1.	DQL basic query statements .....	338
Table 4–2.	Default attribute values by datatype .....	340
Table 4–3.	Predicates that test for NULL and default values .....	341
Table A–1.	Comparison of logical reads with and without default result sets.....	356
Table A–2.	Example of object-based query results.....	363
Table A–3.	Example of row-based query results .....	364
Table B–1.	Arguments for the Docbasic expression migration methods .....	369
Table B–2.	dmc_SetJavaExprEnabled arguments.....	371
Table B–3.	Docbasic operators supported by Java evaluation .....	372
Table B–4.	Docbasic functions supported for Java evaluation .....	373
Table B–5.	Docbasic functions not supported for Java evaluation .....	378
Table B–6.	Docbasic constants supported for Java evaluation.....	379
Table B–7.	Constants not supported for Java evaluation .....	380
Table C–1.	Summary of FTDQL query rules.....	387
Table C–2.	DQL reserved words.....	390



# Preface

---

This manual is the reference manual for Documentum's Document Query Language, supported by Content Server. It is a companion to *Content Server API Reference Manual*, *EMC Documentum Object Reference Manual*, *Content Server Fundamentals*, *Content Server Administrator's Guide*, and *Documentum Distributed Configuration Guide*.

## Intended audience

This manual is written for application developers and system administrators and any others who want to build a content or work-group management application that uses DQL. It assumes that you are familiar with the concepts of document processing, object-oriented programming, and client-server applications. It also assumes working knowledge of SQL.

## Conventions

This manual uses the following conventions in the syntax descriptions and examples.

### Syntax conventions

Convention	Identifies
<i>italics</i>	A variable for which you must provide a value.
[ ] square brackets	An optional argument that may be included only once
{ } curly braces	An optional argument that may be included multiple times

## Terminology changes

Two common terms are changed in 5.3 and later documentation:

- Docbases are now called repositories, except where the term “docbase” is used in the name of an object or attribute (for example, docbase config object).
- DocBrokers are now called connection brokers.

## Revision history

The following changes have been made to this document.

### Revision history

Revision Date	Description
September 2005	Initial publication

## DQL Language Elements

This chapter describes the building blocks of a DQL statement, including:

- [Literals, page 13](#), which describes the literal formats for the Documentum datatypes
- [Special keywords, page 19](#), which describes the special keywords that you can use in DQL queries
- [Functions, page 20](#), which describes the functions that you can use in DQL queries
- [Predicates, page 29](#), which describes the predicates that you can use in expressions in queries
- [Logical operators, page 38](#), which describes the logical operators supported by DQL
- [DQL reserved words, page 40](#), which cross-references you to a list of the words reserved in DQL

### Literals

Literals are values that are interpreted by the server exactly as they are entered. Content Server recognizes five types of literals:

- [Integer literals, page 13](#)
- [Floating point literals, page 14](#)
- [Character string literals, page 15](#)
- [ID literals, page 15](#)
- [Date literals, page 15](#)

### Integer literals

An integer literal specifies any whole number and is expressed in the following format:

[+ | -] *n*

where *n* is any number between -2147483647 and +2147483647.

DQL does not support the negative integer value -2147483648 because this number is not supported in a number of relational databases. If you enter this number, your results are unpredictable.

## Floating point literals

A floating point literal specifies any number that contains a decimal point and is expressed in the following format:

```
5.347
21.
0.45
.66
-4.12
```

A floating point literal can also be expressed in scientific notation. For example:

```
10.4e-6
```

or

```
-3.6E7
```

or

```
12e-3
```

DQL accepts either uppercase or lowercase in the notation.

If you assign an integer literal to an attribute that has a floating point datatype, the system automatically converts the integer to a floating point number.

The underlying RDBMS determines the maximum and minimum values that you can assign as a floating point literal. [Table 1-1, page 14](#) lists the ranges for supported databases.

**Note:** Do not reset the decimal symbol at the operating system level to a comma. Doing so results in incorrect execution of some Documentum Administration jobs.

**Table 1-1. Valid ranges for floating point literals**

RDBMS	Range	Significant digits
Oracle	1.0 x 10 <sup>-129</sup> to 9.99 x 10 <sup>-129</sup>	15
DB2	1.0 x 10 <sup>-307</sup> to 1.798 x 10 <sup>+308</sup>	15
MS SQL Server and Sybase	1.7e-308 to 1.7e+308	15

## Character string literals

Character string literals are strings of printable characters and are enclosed in single quotes. You cannot place non-printable characters, such as line feeds or carriage returns, in a character string literal. To include a single quote as part of the literal, include it twice. For example:

```
'The company's third quarter results were very good.'
```

If you are running against an Oracle RDBMS, a character string literal can contain a maximum of 2,000 characters.

If you are running against an MS SQL Server or Sybase RDBMS, a character string literal can contain a maximum of 255 characters.

If you are running against a DB2 RDBMS, a character string literal contain a maximum of 32,672 characters.

## ID literals

An ID literal specifies an object ID as a 16-character string enclosed in single quotes. ID literals are generally used in DQL queries. For example:

```
SELECT "object_name" FROM "dm_document"
WHERE "r_object_id" = '099af3ce800001ff'
```

Note that you cannot use ID literals to assign object IDs to objects. Object IDs are assigned automatically by the server when the object is created.

## Date literals

A date literal specifies a date using the following syntax:

```
DATE(date_value [, 'pattern'])
```

*date\_value* can be defined using any of the valid character string formats representing a date, or it can be one of the keywords that represent dates.

There are some date formats that the server can interpret as more than one date. For example, 03/04/96 can be interpreted as March 4, 1996 or as April 3, 1996. To resolve this, some formats require you to specify what pattern the server uses to interpret the date.

Valid dates range from 01/01/1753 (January 1, 1753) to 12/31/9999 (December 31, 9999).

The following paragraphs describe the character string formats and keywords you can use to specify a date. When you use a character string format, you must enclose *date\_value* in single quotes.

## Default formats

Table 1–2, page 16 lists the character string formats for *date\_value* that are accepted by default:

**Table 1-2. Default character string formats for dates**

Format	Examples
<i>mm/dd/[yy]yy</i>	DATE('03/24/1989') DATE('4/7/1992')
<i>dd-mon-[yy]yy</i>	DATE('4-Apr-1975')
<i>month dd[, ] [yy]yy</i>	DATE('January 1, 1993')
<i>mon dd [yy]yy</i>	DATE('March 23 1990')
the client's localized short date format	DATE('30-11-1990') (Assumes short date format <i>dd-mm-yyyy</i> is defined on the client machine)
ANSI format ( <i>dow mon dd hh:mm:ss yyyy</i> )	No example

Although it is not required, you can specify a pattern for any of these formats except the short date format and the ANSI format. You cannot specify a pattern when you enter a date using either of those two formats.

When using the formats listed in Table 1–2, page 16, the following rules apply:

- It is not necessary to include leading zeros for months or days that are represented as a single digit.
- You can abbreviate a month's name to three letters.
- If you enter only the year and not the century, the server uses the century defined by the current date in the server machine. For example, 03/23/95 is interpreted as March 23, 1995 before the year 2000 and is interpreted as March 23, 2095 after the year 2000.

## Short date formats

The server accepts a client's localized short date format as a valid character string format as long as the format contains only numeric characters. For example, *dd-mmm-yy* is not an accepted short date format.

Windows, Solaris, and AIX client platforms provide a default short date format. The Windows platform also lets you define your own default short date format. (Refer to [Defining short date formats, page 182](#) of the *Content Server Administrator's Guide* for information about defining short date formats.) For Sun and HP clients, which do not have a default short date format, the server assumes the default format is *mm/dd/yy hh:mi:ss*.



If the locally defined short date format conflicts with one of the default input formats, the locally defined format takes precedence. For example, assume that the locally defined format is *dd/mm/yyyy*. If a user wants to enter March 14, 1994 and enters 03/14/1994 (*mm/dd/yyyy*), the server interprets this as the 3rd day of the 14th month of the year 1994 and returns an error because there is no 14th month.

If you use the pattern argument to specify a format for a date, that pattern takes precedence over the short date format.

## ANSI format

You can also specify the date using the ANSI format (*dow mon dd hh:mm:ss yyyy*). However, DQL ignores the time fields. (To enter a date and time using the ANSI format, use the Set method.)

## Other character string formats

The following character string formats for *date\_value* are also acceptable but require a pattern argument:

```
[dd/]mm/[yy]yy[hh:mi:ss]
[yy]yy/mm[/dd] [hh:mi:ss]
[mon-][yy]yy [hh:mi:ss]
month[,] [yy]yy [hh:mi:ss]
```

If you do not enter the time (hh:mi:ss), the server assumes the time is 00:00:00. If you do not enter the day or month, the server assumes the first day of the month or the first month of the year, respectively.

For example, suppose the pattern argument is *[dd/]mm/[yy]yy [hh:mi:ss]*. The following date literal is interpreted as April 1, 1996 00:00:00 (assuming the current century is 1900).

```
DATE ('04/96', 'mm/yy')
```

When you specify *date\_value*, you can use any character except a space or an alphanumeric character as the delimiter. You must use the same delimiter to separate all elements of the date or time. For example, if you use a forward slash to separate *dd* and *mm*, you must use a forward slash to separate *mm* and *yy*. Similarly, if you use a period to separate *hh* and *mi*, you must use a period to separate *mi* and *ss*.

The delimiter that you use in the date can be different from the delimiter that you use in the time portion. For example, the following is valid:

```
'23-04-96 12.32.04'
```

It is not necessary to use the same delimiter for the *date\_value* and pattern arguments. For example, the following function call is valid:

```
DATE('23-04-96 12.32.04','dd/mm/yy hh:mi:ss')
```

## Date literal keywords

Four keywords represent dates. They are:

- **TODAY**, which returns the current date. The time defaults to 00:00:00. For example,

```
SELECT "supervisor_name", "object_name" FROM "dm_workflow"
WHERE "r_start_date" <= DATE(TODAY)
AND "r_runtime_state"=1
ORDER BY 1
```

- **NOW**, which returns the current date and time. For example,

```
SELECT "supervisor_name", "object_name" FROM "dm_workflow"
WHERE "r_start_date" <= DATE(NOW) AND "r_runtime_state"=1
ORDER BY 1
```

- **YESTERDAY**, which returns the current date minus one day. The time defaults to 00:00:00. For example,

```
SELECT * FROM dm_document
WHERE "r_creation_date" >= DATE(YESTERDAY)
AND
"r_creation_date" <= DATE(TODAY)
AND
ANY "authors" IN (john,henry,jane)
```

- **TOMORROW**, which returns the current date plus one day. The time defaults to 00:00:00. For example,

```
SELECT "r_act_name", "object_name", "process_id"
FROM "dm_workflow"
WHERE ANY "r_pre_timer" >= DATE(TOMORROW)
ORDER BY 2
```

## Date output formats

By default, the server uses the client's localized short date format to output dates to the client. If the client's format represents years using two digits, dates in the current century are displayed using two digits. However, to avoid ambiguity, years in other centuries are displayed using all four digits (1834 or 1792).

The default short date formats, by platform, are:

- Windows NT 4.0 and Windows 2000: The default short date format is the format specified by selecting Control Panel>Regional Settings>Date/Time.
- UNIX/SunOS: The format is assumed to be *mm/dd/yy hh:mi:ss*.
- UNIX/Solaris and UNIX/AIX: The default short date format is the format defined by the machine's locale. (Locale is set by the UNIX `setlocale` command. Refer to

[Defining short date formats, page 182](#) in the *Content Server Administrator's Guide* for information about setting this value.)

- UNIX/HP: The format is assumed to be *mm/dd/yy hh:mi:ss*.

**Note:** The session config attribute `r_date_format` contains the date format that the server returns to a client for a given session.

If the date is entered as `NULLDATE`, then the output is the string `NULLDATE`.

You can override the default format using an argument to the `Get` method. The `Get` method returns the value of a specified attribute. If that attribute has a `Date` datatype, you can include a pattern argument for the `Get` method that defines the format in which the date is returned. For more information, refer to [Get, page 224](#) of the *Content Server API Reference Manual*.

Dates that are included in error messages or a dump file are displayed in ANSI date format.

## Special keywords

DQL includes some special keywords for use in queries. These keywords have special meanings for Content Server when used in a DQL query. (They cannot be used in API methods.) The keywords are:

- `USER`, which identifies the current user.

You can use `USER` in comparison operations in queries. For example,

```
SELECT "process_id", "object_name" FROM dm_workflow
WHERE supervisor_name=USER
```

- `TRUE` and `FALSE`, which represent the Boolean true and false.

You can use `TRUE` and `FALSE` in comparison operations involving any attribute having a `BOOLEAN` datatype. For example,

```
SELECT * FROM "dm_user"
WHERE "r_is_group" = TRUE
```

- `DM_SESSION_DD_LOCALE`, which represents the data dictionary locale most appropriate for the client's session locale.

This keyword is particularly useful if the client session locale has no exact match recognized by the server. If you use this keyword, the server will return the information from the best matching locale. For example, suppose the server recognizes three locales, English (`en`), French (`fr`), and Spanish (`es`) and the client session locale is French Canadian (`fr_cn`). Suppose the client issues the following query:

```
SELECT type_name, label_text from dmi_dd_type_info where
nls_key='fr_cn'
```

The query returns nothing because the locale `fr_cn` is not recognized by the server. However, the following query returns the information from the French locale:

```
SELECT type_name, label_text from dmi_dd_type_info where  
nls_key=DM_SESSION_DD_LOCALE
```

The server checks whether the locale identified in the client's session locale, `fr_cn`, exists in the repository. Since `fr_cn` isn't found, the server attempts to find a good match among the existing locales. It determines that `fr` (French) is a good match and uses that locale to execute the query.

If a good match isn't found, the server uses the default locale to execute the query.

## Functions

Functions are operations on values. DQL recognizes three groups of functions and two unique functions:

- [Scalar functions, page 20](#)

Scalar functions operate on one value and return one value.

- [Aggregate functions, page 22](#)

Aggregate functions operate on a set of values and return one value.

- [Date functions, page 24](#)

Date functions operate on date values.

- [The ID function, page 27](#)

The ID function, a unique function recognized by DQL, is used in the FOLDER and CABINET predicates and in the IN DOCUMENT and IN ASSEMBLY clauses.

- [The MFILE\\_URL function, page 28](#)

The MFILE\_URL function returns URLs to content files and renditions in particular format.

## Scalar functions

The three scalar functions are:

- UPPER
- LOWER
- SUBSTR

## UPPER

The UPPER function takes one argument and returns the uppercase of that value. The value supplied as the argument must be a character string or an attribute that has a character string datatype.

For example, the following statement returns the object names of all documents that have the word government in their title. Because LIKE returns only exact matches, the UPPER function is used to ensure that all instances are found, regardless of the case (upper, lower, or mixed) in which the word appears in the title.

```
SELECT "object_name" FROM "dm_document"  
WHERE UPPER("title") LIKE '%GÖVERMENT%'
```

## LOWER

The LOWER function takes one argument and returns the lowercase of that value. The value supplied as the argument must be a character string or an attribute that has a character string datatype.

For example, the following statement returns the subjects in lowercase of all documents owned by regina:

```
SELECT LOWER("subject") FROM "dm_document"  
WHERE "owner_name" = 'regina'
```

## SUBSTR

The SUBSTR function returns some or all of a particular string. Its syntax is:

```
substr(string_value, start[, length])
```

The value of *string\_value* can be a literal string or a column or attribute name. If you specify a column or attribute, the server uses the value in that column or attribute as the string value. The attribute you specify can be either a single-valued attribute or a repeating attribute.

The *start* argument identifies the starting position, in the specified string, of the substring you want returned. Positions are numbered from the left, beginning at 1. (If you specify 0 as the start, the server automatically begins with 1.) The argument must be an integer.

The *length* argument defines how many characters should be returned in the substring. Length is an optional argument. If you don't specify a length, the default behavior of the function differs depending on the RDMBS you are using. For Oracle and DB2, the default is the entire string value or the full width of the column if an attribute or column is specified. For all other databases, the function returns 1 character.

For example, suppose you have a document subtype called `purchase_order`, which has an attribute called `order_no`. The values in this attribute are 10-digit numbers, with the last three digits representing a specific sales region. The following statement returns all documents for which the last three digits in the `order_no` attribute are 003:

```
SELECT "r_object_id","order_no" FROM "purchase_order"  
WHERE SUBSTR("order_no",8,3) = '003'
```

You can also use the `SUBSTR` function with the `SELECT` DQL statement (refer to [Select, page 115](#)). For example:

```
SELECT SUBSTR("emp_name",1,4) AS short_name FROM "employee"
```

You must use the `AS` clause option to rename the query result attribute that holds the return value of the `SUBSTR` function. If you do not, the server cannot return the result of the function.

If you specify a repeating attribute, the function returns one value for each value in the attribute.

You cannot use the `SUBSTR` function in assignment statements.

Additionally, you cannot specify `SUBSTR` in a `LIKE` predicate. Refer to [Pattern matching with LIKE, page 34](#) for more information about using `LIKE`.

## Aggregate functions

The five aggregate functions are:

- `COUNT`
- `MIN`
- `MAX`
- `AVG`
- `SUM`

### COUNT

The `COUNT` function counts values. The syntax is:

```
COUNT ([DISTINCT] name | *)
```

The `name` argument must identify an attribute or column. You can count all values for the attribute or column or you can include the `DISTINCT` keyword to count the number of unique values.

Using an asterisk directs the server to count all items that match specified criteria. For example, the following statement counts all documents that belong to the user named grace:

```
SELECT COUNT(*) FROM "dm_document"
WHERE "owner_name" = 'grace'
```

## MIN

The MIN function returns the minimum value in a given set of values. The syntax is:

```
MIN(DISTINCT name | [ALL] value_expression)
```

The DISTINCT *name* argument directs the server to first select all distinct values from the set (no duplicates) and then return the minimum value. *name* can be an attribute or column name. Note that for this function, the DISTINCT name option has little meaning.

[ALL] *value\_expression* directs the server to return the minimum value found in the set of values specified by the *value\_expression* argument. *value\_expression* can be any valid numeric expression or an attribute or column name. The keyword ALL is optional; whether it is specified or omitted, all of the values in the set are evaluated.

For example, assuming that rental\_charges is a user-defined object type, the following statement returns the minimum rent charged for a two-bedroom apartment:

```
SELECT MIN("charge") FROM "rental_charges"
WHERE "style" = 'apt' AND "bedroom" = 2
```

## MAX

The MAX function returns the maximum value in a given set of values. The syntax is:

```
MAX(DISTINCT name | [ALL] value_expression)
```

The DISTINCT *name* argument directs the server to first select all distinct values from the set (no duplicates) and then return the maximum value. *name* can be an attribute or column name. Note that for the MAX function, this option has little meaning.

[ALL] *value\_expression* directs the server to return the maximum value found in the set of values specified by the *value\_expression* argument. *value\_expression* can be any valid numeric expression or an attribute or column name. The keyword ALL is optional; whether it is specified or omitted, all of the values in the set are evaluated.

For example, assuming that rental\_charges is a user-defined object type, the following statement returns the maximum rent charged for a two-bedroom apartment:

```
SELECT MAX("charge") FROM "rental_charges"
WHERE "style" = 'apt' AND "bedroom" = 2
```

## AVG

The AVG function returns an average. The syntax is:

```
AVG(DISTINCT name | [ALL] value_expression)
```

The DISTINCT *name* option directs the server to select all distinct values from the set (no duplicates) and return the average of those distinct values. *name* can be an attribute or column name.

[ALL] *value\_expression* directs the server to return the average value found in the set of values specified by the *value\_expression* argument. The value of *value\_expression* can be any valid numeric expression or an attribute or column name. Use the optional keyword ALL to include all of the values in the set in the operation.

For example, assuming that rental\_charges is a user-defined object type, the following statement returns the average rent charged on a two-bedroom apartment:

```
SELECT AVG("charge") FROM "rental_charges"  
WHERE "style" = 'apt' AND "bedroom" = 2
```

## SUM

The SUM function returns a total. The syntax is:

```
SUM(DISTINCT name | [ALL] value_expression)
```

The DISTINCT *name* option directs the server to select all distinct values from the set (no duplicates) and return the sum of those distinct values. *name* can be an attribute or column name.

[ALL] *value\_expression* directs the server to return the sum of the values in the set of values specified by the *value\_expression* argument. The value of *value\_expression* can be any valid numeric expression or an attribute or column name. Use the optional keyword ALL to include all of the values in the set in the operation.

For example, assuming that rent\_records is a user-defined object type that contains payment records of tenants, the following statement provides a total of the rents paid in May:

```
SELECT SUM("rent_amt") FROM "rent_records"  
WHERE "mon" = 'may' AND UPPER("paid") = 'Y'
```

## Date functions

The four date functions are:

- DATEDIFF (refer to [DATEDIFF](#) , page 25)



- DATEADD (refer to [DATEADD](#), page 26)
- DATEFLOOR (refer to [DATEFLOOR](#), page 26)
- DATETOSTRING (refer to [DATETOSTRING](#), page 27)

## DATEDIFF

The DATEDIFF function subtracts two dates and returns a number that represents the difference between the two dates. The syntax is:

```
DATEDIFF(date_part, date1, date2)
```

The *date\_part* argument defines the units in which the return value is expressed. Valid values are year, month, week, and day.

The *date1* and *date2* arguments specify the dates to be subtracted. *date1* is subtracted from *date2*. You can specify the dates as date literals or the names of single-valued attributes with a date datatype.

For example, the following statement uses the DATEDIFF function to return all tasks that were started a month or more late:

```
SELECT "task_number", "supervisor_name"
FROM dm_tasks_queued
WHERE
DATEDIFF(month, "plan_start_date", "actual_start_date")>=1
```

This next example returns all tasks that were started more than one week ago:

```
SELECT "task_number", "r_task_user"
FROM dm_tasks_queued
WHERE DATEDIFF(week, "actual_start_date", DATE(TODAY))>=1
```

If the repository is using Oracle or DB2, the return value is a floating point number.

If the repository is using DB2, the server assumes 365 days per year and 30.42 days per month (the mean number of days in a month). These assumptions can cause return values that differ from the expected value. To illustrate, the following example, which asks for the number of days between March 1, 1996 and Feb 1, 1996, returns 30.42 instead of 28:

```
DATEDIFF(day, date('03/01/1996 0:0:0'),
date('02/01/1996 0:0:0'))
```

If the repository is using MS SQL Server or Sybase, the return value is an integer for all units except day. If you specify day in the function, the return value is a floating point.

The MS SQL Server and Sybase implementations round up if the difference is one half or greater of the specified unit. The implementations round down if the difference is less than one half of the specified unit. To illustrate, the following example, which asks for the difference between March 1, 1996 and July 1, 1996 expressed as years, returns 0 because the difference is only 4 months.

```
DATEDIFF(year, date('03/01/1996'), date('07/01/1996'))
```

## DATEADD

The DATEADD function adds a number of years, months, weeks, or days to a date and returns the new date. The syntax is:

```
DATEADD(date_part, number, date)
```

The *date\_part* argument defines the units that are being added to the specified date. Valid date parts are year, month, week, and day.

The *number* argument defines how *date\_part* values are being added to the date. For example, the following statement uses the DATEADD function to obtain all tasks that started more than a week ago but are not yet finished:

```
SELECT "task_number", "supervisor_name" FROM dm_tasks_queued
WHERE DATEADD(week, 1, "actual_start_date") < DATE(TODAY)
AND "actual_completion_date" IS NULLDATE
AND NOT("actual_start_date" IS NULLDATE)
```

## DATEFLOOR

The DATEFLOOR function rounds a given date down to the beginning of the year, month, or day. The syntax is:

```
DATEFLOOR(date_part, date)
```

The *date* argument is the name of a date attribute. The function rounds the value of the attribute down to the beginning of the unit specified as the value of *date\_part*. Valid *date\_part* values are year, month, and day.

For example, suppose that a document's *r\_creation\_date* attribute has a value of March 23, 1996 at 9:30:15 am. Using the DATEFLOOR function on this attribute returns these values:

Specifying	Returns
<pre>DATEFLOOR(year, "r_creation_ date")</pre>	January 1, 1996 at 00:00:00
<pre>DATEFLOOR(month, "r_creation_ date")</pre>	March 1, 1996 at 00:00:00
<pre>DATEFLOOR(day, "r_creation_date")</pre>	March 23, 1996 at 00:00:00

When you include DATEFLOOR in the selected values list of a SELECT statement, you must use the AS clause to assign a name to the column representing the returned values. For example:

```
SELECT DATEFLOOR(month, "r_creation_date") AS Created...
```

## DATESTRING

The DATESTRING function returns a date as a character string in a particular format. The syntax is:

```
DATESTRING(date, 'format')
```

The *date* argument is the name of a date attribute. The *format* argument defines how you want the character string formatted. The format can be any of the valid character string input formats for date literals, subject to the following RDBMS restrictions and qualifications:

- MS SQL Server and Sybase do not support month. If you specify a format that contains month, such as *month dd yyyy*, the date is returned in the default short date format.
- Oracle and DB2 always return the month name as a fixed-length 9-character string. This means that if the name of the month is shorter than 9 characters, the string value is padded.

Suppose that a document's *r\_creation\_date* attribute has a value of May 14, 1995. Here are some examples of the values returned by the DATESTRING function using this date and a variety of formats.

Specifying	Returns
<code>DATESTRING("r_creation_date", 'dd-mon-yy')</code>	14-May-95
<code>DATESTRING("r_creation_date", 'mm/dd/yy')</code>	05/14/95
<code>DATESTRING("r_creation_date", 'month yy')</code>	May 95

## The ID function

The ID function is used to identify a folder or cabinet whose contents you want to search. You can use the ID function in FOLDER and CABINET predicates and IN DOCUMENT and IN ASSEMBLY clauses. The syntax is:

```
ID('object_id')
```

The *object\_id* argument identifies the folder, cabinet, or virtual document you want to search. The object must reside in the current repository.

Use the literal 16-character representation of the object's ID. For example, the following statement returns all documents in the specified folder with titles containing the characters Pond:

```
SELECT * FROM "dm_document"  
WHERE FOLDER (ID('099af3ce800001ff'))  
AND "title" LIKE '%Pond%'
```

## The MFILE\_URL function

The MFILE\_URL function returns URLs to content files or renditions associated with a document. Only those objects on which the user has at least Read permission are returned if MFILE\_URL is included in a selected values list.

The syntax of MFILE\_URL is:

```
MFILE_URL('format',page_no,'page_modifier')
```

The arguments are ANDed together to determine which URLs are returned. Only the format argument is required. page\_no and page\_modifier are optional.

The format argument restricts the returned URLs to those pointing to files in the specified format. Enclose the format value in single quotes. If you specify format as an empty string, Content Server takes the format value from the returned object's a\_content\_type attribute.

The page\_no argument restricts the returned URLs to those pointing to content files associated with given page number. If you don't include page\_no, the value defaults to 0, which represents the first content file associated with an object. To indicate that all content pages should be considered, define page\_no as -1. If you define page\_no as -1, Content Server returns all URLs that satisfy the other arguments regardless of the page with which the content files are associated.

The page\_modifier argument restricts the returned URLs to those pointing to content files that have the specified page\_modifier. The page\_modifier argument must be enclosed in single quotes. If you specify a value for page\_modifier (other than an empty string), all the returned URLs point to renditions, as content files for primary pages have no page modifier. To return URLs for content files that have no page modifier, define page\_modifier as an empty string. (Content Server sets the page\_modifier attribute to an empty string by default when a user adds a content file to an object or saves a rendition without a page modifier.) If you don't include page\_modifier, Content Server returns the URLs that satisfy the other arguments regardless of whether their content files have a page modifier or not.

## Examples

The following statement returns the URLs to the jpeg\_th renditions of the first content pages of documents owned by ronaldh that have a page modifier of image1.

```
SELECT MFILE_URL('jpeg_th',0,'image1')
FROM "dm_document" WHERE "object_owner"='ronaldh'
```

The following example returns the owner name and object ID of all documents that have the subject prairie chickens. For each returned document, it also returns URLs to the primary content files associated with the document.

```
SELECT owner_name,r_object_id,MFILE_URL('',-1,'')
FROM "dm_document" WHERE "subject"='prairie_chickens'
```

## Predicates

Predicates are used in WHERE clauses to restrict the objects returned by a query. The WHERE clause defines criteria that each object must meet. Predicates are the verbs within the expression that define the criteria.

For example, the following statement returns only documents that contain the value approved in their keywords attribute:

```
SELECT "r_object_id", "object_name", "object_owner"
FROM "dm_document"
WHERE ANY "keywords" = 'approved'
```

In this example, the equal sign (=) is a predicate. It specifies that any object returned must have a keywords attribute value that equals (matches) the specified word (in this case, the word approved).

DQL recognizes these predicates:

- Arithmetic operators (refer to [Arithmetic operators, page 29](#))
- Comparison operators (refer to [Comparison operators, page 30](#))
- Column and attribute predicates (refer to [Column and attribute predicates, page 30](#))
- SysObject predicates (refer to [SysObject predicates, page 35](#))

## Arithmetic operators

Arithmetic operators perform arithmetic operations on numerical expressions. [Table 1-3, page 30](#) lists the arithmetic operators that you can use as predicates.

**Table 1-3. Arithmetic operators**

Operator	Operation
+	Addition
-	Subtraction
/	Division
*	Multiplication
**	Exponentiation

## Comparison operators

Comparison operators compare one expression to another. [Table 1-4, page 30](#) lists the comparison operators that can be used as predicates.

**Table 1-4. Valid comparison operators for DQL statements**

Operator	Relationship
=	Equal
>=	Greater than or equal to
<=	Less than or equal to
>	Greater than
<	Less than
<>	Not equal
!=	Not equal

## Column and attribute predicates

Column and attribute predicates let you compare values in a registered table's columns or an object's attributes to defined criteria. These predicates are divided into two groups. One group is used only when the search criterion specifies a column or single-valued attribute. The other group is used when the search criteria specifies a repeating attribute.

## Predicates for columns and single-valued attributes

The predicates in this group allow you to compare a value in a table column or single-valued attribute to defined criteria. [Table 1–5, page 31](#) lists the predicates in this group.

**Table 1-5. Predicates for columns and single-valued attributes**

Predicate	Description
IS [NOT] NULL	Determines whether the attribute is assigned a value. This predicate is useful only for registered table columns.
IS [NOT] NULLDATE	Determines whether the attribute is assigned a null date.
IS [NOT] NULLSTRING	Determines whether the attribute is assigned a null string.
IS [NOT] NULLINT	Determines whether the attribute is assigned a null integer value.
[NOT] LIKE <i>pattern</i> [ESCAPE <i>character</i> ]	Determines whether the attribute is like a particular pattern. (Refer to <a href="#">The ESCAPE character, page 35</a> for information about the ESCAPE clause option.)
[NOT] IN <i>value_list</i>	Determines whether the attribute is in one of a particular list of values.
[NOT] EXISTS ( <i>subquery</i> )	Determines whether the attribute matches a value found by the subquery.
<i>comparison_op</i> SOME ( <i>subquery</i> ) <i>comparison_op</i> ANY ( <i>subquery</i> )	Determines whether the comparison between the attribute value and the results of the subquery evaluate to TRUE in any case.
<i>comparison_op</i> ALL ( <i>subquery</i> )	Determines whether the comparison between the attribute value and the results of the subquery evaluate to TRUE in all cases.

For example, the following statement selects all documents that have a title that starts with the word Breads:

```
SELECT * FROM "dm_document"
WHERE "title" LIKE 'Breads%'
```

This next example selects all workflows that are supervised by one of the users named in the predicate:

```
SELECT "r_object_id", "supervisor_name" FROM "dm_workflow"
WHERE "supervisor_name" IN ('carrie', 'davidk', 'holly')
ORDER BY 2
```

## Predicates for repeating attributes

The predicates for repeating attributes let you compare the values in repeating attributes to some defined criteria. The basic syntax for repeating attribute predicates is:

```
[NOT] [ANY] predicate
```

The ANY keyword must be used whenever an expression references a repeating attribute predicate unless the query also includes the DQL hint, ROW\_BASED. If ROW\_BASED is included in the query, it is not necessary to include the ANY keyword with repeating attribute predicates.

Table 1–6, page 32 lists the predicates in this group.

**Table 1-6. Predicates for repeating attributes**

Predicate	Description
<i>attr_name</i> [NOT] LIKE <i>pattern</i> [ESCAPE <i>character</i> ]	Evaluates to TRUE if any value of the repeating attribute is [not] like a particular pattern. (Refer to <a href="#">The ESCAPE character, page 35</a> for information about the optional ESCAPE clause.)
<i>attr_name</i> IN ( <i>value_list</i> )	Evaluates to TRUE if any value of the attribute matches a value in the <i>value_list</i> . <i>Value_list</i> is a comma-separated list of values.
[IN EXISTS] <i>attr_name</i> IN ( <i>subquery</i> )	Evaluates to TRUE if any value of the attribute matches a value returned by the subquery.  IN and EXISTS are database hints that may enhance performance by changing the structure of the generated SQL query. Refer to <a href="#">Appendix A, Using DQL Hints</a> , for more information about these optional hints.
<i>attr_name</i> IS [NOT] NULL	Evaluates to TRUE if any value of the attribute is [not] null.
<i>attr_name</i> IS [NOT] NULLDATE	Evaluates to TRUE if any value of the attribute is [not] a nulldate.
<i>attr_name</i> IS [NOT] NULLSTRING	Evaluates to TRUE if any value of the attribute is [not] a null string.
<i>attr_name</i> IS [NOT] NULLINT	Evaluates to TRUE if any value of the specified repeating attribute is [not] a null integer value.
<i>attr_name comparison_op</i> <i>value_expression</i>	Evaluates to TRUE if the comparison operation is TRUE for any value of the attribute.



For example, the following statement returns the object names of all documents with an author whose name begins with a letter in the first half of the alphabet:

```
SELECT "object_name" FROM "dm_document"
WHERE ANY "authors" <= 'M'
```

You can use logical operators to build more complex predicate expressions, such as:

```
[NOT] ANY predicate AND|OR [NOT] ANY predicate {AND|OR
[NOT] ANY predicate}
```

For example, the following statement selects all documents that have Ingrid as an author and a version label of 1.0:

```
SELECT "r_object_id", "object_name" FROM "dm_document"
WHERE ANY "authors" = 'Ingrid' AND
ANY "r_version_label" = '1.0'
```

The predicate statement returns all objects that meet the specified criteria regardless of where the qualifying value is found in the attribute's value list. For example, look at the following statement:

```
SELECT "object_name" FROM "dm_document"
WHERE ANY "authors" IN ('jeanine','harvey') AND
ANY "keywords" = 'working'
```

This statement returns the name of all documents with either jeanine or harvey as a value anywhere in the authors attribute values list and with the keyword working as a value in any position in the list of values in the keywords attribute.

In some circumstances, you may want only objects for which the specified values are found in the same respective positions. For example, assume that the search finds an object for which jeanine is the value at position 3 in the authors attribute. You want the name of this object only if the keyword working is in position 3 of the keywords attribute.

You can impose this restriction by enclosing the predicate expression in parentheses:

```
[NOT] ANY (predicate AND|OR [NOT]predicate {AND|OR [NOT]predicate})
```

For example, the statement below returns the names of documents for which either jeanine and working or harvey and working (or both) are values of the authors and keywords attributes, respectively, and occupy corresponding positions in the attributes.

```
SELECT "object_name" FROM "dm_document"
WHERE ANY ("authors" IN ('jeanine','harvey')
AND "keywords" = 'working')
```

#### Notes:

- To return values at a matching index position level, the keyword ANY is outside the parentheses.
- You cannot specify an index position at which to look for values. You can specify only that the values be found in the same respective positions.
- Using the logical operator OR returns the same results as using the basic syntax, because the server returns the object if either predicate is true.

For more information and examples of querying repeating attributes, refer to [Repeating attributes in queries, page 343](#).

## Pattern matching with LIKE

The [NOT] LIKE predicate for both single-valued and repeating attributes lets you identify a pattern to match the attribute value. This pattern is specified as a character string literal. For example, the following predicate returns all objects whose subject contains the word Cake:

```
subject LIKE '%Cake%'
```

When you use LIKE, the value in the attribute must match the string exactly, including the case of each character. If you use NOT LIKE, then the comparison is TRUE for any value that does not match exactly.

Sometimes, however, you may not know the text of the character string. For example, you might want to retrieve all documents that have the word Cake in their title but not know all of the titles. For these instances, Documentum provides two pattern-matching characters. These characters serve as wildcards. The two characters are:

- The percent sign (%)
- The underbar (\_)

You can include the pattern-matching characters anywhere in a character string.

### The percent sign

The percent sign replaces 0 or more characters. For example, suppose the predicate contains:

```
"subject" LIKE 'Breads%'
```

This returns any object whose subject begins with the word Breads, regardless of the how many characters or words follow Breads.

### The underbar

The underbar replaces one character. For example, suppose the predicate contains:

```
"subject" LIKE 'Bread_'
```

This returns any object whose subject is the single word *Bread*, followed by a space, an *s* (Breads), or any other single, printable character.

## Matching cases

You can use the UPPER and LOWER functions to retrieve an object that matches a value regardless of the case of the letters in the value. For example, to retrieve any object with the word cake in the title, regardless of its position in the title or the case:

```
UPPER("title") LIKE '%CAKE%'
```

Using the UPPER function changes all titles to uppercase for the comparison so that the case of the characters is not a factor. You can use the LOWER function the same way:

```
LOWER("title") LIKE '%cake%'
```

## The ESCAPE character

There may be occasions when the pattern you want to match includes a percent sign (%) or an underbar (\_). To match either of those characters literally, you must specify an escape character and insert that character before the percent sign or underscore in the pattern. Use the optional ESCAPE clause to specify the escape character.

For example, suppose you want to find all documents whose object names contain an underscore. Because the underscore is interpreted as a wild card by default, you must define and use an escape character in the query:

```
SELECT "r_object_id" FROM "dm_document"
WHERE "object_name" LIKE '%\_%' ESCAPE '\'
```

In the above example, the backslash is defined as the escape character. Placed in the pattern directly ahead of the underscore, it tells the server to treat the underscore as a literal character rather than a wild card.

In addition to the wild-card characters, an escape character can also escape itself. For example, suppose you want to match the string %\\_ and that you have defined the backslash as the escape character. Your LIKE predicate would look like this:

```
LIKE '%\_\_%' ESCAPE '\'
```

In the above example, the backslash character escapes not only the percent sign and underscore but also itself.

Escape characters can escape only the two wild-card characters and themselves. If you insert an escape character before any other character in a pattern, it generates an error.

You can use any printable character as an escape character.

## SysObject predicates

Three predicates restrict the search specified in a FROM clause.

When you specify an object type in a FROM clause, the server examines that type and its subtypes for any objects that fulfill the conditions specified in the rest of the query. However, sometimes you may want to limit the search to specific subtypes, folders, or cabinets. These three predicates allow you to do so. They are:

- TYPE
- FOLDER
- CABINET

## The TYPE predicate

The TYPE predicate restricts the search to objects of a single type or one of its subtypes. The syntax is:

```
TYPE(type_name)
```

The *type\_name* argument must identify a subtype of a type specified in the FROM clause.

For example, the following statement retrieves all documents of the type `legal_doc` or `accounting_doc` or a subtype:

```
SELECT * FROM "dm_document"  
WHERE TYPE("legal_doc") OR TYPE("accounting_doc")
```

Using the TYPE predicate provides one way to select from more than one type. For example, to retrieve all documents or workflow process definitions that have been created after a particular date, you could use the following statement:

```
SELECT "r_object_id", "object_name", "owner_name", "r_creation_date"  
FROM "dm_sysobject" WHERE TYPE("dm_document") OR TYPE("dm_process")
```

## The FOLDER predicate

The FOLDER predicate identifies what folders to search. The syntax is:

```
[NOT] FOLDER(folder_expression {,folder_expression} [,DESCEND])
```

The *folder\_expression* argument identifies a folder in the current repository. You can't search a remote folder (a folder that doesn't reside in the current repository). Valid values are:

- An ID function
- The ID function (described in [The ID function, page 27](#)) identifies a particular folder.
- A folder path

A folder path has the format:

```
/cabinet_name{/folder_name}
```

Enclose the path in single quotes. Because cabinets are a subtype of folder, you can specify a cabinet as the folder.

- The keyword DEFAULT

The keyword DEFAULT directs the server to search the user's default folder. Note that a user's default folder is the same as the user's default cabinet (because cabinets are a subtype of folders).

The DESCEND keyword directs the server to search the specified folder or folders and any local folders directly or indirectly contained within that folder or folders. The predicate does not search any contained, remote folders. The specified folder or folders may have no more than 25,000 nested folders if you include the DESCEND keyword.

When the search finds a remote folder, it returns the object ID of the associated mirror object in the current repository. It does not return the object's ID in the remote repository.

DESCEND applies to all folders specified in the predicate. If you want to search one folder without descending but descend through another folder, include two folder predicates in your statement and OR them together. For example:

```
FOLDER ('/cakes/yeasted', DESCEND) OR FOLDER (DEFAULT)
```

The keyword NOT directs the server not to search a particular folder.

## The CABINET predicate

The CABINET predicate restricts the search to a particular cabinet. Its syntax is:

```
[NOT] CABINET(cabinet_expression [,DESCEND])
```

The *cabinet\_expression* argument must identify a cabinet that resides in the current repository. Valid values are:

- An ID function

The ID function (described in [The ID function, page 27](#)) must specify a cabinet ID.

- A folder path

The folder path must identify a cabinet. Its format is:

```
/cabinet_name
```

Enclose the path in single quotes.

The keyword DESCEND directs the server to search the specified cabinet and any folders directly or indirectly contained in the cabinet that reside in the current repository. The predicate does not search any contained, remote folders.

When the search finds a remote folder, it returns the object ID of the associated mirror object in the current repository. It does not return the object's ID in the remote repository.

The keyword NOT directs the server not to search a particular cabinet.

## Logical operators

Logical operators are used in WHERE clauses. DQL recognizes three logical operators:

- AND
- OR
- NOT

### AND and OR

The AND and OR operators join two or more comparison expressions to form a complex expression that returns a single Boolean TRUE or FALSE. The syntax for these operators is:

```
expression AND | OR expression
```

Two expressions joined using AND return TRUE only if both expressions are true. For example, the following complex expression returns TRUE when both of its component expressions are true, and FALSE when only one of them is true:

```
"title" = 'Yeast Breads of England' AND "subject" = 'Breads'
```

Similarly, the following expression also returns TRUE only if both conditions are true:

```
"subject" = 'Breads' AND ANY "authors" IN ('clarice', 'michelle', 'james')
```

Two expressions joined using OR return TRUE if either expression is true. For example, the following expression returns true when either comparison is true:

```
"subject" = 'Cheeses' OR "owner_name" = 'donaldrm'
```

### NOT

The NOT operator reverses the logic of an expression. The following expression is a simple example:

```
"subject" = 'Cheeses'
```

When the server encounters this expression, it is looking for objects that have a subject attribute with the value Cheeses. The server returns TRUE if the subject value is Cheeses and FALSE if it is not.

Now, look at the same expression with the NOT operator:

```
NOT("subject" = 'Cheeses')
```

In this instance, the server looks for the objects with subject attributes that do not contain the value Cheeses. The server returns TRUE if the subject is not Cheeses and FALSE if the subject is Cheeses.

## Order of precedence

You can join any number of expressions together with logical operators. Content Server™ imposes no limit. (Note that your underlying RDBMS may impose a limit.) The resulting complex expression is evaluated from left to right in the order of precedence of the operators. This order, from highest to lowest, is:

NOT  
AND  
OR

To illustrate, look at the following example:

```
NOT expression1 AND expression2 OR expression3 AND expression4
```

The server resolves this expression as follows:

1. Evaluates NOT expression1
2. ANDs the result of Step 1 with the result of expression2
3. ANDs the results of expression3 and expression4
4. ORs the results of Steps 2 and 3

You can change the order of evaluation by using parentheses. For example:

```
NOT expression1 AND (expression2 OR expression3) AND expression4
```

The server resolves this expression as follows:

1. Evaluates NOT expression1
2. ORs the results of expression2 and expression3
3. ANDs the results of Step 1 and Step 2
4. ANDs the results of Step 3 with the result of expression4

Similarly, you can use parentheses to change the precedence of the NOT operator:

```
NOT(expression1 AND expression2 AND expression3)
```

The complex expression inside the parentheses is evaluated first and the NOT operator applied to its result.

## DQL reserved words

[Table C-2, page 390](#) lists the DQL reserved words. If you use any of these as object or attribute names, you must enclose the name in double quotes whenever it is used in a DQL statement.



## DQL Statements

This chapter contains descriptions of the DQL statements. The description of each statement includes:

- Syntax
- Argument descriptions (if any)
- Return value (if appropriate)
- Detailed information about required permissions and usage
- Related statements
- Example

**Quoting Object Type Names and Attribute Names:** Documentum recommends that you put double quotes around all object type names and attribute names referenced in applications. Doing that will ensure that the names will not conflict with any DQL reserved words or any words reserved by your underlying RDBMS. Documentum object type names and attribute names will not generate conflicts, but using the quotes in applications will make sure that no conflicts are generated by user-defined type or attribute names. To encourage this best practice, the DQL examples in our manuals use the double quotes.

**Comments within DQL Statements:** DQL does not accept inline comment. An inline comment is a comment embedded within a DQL statement.

# Abort

**Purpose** Cancels an explicit transaction.

## Syntax

```
ABORT [TRAN[SACTION]]
```

## General notes

The ABORT statement terminates a transaction that was opened with the BEGIN TRAN statement. Any changes made while the transaction was open are rolled back. They are not saved to the repository.

The ALTER TYPE statement does not participate in transactions. If you alter a type as part of an explicit transaction and then issue an ABORT statement, the changes you made to the type are not reversed.

You can include either keyword, TRAN or TRANSACTION.

## Related statements

[Begin Tran, page 60](#)

[Commit, page 64](#)

# Alter Group

**Purpose**        Modifies a user group.

## Syntax

```
ALTER GROUP group_name ADD members
ALTER GROUP group_name DROP members
ALTER GROUP group_name SET ADDRESS email_address
ALTER GROUP group_name SET PRIVATE TRUE|FALSE
```

## Syntax description

**Table 2-1. ALTER GROUP syntax description**

Variable	Description
<i>group_name</i>	Identifies the group you want to modify. Use the group's name. The name can be a string, which must be enclosed in single quotes, or an identifier, which need not be in quotes.
<i>members</i>	Identifies users, groups, or both to add or drop from the group. You can specify user names representing individual users or names representing groups. Use a comma-separated list to specify multiple names. Alternatively, you can use a SELECT statement to identify the members (illustrated in <a href="#">Examples, page 44</a> ).  When you are identifying an individual user, the name must be the value of the user's user_name attribute.
<i>email_address</i>	Defines an electronic mail address for a group. Specify this as a character string literal. You can use any email address that is valid for your environment. To remove a group's email address, specify email_address as an empty string.

## Permissions

To alter a group, you must be either the group's owner or a user with Superuser user privileges.

## General notes

Each ALTER GROUP statement can perform only one type of operation. That is, you cannot add and drop members in the same statement. Nor can you set an email address in the same statement that adds or drops members.

## The SET PRIVATE clause

SET PRIVATE sets the `is_private` attribute for the group. Setting the attribute to TRUE makes the group a private group; setting it to FALSE makes the group a public group.

## Related statements

[Create Group, page 65](#)

[Drop Group, page 93](#)

## Examples

The following example adds two users to the group called superusers:

```
ALTER GROUP superusers ADD steve,chip
```

The next example uses a SELECT statement to identify which users to add to the engineering group:

```
ALTER GROUP engineering ADD (SELECT "user_name"  
FROM "dm_user" WHERE "user_os_name" LIKE '%eng%')
```

The final example defines an email address for the engineering group:

```
ALTER GROUP engineering  
SET ADDRESS 'engineering@lion.stellar.com'
```

# Alter Type

**Purpose**        Modifies an object type.

## Syntax

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
type_modifier_list [PUBLISH]

ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_modifier_clause) [PUBLISH]

ALTER TYPE type_name
ADD attribute_def {,attribute_def} [PUBLISH]

ALTER TYPE type_name
DROP attribute_name {,attribute_name} [PUBLISH]
```

## Syntax description

**Table 2-2. ALTER TYPE syntax description**

Variable	Description
<i>type_name</i>	Identifies the type to alter. Use the name defined for the type when it was created.
<i>policy_id</i>	Identifies the default lifecycle for the type. Use the policy's object ID.  Including the FOR POLICY clause requires at least Version permission on the lifecycle identified by <i>policy_id</i> .
<i>state_name</i>	Identifies one state in the default lifecycle. Use the state's name as defined in the dm_policy object.

Variable	Description
<i>type_modifier_list</i>	<p>Lists one or more specifications that set or alter type data dictionary information for the type. Valid type modifiers are:</p> <p><i>update_modifier</i>  <i>mapping_table_specification</i>  <i>constraint_specification</i>  <i>component_specification</i>  <i>type_drop_clause</i></p> <p>Separate multiple type modifiers with commas.</p> <p>Refer to <a href="#">Type modifiers, page 81</a> for information about the syntax and use of each specification.</p>
<i>attribute_modifier_clause</i>	<p>Defines the change you want to make to the attribute. The syntax for this clause is:</p> <p><i>attribute_name domain</i></p> <p>or</p> <p><i>attribute_name (attribute_modifier_list)</i></p> <p><i>domain</i> is any valid DQL datatype.</p> <p><i>attribute_modifier_list</i> lists one or more modifiers that set or alter data dictionary information for the attribute. Valid attribute modifiers are:</p> <p><i>update_modifier value_assistance_specification</i>  <i>mapping_table_specification default_specification</i>  <i>constraint_specification attribute_drop_clause</i></p> <p>Separate multiple modifiers with commas.</p> <p>Refer to <a href="#">Attribute modifiers, page 74</a> for information about the syntax and use of each attribute modifier.</p>

Variable	Description
<i>attribute_def</i>	<p>Defines the attributes you want to add to the type. <i>attribute_def</i> has the following syntax:</p> <pre><i>attribute_name</i> <i>domain</i> [REPEATING] (<i>attribute_modifier_list</i>)</pre> <p><i>attribute_name</i> is the name of the attribute and <i>domain</i> is any valid DQL datatype. The keyword REPEATING defines the attribute as a repeating attribute.</p> <p><i>attribute_modifier_list</i> lists one or more modifiers that define data dictionary information for the attribute. Refer to the preceding description of <i>attribute_modifier_clause</i> for a list of valid attribute modifiers.</p> <p>Separate multiple modifiers with commas.</p>
<i>attribute_name</i>	<p>For the DROP option, <i>attribute_name</i> identifies the attribute that you want to remove from the specified type.</p>

## Permissions

To issue an ALTER TYPE statement that changes locale-specific data dictionary information, your session locale must match exactly one of the locales identified in the `dd_locales` attribute of the `docbase` config.

[Table 2-3, page 47](#) lists the ALTER TYPE operations and describes who can perform them and on which types.

**Table 2-3. Alter Type operations**

Alteration	Who can do it	To which types
Set default ACL	Type owner or Superuser	All types
Set default storage area	Type owner or Superuser	All types
Set or drop data dictionary information	Type owner or Superuser	All types, with two exceptions. See * below table.
Add read/write attributes	Type owner or Superuser	User-defined only
Add read-only attributes	Superuser	User-defined only

Alteration	Who can do it	To which types
Drop read/write attributes	Type owner or Superuser	User-defined only
Lengthen character string attributes	Type owner or Superuser	User-defined only

\*It is not possible to add value assistance or constraints to a system-defined object type.

(For the definition of a read and write or read-only attribute, refer to [Chapter 1, Object Basics](#), in the *Object Reference Manual*. For details about using indexed attributes, refer to [Chapter 8, Full-Text Indexing](#), in the *Content Server Administrator's Guide*.)

## General notes

You cannot change the definition of a system-defined type. You can only set a variety of defaults and data dictionary information. You can change the definition of a user-defined type. (Refer to [Table 2–3, page 47](#) for a summary of valid operations, the types on which they can be performed, and who can perform them.)

You cannot issue an ALTER TYPE statement in an explicit transaction.

After you execute an ALTER TYPE statement, the changes are not reflected to users until the global type cache is updated. The updating is automatic but may take several minutes.

## Localization

If you change data dictionary information that is locale-specific, such as label\_text, it is changed only in the current locale (defined in the session's session\_locale attribute).

If you change data dictionary information that isn't locale-specific, such as a constraint definition, it is changed in all locales.

## PUBLISH

Including PUBLISH causes the server to publish the data dictionary information for the object type or attribute immediately. Including PUBLISH publishes the changes made in the ALTER TYPE statement and any other changes to the type or attribute that are not yet published.

If you don't include PUBLISH, the information is published the next time the Data Dictionary Publish job runs.

You can include PUBLISH with any ALTER TYPE statement except those that add or drop attributes from the list of indexed attributes.



It isn't necessary to include PUBLISH if you are dropping an attribute from the type definition. The data dictionary information for the attribute is automatically removed in such cases.

## Type modifiers

Type modifiers set or change some attributes of the type info object and data dictionary information for the type. For example, you can set the default storage area, ACL, or default lifecycle or add constraints to the type. You can also drop data dictionary information.

You cannot include a type modifier if the ALTER TYPE statement is executing inside an explicit transaction.

### update\_modifiers

Update modifiers set or remove attribute values in the dm\_nls\_dd\_info, dm\_dd\_info, and dmi\_type\_info objects for the type. The dm\_nls\_dd\_info and dm\_dd\_info objects hold data dictionary information for the type. The dmi\_type\_info object holds non-structural information about the type. You can set attributes in the dmi\_type\_info object that define the type's default ACL, default permissions, default storage area, and default group.

An update modifier is also used to define a default lifecycle for the type.

### Setting or removing data dictionary values

You can set any attribute in the dm\_nls\_dd\_info and dm\_dd\_info objects that are applicable to types. Use one of the following statement clauses:

```
SET attribute_name[[index]]=value
APPEND attribute_name=value
INSERT attribute_name[[index]]=value
REMOVE attribute_name[[index]]
TRUNCATE attribute_name[[index]]
```

The *attribute\_name* is the name of an attribute defined for the dm\_nls\_dd\_info or dm\_dd\_info object type. The nls\_dd\_info or dm\_dd\_info attribute must be applicable to object types and settable by users.

Include *index* if the attribute is a repeating attribute. Its interpretation varies:

- For a SET operation, the index defines which value to set.
  - If the operation is adding a new value to the attribute, the index must identify the next available position in the list. If the SET operation is replacing an existing value, specify the index of the existing value.
- For an INSERT operation, the index defines where to insert the new value.
  - Existing values are renumbered after the insertion.
- For a REMOVE operation, the index defines which value to remove.
- For a TRUNCATE operation, the index defines the starting position for the truncation.

All values in the list, beginning at that position, are truncated.

The APPEND statement clause doesn't require an index value for repeating attributes because it automatically puts the new value at the end of the repeating attribute's value list.

You must enclose the index in square brackets.

*value* is a literal value. If you use a SET statement clause and the attribute is single-valued, *value* can be NULL.

You cannot use the SET, INSERT, or APPEND statement clauses to alter any of the following attributes:

- From dm\_dd\_info type
  - parent\_id, default\_value, cond\_value\_assist, cond\_computed\_expr, val\_constraint, unique\_keys, foreign\_keys, primary\_key
- From the dm\_nls\_dd\_info type
  - parent\_id

### Defining the default ACL

A type's default ACL does not define access to the type. Users can assign a type's default ACL to any object of the type they create. Use the following syntax to set a type's default ACL:

```
ALTER TYPE type_name
SET DEFAULT ACL acl_name [IN acl_domain]
```

The value of *acl\_name* is the ACL's object name. The *acl\_domain* value is the name of its owner. The owner's name is either the user who created the ACL or the alias dm\_dbo, representing the repository owner. The combination of the ACL name and domain uniquely identifies the ACL within the repository. (For more information about ACLs and their names and implementation, refer to [Managing ACLs, page 387](#) in the *Content Server Administrator's Guide* or [Assigning an ACL, page 138](#) in *Content Server Fundamentals*.)

If either the name or domain includes a character (such as a space) that requires you to enclose the string in single quotes, then you must enclose both strings in single quotes.

If the default ACL is NULL, the server uses the default ACL defined for the type's supertype as the default. To set the default ACL to NULL, use the following syntax:

```
ALTER TYPE type_name SET DEFAULT ACL NULL
```

If you set the default ACL to NULL, the server automatically sets the default\_owner\_permit, default\_group\_permit, and default\_world\_permit attributes for the type to NULL also.

### Defining the default storage area

The default storage area for a type is where Content Server stores all content files associated with objects of that type. (Users can change the storage area for an individual object.)

To set the default storage area for a type, use the following syntax:

```
ALTER TYPE type_name SET DEFAULT STORAGE [=] storage_area_name
```

*storage\_area\_name* is the name of the storage object representing the storage area.

### Defining the default group

To set the default group for a type, use the following syntax:

```
ALTER TYPE type_name SET DEFAULT GROUP [=] group_name
```

*group\_name* can be a character string literal or an identifier. If it is a character string literal, you must enclose it in single quotes.

To set the default group to NULL, use:

```
ALTER TYPE type_name SET DEFAULT GROUP [=] NULL
```

### Defining a default lifecycle

A lifecycle describes the life cycle of an object. Typically, the SysObject type and its subtypes have default lifecycles. If an object type has a default lifecycle defined for it, users or applications can attach the default simply by specifying the keyword Default in the Attach method. (An object is not attached to a default lifecycle automatically. Users or the application must explicitly issue an Attach method. Defining a default lets users or applications attach an object to the default without requiring them to know the name or object ID of the default. Also, it allows you to change the default without requiring you to rewrite and recompile any applications that reference the default.)

To set the default lifecycle for a type, use the following syntax:

```
ALTER TYPE type_name  
SET DEFAULT BUSINESS POLICY [=] chronicle_id  
[VERSION version_label]
```

*chronicle\_id* is the object ID of the root version of the lifecycle. The VERSION clause identifies which version of the lifecycle you want to apply to the type. If you do not specify a version, the default is the CURRENT version.

### mapping\_table\_specification

A mapping table specification is typically used to map a user-friendly character string value to an underlying value that may not be as readable or easy to understand. A mapping table can also be used to map localized or type-specific values to an underlying value.

For example, Desktop Client uses a mapping table defined at the type level to display a list of the display config objects appropriate for the object type. The actual names of the

display config objects are mapped to more user-friendly strings. (Display config objects represent subsets of type attributes that have a common display definition.)

If you define a mapping table at the type level, the mappings apply to that object type and its subtypes.

### **constraint\_specification**

You can define the following constraints in a type modifier list:

- Primary key
- Unique key
- Foreign key
- Check

You must have Sysadmin or Superuser user privileges to create a foreign key constraint. You cannot add a constraint to a system-defined object type.

The syntax for adding a constraint to a type is:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
ADD constraint_specification
```

You cannot include the FOR POLICY clause for primary, unique, or foreign constraints. The FOR POLICY clause can only be included when you are defining a check constraint. The clause defines the check constraint as applicable only when instances of the type are in the specified lifecycle state. You must have at least Version permission on the lifecycle identified by *policy\_id*.

If you include the FOR POLICY clause, *type\_name* must be the primary type for the lifecycle identified in *policy\_id*. The *policy\_id* is the object ID of the dm\_policy object for the lifecycle. *state\_name* is the name of the state for which you are defining the constraint. State names are defined in the lifecycle's dm\_policy object.

To define a check constraint that applies to all states for an object, do not include the FOR POLICY clause. Any string literals in the check constraint must be ASCII characters.

For an explanation of each constraint specification, refer to [constraint\\_specification, page 78](#).

### **component\_specification**

Component specifications identify which components can operate on objects of the type. Use the following syntax to identify a component for a type:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
component_specification
```

Include the optional FOR POLICY clause to define the component as applicable only when objects of the type are in the specified state. You must have at least Version permission on the lifecycle identified by *policy\_id* to include the FOR POLICY clause.

If you include the clause, *type\_name* must be the primary type for the lifecycle identified in *policy\_id*. *policy\_id* is the object ID of the *dm\_policy* object for the lifecycle. *state\_name* is the name of the state for which you are defining the component. State names are defined in the lifecycle's *dm\_policy* object.

The format of a valid component specification is described in [component\\_specification](#), page 84.

## type\_drop\_clause

A type drop clause removes constraint and component definitions defined for a type. You cannot remove inherited constraints and component definitions.

- To drop a primary key constraint, use:

```
ALTER TYPE type_name
DROP PRIMARY KEY
```

- To drop a unique key constraint, use:

```
ALTER TYPE type_name
DROP UNIQUE KEY [(attribute_list)]
```

*attribute\_list* identifies the type attributes that constitute the key. If there is more than one attribute, separate them with a comma.

- To drop a foreign key constraint, you must have Sysadmin or Superuser user privileges. Use the following syntax:

```
ALTER TYPE type_name
DROP FOREIGN KEY [(attribute_list)]
REFERENCES type_name [(attr_list)]
```

*attribute\_list* identifies the type attributes that constitute the key. If there is more than one attribute, separate them with a comma. In the REFERENCES clause, *type\_name* identifies the object type that contains the referenced attributes and *attr\_list* names the referenced attributes.

- To drop check constraints, use:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
DROP CHECK
```

DROP CHECK drops all check constraints defined for the type. If you include the FOR POLICY clause, the statement removes the check constraints defined for the specified state. You must have at least Version permission on the lifecycle identified by *policy\_id* to include the FOR POLICY clause.

**Note:** To drop a single check constraint, use an update modifier to remove or truncate the *val\_constraint[x]*, *val\_constraint\_enf[x]*, and *val\_constraint\_msg[x]* attributes, where *x* is the index position of the check constraint you want to remove.

- To drop a component, use:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
DROP COMPONENTS
```

DROP COMPONENTS drops all components defined for the type. If you include the FOR POLICY clause, the statement removes the components defined for the specified state. You must have at least Version permission on the lifecycle identified by *policy\_id* to include the FOR POLICY clause.

## Attribute modifiers

Attribute modifiers set or drop data dictionary information for an attribute. Data dictionary information includes constraints, value assistance, default value definitions, and mapping information. Attribute modifiers set attributes of *dm\_nls\_dd\_info* and *dm\_dd\_info* objects.

You cannot include an attribute modifier if the ALTER TYPE statement is executing in an explicit transaction.

### update\_modifiers

An update modifier lets you directly set an attribute in a *dm\_dd\_info* or *nls\_dd\_info* object. Use the following syntax to include an update modifier:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_name (update_modifier))
```

Including the FOR POLICY clause makes the change to the attribute effective only when instances of the type are in the specified lifecycle state. You must have at least Version permission on the lifecycle identified by *policy\_id*.

If you include the clause, *type\_name* must be the primary type for the lifecycle identified in *policy\_id*. The *policy\_id* is the object ID of the *dm\_policy* object for the lifecycle. The *state\_name* is the name of the state for which you are defining the assistance. State names are defined in the lifecycle's *dm\_policy* object.

For a description of valid update modifiers for an attribute, refer to [update\\_modifier](#), page 74.

### value\_assistance\_specification

A value assistance specification identifies one or more values for an attribute. Typically, value assistance is used to populate a pick list of values for the attribute when it appears as a field on a dialog box.

You cannot add value assistance to a system-defined object type.

Use the following syntax to add or change the value assistance specification for an attribute:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_name (value_assistance_specification))
```

Including the FOR POLICY clause provides value assistance only when objects of the type are in the specified lifecycle state. You must have at least Version permission on the lifecycle identified by *policy\_id*.

If you include the clause, *type\_name* must be the primary type for the lifecycle identified in *policy\_id*. The *policy\_id* is the object ID of the dm\_policy object for the lifecycle. The *state\_name* is the name of the state for which you are defining the assistance. State names are defined in the lifecycle's dm\_policy object.

For a description of valid value assistance specifications, refer to [value\\_assistance\\_modifier](#), page 75.

### mapping\_table\_specification

A mapping table specification typically maps a list of character string values to a list of integer values. This is useful when you want to provide users with an easily understood list of values for an attribute that is an integer data type. Use the following syntax to add or change a mapping table specification:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_name (mapping_table_specification))
```

If you include the FOR POLICY clause, the mapped values are only available when instances of the type are in the specified lifecycle state. You must have at least Version permission on the lifecycle identified by *policy\_id*.

When you include the clause, *type\_name* must be the primary type for the lifecycle identified in *policy\_id*. The *policy\_id* is the object ID of the dm\_policy object for the lifecycle. The *state\_name* is the name of the state for which you are defining the mapping. State names are defined in the lifecycle's dm\_policy object.

For a description of valid mapping table specifications, refer to [mapping\\_table\\_specification](#), page 77.

### default\_specification

A default specification defines the default value for an attribute. When a new object of the type is created, the server automatically sets the attribute to the default value if no other value is specified by the user.

You cannot add an attribute and specify a default value for it in the same ALTER TYPE statement. You must use two ALTER TYPE statements: one to add the attribute and one to set its default value.

Use the following syntax to set or change a default specification:

```
ALTER TYPE type_name
MODIFY (attribute_name ([SET] default_specification))
```

The default value can be specified as:

- A literal value
- One of the following keywords: USER, NOW, TODAY, TOMORROW, YESTERDAY
- NULL (This option is not valid for repeating attributes.)

For a single-valued attribute, the default specification syntax is:

```
DEFAULT [=] default_value
```

For a repeating attribute, the default specification syntax is:

```
DEFAULT [=] (default_value {,default_value})
```

For example, the following ALTER TYPE statement sets the default value for one single-valued attribute and one repeating attribute in the object type mytype:

```
ALTER TYPE "mytype"  
MODIFY ("single_attr" (SET default=NOW)),  
MODIFY ("rep_attr" (SET default=(USER)))
```

The default value must be appropriate for the datatype of the attribute. For example, you cannot specify the keyword USER for a date attribute. You cannot specify NULL for a repeating attribute.

### **constraint\_specification**

You can specify the following constraints in an attribute modifier list:

- Primary key
- Unique key
- Foreign key
- Not null
- Check

You must have Sysadmin or Superuser user privileges to create a foreign key constraint. You cannot add a constraint to a system-defined object type.

The syntax for adding a constraint to an attribute is:

```
ALTER TYPE type_name  
[FOR POLICY policy_id STATE state_name]  
MODIFY (attribute_name (ADD constraint_specification))
```

You cannot include the FOR POLICY clause for primary key, unique key, or foreign key constraints. The FOR POLICY clause can only be included when you are defining a NOT NULL or check constraint. The clause defines the constraint as applicable only when instances of the type are in the specified lifecycle state. To define a NOT NULL or check constraint that applies to all states for the object, do not include the FOR POLICY clause.

If you include the clause, you must have at least Version permission on the lifecycle identified by *policy\_id*. The *type\_name* must be the primary type for the lifecycle identified in *policy\_id*. The *policy\_id* is the object ID of the dm\_policy object for the lifecycle. The *state\_name* is the name of the state for which you are defining the constraint. State names are defined in the lifecycle's dm\_policy object.

For an explanation of each constraint specification, refer to [constraint\\_specification](#), page 78.



## attribute\_drop\_clause

An attribute drop clause removes constraints, value assistance, or mapping information defined for the attribute. You cannot drop inherited constraints, value assistance, or mapping information.

- To drop a primary key constraint, use:

```
ALTER TYPE type_name
MODIFY (attribute_name (DROP PRIMARY KEY))
```

- To drop a unique key constraint, use:

```
ALTER TYPE type_name
MODIFY (attribute_name (DROP UNIQUE KEY
[(attribute_name)]))
```

*attribute\_name* identifies the attribute you are removing from the unique key.

- To drop a foreign key constraint, you must have Sysadmin or Superuser user privileges. Use the following syntax:

```
ALTER TYPE type_name
MODIFY (attribute_name (DROP FOREIGN KEY
[(attribute_name)] REFERENCES type_name [(attr_name)]))
```

*attribute\_name* identifies the attribute you are removing from the foreign key.

In the REFERENCES clause, *type\_name* identifies the object type that contains the referenced attributes and *attr\_name* identifies the referenced attribute. If *attr\_name* is not included, *r\_object\_id* is the default.

- To drop check constraints, use:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
DROP CHECK
```

DROP CHECK drops all check constraints defined for the type. If you include the FOR POLICY clause, the statement removes the check constraints defined for the specified state. You must have at least Version permission on the lifecycle identified by *policy\_id*.

**Note:** To drop a single check constraint, use an update modifier to remove or truncate the *val\_constraint[x]*, *val\_constraint\_enf[x]*, and *val\_constraint\_msg[x]* attributes, where *x* is the index position of the check constraint you want to remove.

- To drop value assistance, use the following syntax:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_name (DROP VALUE ASSISTANCE))
```

- To drop mapping information, use the following syntax:

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_name (DROP MAPPING TABLE))
```

## Modifying the type definition

You can modify the definition of a user-defined type by:

- Adding a new attribute
- Lengthening a character string attribute
- Dropping an attribute defined for the type

### Adding a new attribute

Use the following syntax to add a new attribute to a user-defined type:

```
ALTER TYPE type_name ADD attribute_def {,attribute_def}
```

*attribute\_def* defines the new attribute and has the following syntax:

```
attribute_name domain [REPEATING]  
(attribute_modifier_list)
```

The attribute name must follow the naming rules outlined in [Object type and attribute names, page 31](#) in the *Object Reference Manual*.

*domain* is any valid DQL datatype. The keyword REPEATING defines the attribute as a repeating attribute. Valid attribute modifiers are described in [Attribute modifiers, page 74](#).

Do not add more than 30 attributes in one ALTER TYPE statement. If the repository is running against DB2, the sum total of the lengths of all attributes defined for the type is constrained by the page size defined for the tablespace. [Table 2-8, page 73](#) lists the tablespace page sizes and the corresponding maximum row length allowed.

### Lengthening a character string attribute

The character string attribute must be defined for the object type. It cannot be an inherited attribute. The only valid change that you can make is to lengthen the attribute. The change is applied to all the type's subtypes and to all existing objects of the type and its subtypes.

For example, if a string attribute called `employee_name` is currently 24 characters, the following statement would change the length to 32 characters:

```
ALTER TYPE "test_result" MODIFY ("patient_name" char(32))
```

If you are running against Sybase, you can lengthen only one attribute in each execution of the statement.

### Dropping attributes from the type

Only attributes that are defined for the type can be dropped. You cannot drop an inherited attribute. Dropping an attribute also removes the attribute from all subtypes of the type and removes all data dictionary information for the attribute.

The drop operation fails if the attribute is an indexed attribute for the type and there are existing indexed objects of the type or its subtypes. (An indexed attribute is an attribute whose value is stored in the full-text index.)

**Note:** Some relational database management systems (for example, DB2) don't support dropping columns from a table (an attribute is a column in a table in the underlying RDBMS). For those databases, if you drop an attribute, the corresponding column in the table representing the type is not actually removed. If you later try to add an attribute to that type that has the same name as the dropped attribute, you will receive an error message.

## Related statements

[Create Type, page 71](#)

[Drop Type, page 94](#)

## Examples

This example sets the owner's object-level permission and the default group for the user-defined type called `legal_document`:

```
ALTER TYPE "legal_document"  
SET OWNER PERMIT browse,  
DEFAULT GROUP lawyers
```

This example adds an attribute to the user-defined type called `report_doc`:

```
ALTER TYPE "report_doc"  
ADD "monthly_total" integer
```

The next example changes the length of the `client_comments` attribute in the user-defined type called `case_report`:

```
ALTER TYPE "case_report"  
MODIFY ("client_comments" string(255))
```

This next example sets the default ACL for the document object type:

```
ALTER TYPE "dm_document"  
SET DEFAULT ACL generic_doc IN dm_dbo
```

This final example demonstrates the use of single quotes in setting the ACL:

```
ALTER TYPE "dm_document" SET DEFAULT ACL 'mktg default'  
IN 'marketing'
```

# Begin Tran

**Purpose**        Opens an explicit transaction.

## Syntax

```
BEGIN TRAN[SACTION]
```

## Permissions

Anyone can execute the BEGIN TRAN statement.

## General notes

The BEGIN TRAN or BEGIN TRANSACTION statement opens an explicit transaction. When you are working in an explicit transaction, none of the changes you make to files or attribute values in objects are saved until you issue a COMMIT statement to commit the changes.

## Related statements

[Abort, page 42](#)  
[Commit, page 64](#)

# Change...Object

**Purpose** Changes the object type of one or more objects.

## Syntax

```
CHANGE current_type OBJECT[S]
TO new_type[update_list]
[IN ASSEMBLY document_id [VERSION version_label] [DESCEND]]
[SEARCH fulltext search condition]
[WHERE qualification]
```

## Syntax description

**Table 2-4. CHANGE...OBJECT syntax description**

Variable	Description
<i>current_type</i>	Identifies the object type of the objects to change. Use the type's internal name (for example, dm_document). The type must be a SysObject or SysObject subtype (system- or user-defined).
<i>new_type</i>	Specifies the new object type for the objects. Use the type's internal name. Valid new types depend on the object's current type. Refer to <a href="#">Rules and constraints on changing types, page 63</a> for a complete description.
<i>update_list</i>	Specifies one or more change operations to perform on the objects you are changing. Valid formats are: <pre>set <i>attribute_name</i> = <i>value</i> set <i>attribute_name</i>[[<i>index</i>]] = <i>value</i> append [<i>n</i>] <i>attribute_name</i> = <i>value</i> insert <i>attribute_name</i>[[<i>index</i>]] = <i>value</i> remove <i>attribute_name</i>[[<i>index</i>]] truncate <i>attribute_name</i>[[<i>index</i>]]</pre>

If you specify more than one operation, use commas to separate the clauses. Refer to [The update operations, page 162](#) for details of these options.

Variable	Description
IN ASSEMBLY clause	Limits the statement to those objects that belong to a particular assembly. For details about the syntax and use, refer to <a href="#">The IN ASSEMBLY clause, page 137</a> .
SEARCH clause	Restricts the candidates for change to those that satisfy the full-text search condition. Refer to <a href="#">The SEARCH clause, page 138</a> for a description of its syntax and use.
WHERE clause	Defines a qualification used to restrict what objects are changed. Refer to <a href="#">The WHERE clause, page 142</a> for a description of a valid qualification.

## Return value

The CHANGE...OBJECT statement returns a collection identifier. The collection has one query result object, with one attribute called `objects_changed`. The attribute contains the number of objects changed.

## Permissions

To use this statement, you must have Delete permission for the objects you want to change.

To use the `update_list` option to change an object's owner (`owner_name` attribute), you must be have Superuser user privileges or Change Ownership privilege.

## General notes

The CHANGE...OBJECT[S] statement lets you change one or more objects from one type to another.

Note that when you change an object from a subtype to a supertype, you lose the values of the attributes that are inherited from the supertype.

Objects that meet the criteria in the statement are deleted from the repository and recreated as objects of the specified new type. The recreated objects retain their old object IDs and all their previous relationships. For example, if a document is annotated, changing that document to another type of document does not delete its annotations.

You cannot execute CHANGE...OBJECT on an object that is immutable.

## Rules and constraints on changing types

There are some constraints when you change an object's type:

- The object's current type and the new type must have the same type identifier, and they must be SysObjects or SysObject subtypes.
- The type identifier is the first two characters in an object ID. For example, all documents have 09 as their type identifier.
- The old and new types cannot be at the same level in the type hierarchy.

You can move objects up or down in the object hierarchy but not laterally. For example, the dm\_document and dm\_process object types are peers—they are at the same level because both are direct subtypes of dm\_sysobject. Now, suppose you have two dm\_document subtypes, proposal\_doc and report\_doc. Using CHANGE...OBJECT, you can change a proposal\_doc or report\_doc object to a dm\_document object or a dm\_document object to a proposal\_doc or report\_doc object, but you cannot change any of the document types or subtypes to dm\_process objects.

When you change an object to a type that is higher in the hierarchy, the server drops any attributes from the old type that are not in the new type. Similarly, when you move an object to a type that is lower in the hierarchy, the server adds any attributes in the new type that were not in the old type. Unless you use the update\_list option to set these new attributes, the server assigns them default values appropriate for their datatypes.

The optional clauses are applied in the following order:

- The SEARCH clause
- The IN ASSEMBLY clause
- The WHERE clause

## Related statements

[Create...Object](#), page 67

[Delete...Object](#), page 89

[Update...Object](#), page 158

## Example

This example changes all documents for which the subject is new book proposal to the document subtype book\_proposal:

```
CHANGE "dm_document" OBJECTS TO "book_proposal"
SET "department" = 'marketing'
WHERE "subject" = 'new book proposal'
```

# Commit

**Purpose**            Commits the changes made during an explicit transaction and closes the transaction.

## Syntax

```
COMMIT [TRAN[SACTION]]
```

## Permissions

Anyone can execute the COMMIT statement.

## General notes

The COMMIT statement closes a transaction that was opened with the BEGIN TRAN statement and commits to the repository any changes made to objects or files during the transaction.

You can include either TRAN or TRANSACTION.

## Related statements

[Abort, page 42](#)  
[Begin Tran, page 60](#)



# Create Group

**Purpose** Creates a user group.

## Syntax

```
CREATE [PUBLIC|PRIVATE] GROUP group_name
[WITH] [ADDRESS email_address] [MEMBERS members]
```

## Syntax description

Table 2-5. CREATE GROUP syntax description

Variable	Description
<i>group_name</i>	<p>Defines the name of the group that you are creating. This name must be unique among all group names and user names in the repository. You can specify the name as an identifier or as a character string literal.</p> <p><b>Note:</b> Content Server stores all group names in lowercase.</p>
<i>email_address</i>	<p>Specifies the electronic mail address of the group. Use a character string literal. You can specify any email address that is valid for your environment.</p>
<i>members</i>	<p>Specifies users to be included in the group. You can specify user names representing individual users, names representing other groups, or both. The names must appear as a comma-separated list. Alternatively, you can use a SELECT statement to populate the group.</p> <p>When you are specifying an individual user, the name must be the value of the user's user_name attribute.</p>

## Return value

The CREATE GROUP statement returns a collection whose result object has one attribute, new\_object\_id, which contains the object ID of the new group.

## Permissions

You must have Create Group, Sysadmin, or Superuser user privileges to create a group.

## General notes

When you create a group, you are its owner and can alter or delete the group.

## Public and private groups

The PUBLIC keyword creates the group as a public group. The PRIVATE keyword creates the group as a private group. When a user with Sysadmin or Superuser user privileges creates a group, the group is public by default. When a user with Create Group user privileges creates a group, the group is private by default.

The public or private setting is stored in the `is_private` attribute for the group. The setting is not used by Content Server. All groups, public or private, are visible in the `dm_group` type. This attribute is provided for use by user-written applications.

## Related statements

[Alter Group, page 43](#)

[Drop Group, page 93](#)

## Examples

The following example creates a group called `supers` whose members are all the users with the Superuser user privilege:

```
CREATE GROUP supers MEMBERS
(SELECT "user_name" FROM "dm_user"
 WHERE "user_privileges" >= 16)
```

The next example creates a group that includes all of Ron's co-workers but not Ron:

```
CREATE GROUP rons_baby_gift WITH MEMBERS
(SELECT "user_name" FROM "dm_user"
 WHERE "user_name" != 'Ron')
```

This final example creates a group and defines an email address for that group:

```
CREATE GROUP client_group
WITH ADDRESS 'john@jaguar.docu.com'
MEMBERS john,regina,kendall,maria
```

# Create...Object

**Purpose** Creates an object.

## Syntax

```
CREATE type_name OBJECT update_list
[,SETFILE 'filepath' WITH CONTENT_FORMAT='format_name']
{,SETFILE 'filepath' WITH PAGE_NO=page_number}
```

## Syntax description

**Table 2-6. CREATE...OBJECT syntax description**

Variable	Description
<i>type_name</i>	Identifies the type of object to create. Specify the name of the object type. You can use any valid type in the repository.
<i>update_list</i>	Specifies one or more operations you want to perform on the new object. Valid formats are: <pre>set <i>attribute_name</i> = <i>value</i> set <i>attribute_name</i>[[<i>index</i>]] = <i>value</i> append [<i>n</i>]<i>attribute_name</i> = <i>value</i> insert <i>attribute_name</i>[[<i>index</i>]] = <i>value</i> remove <i>attribute_name</i>[[<i>index</i>]] truncate <i>attribute_name</i>[[<i>index</i>]] [un]link '<i>folder path</i>' move [to] '<i>folder path</i>'</pre> <p>If you include multiple clauses in the update list, use commas to separate the clauses.</p>
SETFILE clause WITH CONTENT_FORMAT option	Adds the first content file to the new object. Refer to <a href="#">WITH CONTENT_FORMAT option, page 69</a> for details.
SETFILE clause WITH PAGE_NO option	Adds additional content to the object. Refer to <a href="#">WITH PAGE_NO option, page 69</a> for details.

## Return value

The CREATE...OBJECT statement returns a collection whose result object has one attribute, `object_created`, which contains the object ID of the new object.

## Permissions

Anyone can issue the CREATE...OBJECT statement. However, you must have Superuser privileges to include the SETFILE clause.

## General notes

The CREATE...OBJECT statement creates and saves a new object. As part of the process, you can set some of the object's attributes, specify a storage location for the object, and associate one or more content files with the object.

If you do not use the link update option to define a storage location, the new object is stored in your default folder or cabinet.

## The SETFILE clauses

A SETFILE clause adds content to the new object. You must have Superuser privileges to include the clause in the statement. Using the SETFILE clause is subject to the following conditions:

- The content must be a file. It cannot be a block of data in memory.
- The content file must be located in a directory visible to Content Server.
- You cannot add a file created on a Macintosh machine.
- The content cannot be stored in content-addressed storage or in a turbo store storage area.

Any object capable of having content may have multiple associated content files. All files must have the same content format. The content format is defined when you add the first content file to the object. All subsequent additions must have the same format. Consequently, specifying the format for content additions after the first file is not necessary. Instead, you must specify the content's position in the ordered list of content files for the object.

To add the first content, use the SETFILE clause with the WITH CONTENT\_FORMAT option.

To add additional content, use the SETFILE clause with the PAGE\_NO option.

You can't include both options in a single SETFILE clause.

## WITH CONTENT\_FORMAT option

Use this SETFILE option to add the first content file to a new object. The syntax is:

```
SETFILE 'filepath' WITH CONTENT_FORMAT='format_name'
```

The filepath must identify a location that is visible to Content Server.

The format name is the name found in the name attribute of the format's dm\_format object.

## WITH PAGE\_NO option

Use this SETFILE option to add additional content to a new object. The syntax is:

```
SETFILE 'filepath' WITH PAGE_NO=page_number
```

The filepath must identify a location that is visible to Content Server.

The page number identifies the file's position of the content file within the ordered contents of the new object. You must add content files in sequence. For example, you cannot add two files and specify their page numbers as 1 and 3, skipping 2. Because the first content file has a page number of 0, page numbers for subsequent additions begin with 1 and increment by 1 with each addition.

You cannot use the SETFILE clause with the PAGE\_NO option in a CREATE...OBJECT statement unless the statement contains a prior SETFILE clause with the CONTENT\_FORMAT option.

## Related statements

[Change...Object, page 61](#)

[Delete...Object, page 89](#)

[Update...Object, page 158](#)

## Examples

The following example creates a new document and sets its title and subject:

```
CREATE "dm_document" OBJECT
SET "title" = 'Grant Proposal',
SET "subject" = 'Research Funding'
```

This example creates a new document, sets its title, and adds two content files. The example is shown for both Windows and UNIX platforms.

On Windows:

```
CREATE "dm_document" OBJECT
SET "title" = 'Grant Proposal',
SETFILE 'c:\proposals\grantreq.doc'
WITH CONTENT_FORMAT='msww',
SETFILE 'c:\proposals\budget.doc' WITH PAGE_NO=1
```

On UNIX:

```
CREATE "dm_document" OBJECT
SET "title" = 'Grant Proposal',
SETFILE 'u12/proposals/grantreq.doc
WITH CONTENT_FORMAT='msww',
SETFILE 'u12/proposals/budget.doc' WITH PAGE_NO=1
```

# Create Type

**Purpose** Creates an object type.

## Syntax

```
CREATE TYPE type_name
[(attribute_def {,attribute_def})]
[WITH] SUPERTYPE parent_type
[type_modifier_list] [PUBLISH]
```

## Syntax description

Table 2-7. CREATE TYPE syntax description

Variable	Description
<i>type_name</i>	Names the new type. Use any valid name that is unique among the other user-defined type names in the repository. Types with names beginning with dm can only be created by a user with Superuser user privileges.
<i>attribute_def</i>	<p>The type name must consist of ASCII characters.</p> <p>Defines an attribute for the new type. You can define up to 30 attributes. Each attribute definition has the following format:</p> <pre><i>attribute_name</i> <i>domain</i> [REPEATING] [(<i>attribute_modifier_list</i>)]</pre> <p><i>attribute_name</i> names the attribute and <i>domain</i> defines its datatype. The attribute name must consist of ASCII characters. The domain can be any valid DQL datatype. If the datatype is a character string datatype, the <i>domain</i> specification must also include the length.</p> <p>REPEATING defines the attribute as a repeating attribute.</p>

Variable	Description
<i>attribute_def</i> continued	<p><i>attribute_modifier_list</i> defines data dictionary information for the attribute. Valid attribute modifiers are:</p> <p><i>update_modifier</i>  <i>value_assistance_specification</i>  <i>mapping_table_specification</i>  <i>default_specification</i>  <i>constraint_specification</i></p> <p>Separate multiple modifiers with commas.</p> <p>You cannot include an attribute modifier if the statement is inside an explicit transaction.</p> <p>Refer to <a href="#">Attribute modifiers, page 74</a> for descriptions of the attribute modifiers.</p>
<i>parent_type</i>	<p>Identifies the supertype of the new type. Valid supertypes are:</p> <p>dm_sysobject and its subtypes  dm_user and its subtypes  dm_relation  dm_state_extension  dm_state_type</p> <p>dm_email_message  user-defined types</p> <p>To create a type with no supertype, specify <i>parent_type</i> as NULL. This requires the Superuser user privilege.</p>
<i>type_modifier_list</i>	<p>Defines data dictionary information for the type. Valid type modifiers are:</p> <p><i>update_modifier</i>  <i>mapping table specification</i>  <i>constraint_specification</i>  <i>component_specification</i></p> <p>Separate multiple modifiers with commas.</p> <p>You cannot include a type modifier if the statement is inside an explicit transaction.</p> <p>Refer to <a href="#">Type modifiers, page 81</a> for descriptions of the type modifiers.</p>



## Return value

The CREATE TYPE statement returns a collection whose result object has one attribute, `new_object_ID`, which contains the object ID of the new object type.

## Permissions

You must have Create Type, Sysadmin, or Superuser privileges to create a new object type.

To define read-only attributes for a new type (attributes whose names begin with `r_`) or to create a type that has no supertype, you must have Superuser privileges.

If the statement sets locale-specific information for the new type or attributes, your session locale must match exactly one of the locales defined in the `dd_locales` attribute of the `docbase` config object.

## General notes

The user who creates a type becomes the type's owner.

You cannot include a CREATE TYPE statement in an explicit transaction.

Do not define more than 30 attributes in a single CREATE TYPE statement. If the type has more than 30 attributes, use ALTER TYPE to add the additional attributes. The total number of attributes in the new type cannot exceed the supported number of columns in a table in the underlying RDBMS. On DB2, the sum total of the lengths of the attributes defined for the type cannot exceed the maximum row length set by the page size of the tablespace. [Table 2-8, page 73](#) lists the tablespace page sizes and the corresponding maximum row length allowed.

**Table 2-8. DB2 tablespace page sizes and associated maximum row lengths**

Tablespace Page Size	Maximum Row Length, in bytes
4K	4005
8K	8101
16K	16,293
32K	32,677

Additionally, be sure to follow the naming rules for attributes described in [Object type and attribute names, page 31](#) in the *Object Reference Manual*.

## Specifying an ACL for the type

If one or more servers in the repository is using type-based ACL inheritance, it is recommended that you make sure that the object type has an ACL defined for the type. The ACL does not determine who can use the type, but serves as the default ACL for objects of the type when users create the objects without naming an ACL for the objects. After you create the object type, use ALTER TYPE to set the `acl_name` and `acl_domain` attributes for the object type if needed.

**Note:** Default ACL inheritance is defined in the server config object, in the `default_acl` attribute. If that attribute is set to 2, for type-based ACL inheritance, when a user creates an instance of the type and does not explicitly assign an ACL to the object, the server will assign the ACL associated with the object type.

## Localization

When you create a new type, any data dictionary information that you define for the type or its attributes in the statement is published in all locales identified in the `dd_locales` attribute of the `doibase` config object.

## PUBLISH

Including PUBLISH causes the server to publish the data dictionary information for the object type immediately. Publishing creates the `dd type info`, `dd attr info`, and `dd common info` objects that store the published data dictionary information for the object type.

If you don't include PUBLISH, the information is published the next time the Data Dictionary Publish job runs.

## Attribute modifiers

Attribute modifiers define data dictionary information for an attribute. Defining attribute modifiers sets attributes in data dictionary-related objects for the type. It does not set any `dm_type` or `dm_type_info` attributes for the type.

### **update\_modifier**

An *update\_modifier* specification can be:

```
SET attribute_name[[index]]=value
APPEND attribute_name=value
INSERT attribute_name[[index]]=value
REMOVE attribute_name[[index]]
TRUNCATE attribute_name[[index]]
```

*attribute\_name* is the name of an attribute defined for the `dm_nls_dd_info` or `dm_dd_info` object type. The `dm_nls_dd_info` or `dm_dd_info` attribute must be applicable to object type attributes and settable by users. If the attribute is a `dm_nls_dd_info` attribute, only the `dm_nls_dd_info` object specific to the current locale is updated.

Include *index* if the attribute is a repeating attribute. Its meaning varies:

- For a SET operation, the index defines which value to set.  
If the operation is adding a new value to the attribute, the index must identify the next available position in the list. If the SET operation is replacing an existing value, specify the index of the existing value.
- For an INSERT operation, the index defines where to insert the new value.  
Existing values are renumbered after the insertion.
- For a REMOVE operation, the index defines which value to remove.
- For a TRUNCATE operation, the index defines the starting position for the truncation.

All values in the list, beginning at that position, are truncated.

The APPEND statement clause doesn't require an index value for repeating attributes because it automatically puts the new value at the end of the repeating attribute's value list.

You must enclose the index in square brackets.

*value* is a literal value. If you use a SET statement clause and the attribute is single-valued, *value* can be NULL.

You cannot use the SET, INSERT, or APPEND formats against the following attributes:

- From `dm_dd_info` type  
parent\_id, default\_value, cond\_value\_assist, cond\_computed\_expr, val\_constraint, unique\_keys, foreign\_keys, primary\_key
- From the `dm_nls_dd_info` type  
parent\_id

### **value\_assistance\_modifier**

A value assistance modifier identifies one or more values for the attribute. Typically, value assistance is used to populate a pick list of values for the attribute when it appears as a field on a dialog box. A value assistance modifier has the following format:

```
VALUE ASSISTANCE IS
[IF (expression_string)
va_clause
(ELSEIF (expression_string)
va_clause)
ELSE]
va_clause
[DEPENDENCY LIST ([attribute_list])
```

You can include multiple ELSEIF clauses.

**expression\_string**

*expression\_string* defines a Boolean condition. If it returns TRUE, the server executes the associated *va\_clause* to return values for the attribute. *expression\_string* can be an expression or a complete user-defined routine. It must be written in Docbasic and must return a Boolean value. Any string literals in the expression or routine must consist of ASCII characters. You must have Sysadmin or Superuser user privileges to provide a user-defined routine for *expression\_string*.

For expressions, the syntax is:

```
expression [LANGUAGE docbasic]
```

For user-defined routines, the syntax is:

```
routine_name
([routine_parameter {,routine_parameter}])
FROM OBJECT(object_id)
[LANGUAGE docbasic]
```

*routine\_parameter* is an attribute name.

The FROM clause must identify an object whose content file contains the source code for the routine.

**va\_clause**

The *va\_clause* provides the values for the attribute. It has two possible formats. One format provides a list of literal values. The second defines a query to return the values. Each format allows you to provide an estimate of the expected number of attribute values and indicate whether the values represent the complete list of allowed values for the attribute.

- To provide a list of literal values, use the syntax:

```
LIST (literal_list)
[VALUE ESTIMATE=number]
[IS [NOT] COMPLETE]
```

- To provide values from a query, use the syntax:

```
QRY 'query_string'
[QRY ATTR = attribute_name]
[ALLOW CACHING]
[VALUE ESTIMATE=number]
[IS [NOT] COMPLETE]
```

*query\_string* is a DQL query. You can use the special token, \$\$, in the query string. When the query is executed, a single dollar sign will appear in the query in place of the token.

The QRY ATTR clause defines which of the attributes in the selected values list to display. Typically, *query\_string* has only one attribute in its selected values list—the attribute for which you are providing the value assistance. However, some situations may require you to put more than one attribute in the selected values list. In such cases, the server assumes that the first attribute selected is the attribute for which

you are providing value assistance. If this is not the case, you must include the QRY ATTR clause to define which attribute is the attribute for which you are providing assistance.

The ALLOW CACHING clause permits clients to cache the query results.

### dependency clause

The DEPENDENCY LIST clause identifies the attributes on which *expression\_string* depends. The default is all attributes named in the *expression\_strings*.

### mapping\_table\_specification

A mapping table specification most commonly maps a list of descriptive character strings to a corresponding list of integers. For example, assume an object type has an integer attribute called country in which each value represents a different country. If you display that attribute in a dialog box, you want users to see a name for each country rather than an integer value. Using a mapping table specification, you can map the integer values to corresponding country names.

The syntax is:

```
MAPPING TABLE (map_element {,map_element} )
```

where *map\_element* is:

```
VALUE=value_string  
[DISPLAY=display_string]  
[COMMENT=description_string]
```

For example:

```
MAPPING TABLE (VALUE=4  
DISPLAY=Spain  
COMMENT='Added to list in first quarter 98')
```

The default for DISPLAY is the data value as a character string literal. The default for COMMENT is a NULL string.

### default\_specification

A default specification provides a default value for an attribute. When a user creates a new instance of the type, Content Server automatically sets the attribute to the default value if no other value is specified by the user.

The default value can be specified as:

- A literal value
- One of the following keywords: USER, NOW, TODAY, TOMORROW, YESTERDAY
- NULL (This option is not valid for repeating attributes.)

For a single-valued attribute, the syntax is:

```
DEFAULT[=]default_value
```

For a repeating attribute, the syntax is:

```
DEFAULT [=] (default_value {, default_value})
```

If you specify the default value with a keyword that indicates a date or time, the keyword is specified using the DATE function:

```
DATE (keyword)
```

For example, the following statement creates object type mytype with one single-valued date attribute and one repeating string attribute and defines a default value for each:

```
CREATE TYPE "mytype" ("single_attr" date (default=date(NOW)),  
"rep_attr" string(32) repeating (default=(USER)))  
WITH SUPERTYPE "dm_document"
```

The default value must be appropriate for the datatype of the attribute. For example, you cannot specify the keyword USER for a date attribute. You cannot specify NULL for a repeating attribute.

### **constraint\_specification**

A *constraint\_specification* defines constraints for an attribute. For example, you can define a check constraint to provide data validation for the attribute. Or you can make the attribute a unique key for the type. Constraints are not enforced by Content Server. If you define a constraint for an attribute, it must be enforced by your client application.

You can define five kinds of constraints in an attribute modifier list:

- Primary key
- Unique key
- Foreign key
- Not null
- Check

If you define a primary key, unique key, foreign key, or check constraint in an attribute modifier list, only the single attribute for which the constraint is defined can be part of the constraint. If the constraint must include multiple attributes, define it in the type modifier list.

For each constraint, you can define an error message to display if the constraint is violated. You can also indicate whether the constraint is enforced by the application or disabled. Enforcement information is stored in the data dictionary but not used by Content Server.

Any constraint you define is inherited by all the type's subtypes.

Refer to [Constraints, page 66](#) in *Content Server Fundamentals* for expanded information about the constraints and what they do.

#### **Primary key**

You can define a primary key constraint only on single-valued attributes. The attribute must also have a NOT NULL constraint. You can define only one primary key for a

type. (A type may have more than one primary key if it inherits a primary key from its supertype.)

To define a primary key constraint in an attribute modifier list, use:

```
PRIMARY KEY
[REPORT 'message_string' [ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

*message\_string* is a string literal that contains the error message you want displayed if the constraint is violated. You can include the following special tokens in the message to parameterize it:

- `$value(attribute_name)`

When the message is displayed, `$value(attribute_name)` is replaced with the value entered by the user in the field associated with the attribute. The attribute name must be enclosed in parentheses with no spaces before the opening parenthesis.

For example:

```
The security type must be folder" or cabinet and you
entered $value(security_val).
```

- `$$`

Use `$$` in a message string to display a single dollar sign in the message.

The `DISABLE|ENFORCE` clause lets you specify that the constraint is either disabled or enforced by the client application. Content Server does not enforce constraints.

### Unique key

The syntax for a unique key constraint specification in an attribute modifier list is:

```
UNIQUE KEY
[REPORT 'message_string' [ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

Refer to the preceding description of a primary key constraint for information about the error message and enforcement clauses.

### Foreign key

You must have Sysadmin or Superuser user privileges to create a foreign key. To define a foreign key constraint in an attribute modifier list, use:

```
FOREIGN KEY REFERENCES type_name [(attribute_name)]
[REPORT 'message_string' [ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

*type\_name* identifies the foreign object type that contains the referenced attribute.

*attribute\_name* identifies the referenced attribute in *type\_name*. The referenced attribute must have the same datatype as the attribute for which you are defining the constraint.

If *attribute\_name* is not included, the server assumes that the referenced attribute is `r_object_id`. In this case, the attribute for which you are defining the foreign key must have an ID datatype.

Refer to [Primary key, page 78](#) for information about the error message and enforcement clauses.

### Not null

A NOT NULL constraint means that the attribute can't have a NULL value. To define a NOT NULL constraint specification, use:

```
NOT NULL
[REPORT 'message_string' [ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

Refer to [Primary key, page 78](#) for information about the error message and enforcement clauses.

### Check

Check constraints are used most commonly for data validation. The constraint consists of an expression or routine that must return TRUE. If it returns FALSE, the attribute's value violates the constraint. To define a check constraint, use:

```
CHECK(expression_string)
```

*expression\_string* can be an expression or a complete user-defined routine. The expression or routine must be written in Docbasic and return a Boolean value. Any string literals in the expression or routine must consist of ASCII characters. To specify a routine, you must have Sysadmin or Superuser user privileges.

For expressions, the syntax is:

```
expression [LANGUAGE docbasic]
[REPORT 'message_string' [ON VIOLATION]]
```

For user-written routines, the syntax is:

```
routine_name
([routine_parameter {,routine_parameter}])
FROM OBJECT(object_id |
PATH folder_path[VERSION version_label])
[LANGUAGE docbasic]
[REPORT 'message_string' [ON VIOLATION]]
```

*routine\_parameter* is an attribute name.

The FROM clause must identify an object whose content file contains the source code for the routine. Identify the object using the object's ID or a folder path for an object. If you use a folder path, you can include a version label. If you don't include a version label, the label CURRENT is assumed.

Refer to [Primary key, page 78](#) for information about the error message clause.



## Type modifiers

Type modifiers define data dictionary information for the type. For example, you can set the type's default lifecycle or define constraints for the type. Type modifier set attributes in the data dictionary objects for the type. No attributes are set in the `dm_type` or `dm_type_info` objects for the type.

### update\_modifier

You can use the same update modifiers for types as for attributes (refer to [Attribute modifiers, page 74](#)). In addition, you can use the following to set the default lifecycle for a type:

```
SET DEFAULT BUSINESS POLICY[=]
chronicle_id [VERSION version_label] |
NULL |
NONE
```

*chronicle\_id* is the object ID of the root version of the lifecycle. The VERSION clause identifies which version of the lifecycle you want to use. The default is the CURRENT version.

If you set the default lifecycle to NULL, the type inherits the default lifecycle from its supertype.

If you set the default lifecycle to NONE, the type has no default lifecycle.

### mapping\_table\_specification

A mapping table specification is typically used to map a user-friendly character string value to an underlying value that may not be as readable or easy to understand. A mapping table can also be used to map localized or type-specific values to an underlying value.

For example, Desktop Client uses a mapping table defined at the type level to display a list of the display config objects appropriate for the object type. The actual names of the display config objects are mapped to more user-friendly strings. (Display config objects represent subsets of type attributes that have a common display definition.)

If you define a mapping table at the type level, the mappings apply to that object type and its subtypes.

### constraint\_specification

You can define four kinds of constraints in a type modifier list:

- Primary key
- Unique key
- Foreign key
- Check

For each constraint, you can define an error message to display if the constraint is violated. You can also indicate whether the constraint is enforced by the application or

disabled. Enforcement information is stored in the data dictionary but not used by the Content Server.

Constraints are defined in a type modifier list if the constraint includes multiple attributes. For example, if two attributes make up the primary key for a type, the primary key constraint must be defined in the type modifier list. If only a single attribute defines the constraint, the constraint is typically defined in the attribute modifier list for the attribute rather than the type modifier list.

Refer to [Constraints](#), page 66 in *Content Server Fundamentals* for more information about the constraints and their use.

### Primary key

You can define only one primary key for a type. (A type may have more than one primary key if it inherits a primary key from its supertype.) To define a primary key constraint in a type modifier list, use:

```
PRIMARY KEY (attribute_list)
[REPORT 'message_string '[ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

*attribute\_list* is a comma-separated list of the attributes that make up the key. The attributes must be single-valued attributes and each must have a NOT NULL constraint. Additionally, all the attributes must be defined for the same object type.

*message\_string* is a string literal that contains the error message you want displayed if the constraint is violated. You can include two special tokens in the message to parameterize it:

- `$value(attribute_name)`

When the message is displayed, the server replaces `$value(attribute_name)` with the value entered by the user in the field associated with the attribute. The attribute name must be enclosed in parentheses with no space before the opening parenthesis. For example:

```
The security type must be folder or cabinet and you
entered $value(security_val).
```

- `$$`

Use `$$` in a message string to display a single dollar sign in the message.

The `DISABLE|ENFORCE` clause lets you specify that the constraint is either disabled or enforced by the client application. Content Server does not enforce constraints.

### Unique key

To define a unique key constraint, use:

```
UNIQUE KEY (attribute_list)
[REPORT 'message_string '[ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

*attribute\_list* is a comma-separated list of the attributes that make up the key. The attributes can be either single-valued or repeating attributes, but all must be defined for the same object type.

Refer to the description of a primary key constraint on [Primary key, page 82](#) for information about the error message and enforcement clauses.

### Foreign key

You must have Sysadmin or Superuser user privileges to define a foreign key constraint. Use the following syntax:

```
FOREIGN KEY (attribute_list)
REFERENCES type_name [(attr_list)]
[REPORT 'message_string' [ON VIOLATION]]
[DISABLE|ENFORCE [BY] APPLICATION]
```

*attribute\_list* is a comma-separated list of the attributes that make up the key.

*type\_name* identifies the foreign object type that contains the referenced attributes and *attr\_list* identifies the attributes in the foreign type that are referenced by those in the foreign key. The attributes defined in *attribute\_list* and *attr\_list* must match in number and datatypes.

If an *attr\_list* is not included, the server assumes that the referenced attribute is *r\_object\_id*. In this case, only one attribute can be specified in *attribute\_list* and it must have an ID datatype.

Refer to the description of a primary key constraint on [Primary key, page 82](#) for information about the optional error message and enforcement clauses.

### Check

Check constraints are used most commonly for data validation. The constraint consists of an expression or routine that must return TRUE. If it returns FALSE, the attribute's value violates the constraint. Check constraints are defined at the type level when *expression\_string* references two or more attributes. To define a check constraint, use:

```
CHECK(expression_string)
```

*expression\_string* can be an expression or a complete user-defined routine. The expression or routine must be written in Docbasic and return a Boolean value. To specify a routine, you must have Sysadmin or Superuser user privileges.

For expressions, the syntax is:

```
expression [LANGUAGE docbasic]
[REPORT 'message_string' [ON VIOLATION]]
```

For user-written routines, the syntax is:

```
routine_name
([routine_parameter {,routine_parameter}])
FROM OBJECT (object_id |
PATH folder_path[VERSION version_label])
[LANGUAGE docbasic]
```

```
[REPORT 'message_string' [ON VIOLATION]]
```

*routine\_parameter* is an attribute name.

The FROM clause must identify an object whose content file contains the source code for the routine. Identify the object using the object's ID or a folder path for an object. If you use a folder path, you can include a version label. If you don't include a version label, the label CURRENT is assumed.

Refer to [Primary key, page 82](#) for information about the error message clause.

### **mapping\_table\_specification**

A mapping table specification is typically used to map a user-friendly character string value to an underlying value that may not be as readable or easy to understand. A mapping table can also be used to map localized or type-specific values to an underlying value.

For example, Desktop Client uses a mapping table defined at the type level to display a list of the display config objects appropriate for the object type. The actual names of the display config objects are mapped to more user-friendly strings. (Display config objects represent subsets of type attributes that have a common display definition.)

If you define a mapping table at the type level, the mappings apply to that object type and its subtypes.

### **component\_specification**

Component specifications identify which components can operate on objects of the type. The syntax is:

```
COMPONENTS (component_id_list)
```

*component\_id\_list* contains one or more entries with the following syntax:

```
component_classifier=object_id|NONE
```

*component\_classifier* is a character string that represents a qualified component (a `dm_qual_comp` object). It consists of the component's class name and an acronym for the build technology used to build the component. For example, an component classifier might be `Checkin.ACX` or `Checkin.HTML`.

*object\_id* is the object ID of the qualified component represented by the classifier. If you specify NONE instead of an object ID, the component is not available for the type.

## **Related statements**

[Alter Type, page 45](#)

[Drop Type, page 94](#)

## Examples

The following example creates a new subtype called `employee`, sets the label text for most attributes and constraints in the attribute modifier list and the type modifier list, and publishes the data dictionary information.

```
CREATE TYPE "employee"
(
  "emp_ssn" string(10)
  (PRIMARY KEY,
   CHECK ('Len(emp_ssn)=10'
    LANGUAGE docbasic)
  REPORT 'The SSN must have exactly 10 digits.'
  ENFORCE BY APPLICATION,
   SET "label_text"='Social Security Number'),
  "emp_first_name" string(32)
  (SET "label_text"='First Name',
   NOT NULL),
  "emp_middle_name" string(32)
  (SET "label_text"='Middle Name'),
  "emp_last_name" string(32)
  (SET "label_text"='Last Name',
   NOT NULL),
  "emp_disambiguator" integer,
  "emp_department" integer
  (FOREIGN KEY REFERENCES "department" ("department_id")
  REPORT 'The dept. code must be a valid code.',
   NOT NULL,
   SET "label_text"='Department Code'),
  "emp_job_title" string(64)
  (CHECK (validate_job_title ("emp_job_title")
  FROM '\Admin\Scripts'
  VERSION 'approved'
  LANGUAGE docbasic
  REPORT '$value is not a valid job title.')
  )
  WITH SUPERTYPE NULL
  UNIQUE ("emp_first_name","emp_middle_name",
   "emp_last_name","emp_disambiguator"),
  SET "label_text"='Employee Record'
  PUBLISH
```

The following example creates a subtype of `dm_document` and publishes its information in the data dictionary:

```
CREATE TYPE "legal"
("lawyer" CHAR(30),"case_number" INT,
 "defendants" CHAR(30) REPEATING)
WITH SUPERTYPE "dm_document" PUBLISH
```

The following example creates a user-defined type with no supertype:

```
CREATE TYPE "my_base_type"
("author" CHAR(30), "title" CHAR(145), "doc_id" ID)
WITH SUPERTYPE NULL
```

The following example creates a subtype of the user-defined type `my_base_type`:

```
CREATE TYPE "acctg" ("accounting" CHAR(30) REPEATING)
WITH SUPERTYPE "my_base_type"
```

# Delete

**Purpose** Removes rows from a registered table.

## Syntax

```
DELETE FROM table_name WHERE qualification
```

## Syntax description

**Table 2-9. DELETE syntax description**

Variable	Description
<i>table_name</i>	Identifies the registered table from which you are removing rows. Use the name of table in the underlying RDBMS.
<i>qualification</i>	Defines the conditions used to restrict the rows that are deleted. Refer to <a href="#">The WHERE clause, page 142</a> for a description of the valid forms of a qualification.

## Return value

The DELETE statement returns a collection whose result object has one attribute, `rows_deleted`, that contains the number of rows deleted.

## Permissions

To delete a row, the following conditions must be true:

- Your object-level permission for the `dm_registered` object in the repository that represents the RDBMS table must be at least Browse.
- Your table permission for the `dm_registered` object that represents the table must be `DM_TABLE_DELETE`.
- The user account under which Content Server is running must have the appropriate RDBMS permission to delete from the specified table. (The actual name of this permission will depend on your RDBMS.)

(For more information about security and object-level and table permissions, refer to [Chapter 4, Security Services](#), in the *Content Server Administrator's Guide*.)

## General notes

The DELETE statement deletes rows from a registered table. A registered table is a table in the underlying RDBMS that has been registered with Content Server. (Refer to [Register, page 108](#) for information about registering tables.) All rows for which the WHERE clause qualification evaluates to TRUE are deleted.

## Related statements

[Insert, page 105](#)

[Register, page 108](#)

[Select, page 115](#)

[Unregister, page 153](#)

[Update, page 155](#)

## Example

The following example deletes the rows in the registered table `authors_table` that contain the name of any author who is not also found in the authors attribute of a document:

```
DELETE FROM "authors_table"  
WHERE NOT EXISTS (SELECT * FROM "dm_document"  
                  WHERE ANY "authors" = authors_table.name)
```



# Delete...Object

**Purpose** Deletes objects from the repository.

## Syntax

```
DELETE [PUBLIC] type_name [(ALL)]
[correlation_variable] OBJECT[S]
[IN ASSEMBLY document_id [VERSION version_label]
[NODE component_id] [DESCEND]]
[SEARCH fulltext search condition]
[WHERE qualification]
```

## Syntax description

**Table 2-10. DELETE...OBJECT syntax description**

Variable	Description
<i>type_name</i>	Identifies the type of object to remove from the repository. Valid <i>type_names</i> are:  dm_assembly and its subtypes dm_user and its subtypes dm_group and its subtypes dm_sysobject and its subtypes user-defined types with NULL supertypes and their subtypes
<i>correlation_variable</i>	Defines a qualifier for the type name that is used to clarify attribute references.
IN ASSEMBLY clause	Restricts the statement to objects in a particular assembly.  <i>document_id</i> identifies the document with which the assembly is associated. Use a literal object ID:  ID('object_id')  <i>version_label</i> specifies a particular version of the document. Use a symbolic or an implicit version label. If you do not include a version label, the server uses the version identified by the <i>document_id</i> argument.  <i>component_id</i> specifies a particular component in the assembly. Including the NODE option restricts the

Variable	Description
	<p>statement to the specified component. Use a literal object ID to identify the component:</p> <pre>ID('object_id')</pre>
	<p>The DESCEND keyword directs the server to delete not only all directly contained node components, but also any indirectly contained components that meet the criteria. If you do not include this keyword, the server deletes only the directly contained components that meet the criteria in the statement.</p>
SEARCH clause	<p>Restricts the statement to objects that meet the SEARCH clause <i>fulltext search condition</i>. Refer to <a href="#">The SEARCH clause, page 138</a> for a detailed description of the SEARCH clause.</p>
WHERE clause	<p>Restricts the statement to objects that meet the <i>qualification</i>. The qualification is an expression that evaluates to TRUE or FALSE. Refer to <a href="#">The WHERE clause, page 142</a> for a description of the valid forms of a qualification.</p>

## Return value

The DELETE...OBJECT statement returns a collection whose result object has one attribute, `objects_deleted`, that contains the number of objects deleted.

## Permissions

You must have Delete permission on an object to delete the object.

## General notes

Deleting objects is subject to the following conditions:

- The object cannot belong to a frozen assembly or a frozen or immutable virtual document.
- If the `compound_integrity` attribute in the server's server config object is set to TRUE, the object cannot belong to any virtual document, regardless of the whether the document is immutable or not.

The *type\_name* argument specifies the type of object to delete. The server searches all objects of the specified type and any subtypes for objects that meet any additional criteria you define in the statement.

The keyword PUBLIC restricts the query to those objects with the *r\_is\_public* attribute set to TRUE. If the query contains a SEARCH clause, the full-text search is restricted to those documents for which ISPUBLIC is TRUE. When the server queries or searches only public objects, it uses only the setting of *r\_is\_public* for security checks.

The keyword ALL directs the server to consider all versions of each object. If you do not include ALL, the server only considers versions with the symbolic label CURRENT. You must enclose ALL in parentheses.

After the server finds all objects that meet the defined criteria, it deletes those for which you have Delete permission. If any of the objects that are otherwise eligible for deletion belong to a frozen assembly or an unchangeable virtual document, then the entire statement is rolled back and no objects are deleted.

The optional clauses, IN ASSEMBLY, SEARCH, and WHERE, restrict the statement's scope. The IN ASSEMBLY clause restricts the operation to a particular virtual document assembly or node (component) within the virtual document. The SEARCH clause restricts the operations to indexed objects that meet the fulltext search condition. The WHERE clause restricts the operations to objects that meet the specified qualification. The clauses are applied in the following order:

- SEARCH
- IN ASSEMBLY
- WHERE

You cannot use DELETE...OBJECT to destroy record objects. Use the Destroy API command instead.

## The IN ASSEMBLY clause

Including the IN ASSEMBLY clause generates an error if the *compound\_integrity* attribute in the server's server config object is set to TRUE. This attribute controls whether objects contained in a virtual document may be deleted from the repository.

If the document identified by *document\_id* does not have an associated assembly, the statement fails with an error.

## Related statements

[Change...Object](#), page 61

[Create...Object](#), page 67

[Update...Object](#), page 158

## Examples

This example deletes all old documents owned by Joe:

```
DELETE "dm_document" OBJECTS
WHERE "r_modify_date" < date('01/01/1970')
AND "owner_name" = 'joe'
```

The following statement deletes all documents that contain the word yeast but not the word rolls:

```
DELETE "dm_document" OBJECTS
SEARCH DOCUMENT CONTAINS 'yeast' AND NOT'rolls'
```

This statement deletes all documents that contain the words yeast and bread and those that contain the words cake and wedding:

```
DELETE "dm_document" OBJECTS
SEARCH DOCUMENT CONTAINS 'yeast' AND 'bread'
OR 'cake' AND 'wedding'
```

The following statement deletes all workflows that have either Janine or Jeremy as a supervisor:

```
DELETE "dm_workflow" OBJECTS
WHERE "supervisor_name" = 'janine' OR
"supervisor_name" = 'jeremy'
```

This example also deletes all workflows that have Janine or Jeremy as a supervisor:

```
DELETE "dm_workflow" OBJECTS
WHERE "supervisor_name" IN ('janine','jeremy')
```

# Drop Group

**Purpose** Removes a group from the repository.

## Syntax

```
DROP GROUP group_name
```

## Syntax description

Table 2-11. DROP GROUP syntax description

Variable	Description
<i>group_name</i>	Identifies the group to remove from the repository. You can specify the name as an identifier or as a character string literal.

## Permissions

You must be the owner of a group or have Superuser user privileges to drop a group.

## General notes

When you drop a group, Content Server also removes any registry objects that identify the group as the subject of an audit or event notification request.

## Related statements

[Alter Group, page 43](#)

[Create Group, page 65](#)

## Example

This example drops the group called `rons_baby_gift`:

```
DROP GROUP rons_baby_gift
```

# Drop Type

**Purpose** Removes a user-defined object type from the repository.

## Syntax

```
DROP TYPE type_name
```

## Syntax description

**Table 2-12. DROP TYPE syntax description**

Variable	Description
<i>type_name</i>	Identifies the object type to remove.

## Permissions

You must own the type or have Superuser privileges to drop a type.

## General notes

The type you identify must meet the following conditions:

- No objects of the type can exist in the repository.
- The type cannot have any subtypes.

Dropping a type also removes data dictionary information for the type and its attributes. The information is removed when one of the following occurs:

- The `publish_dd` method is executed for the entire repository.
- The Data Dictionary Publisher job runs.

**Note:** After you drop a type, you may not be able to create a type by the same name immediately. Allow time for the system to synchronize the type cache before attempting to recreate the type.

**Related statements**

[Alter Type, page 45](#)

[Create Type, page 71](#)

**Example**

```
DROP TYPE "my_base_type"
```

# Execute

**Purpose** Executes administration methods.

## Syntax

```
EXECUTE admin_method_name [[FOR] object_id]
[WITH argument = value {,argument = value}]
```

## Syntax description

**Table 2-13. EXECUTE syntax description**

Variable	Description
<i>admin_method_name</i>	Specifies which administration method you want to execute. <a href="#">Table 2-14, page 97</a> lists the methods.
<i>object_id</i>	Administration method names are not case sensitive. Identifies the object on which you want the function to operate. Use the object's object ID.
WITH clause	Defines one or more arguments and their values for the administration method.  Valid arguments differ for each method. Refer to <a href="#">Chapter 3, Administration Methods</a> , for a description of each administration method and their arguments.

## Return value

The EXECUTE statement returns a collection. The attributes of the query result object in the collection depend on which administration method was executed. Refer to the description of each method in [Chapter 3, Administration Methods](#), for details.



## Permissions

The privileges required depend on the administration method you are executing. Refer to the description of the individual method in [Chapter 3, Administration Methods](#), for information.

## General notes

The EXECUTE statement is the DQL equivalent of the API Apply method. You can use EXECUTE to invoke any of the administration methods that you can invoke using the Apply method except PING and WEBCACHE\_PUBLISH. These cannot be executed using the EXECUTE statement.

[Table 2-14, page 97](#) lists the administration methods that you can invoke with the EXECUTE statement and briefly describes the task that each performs.

**Table 2-14. Administration methods by category for the EXECUTE statement**

Category of Operation	Function	Description
Process Management	<a href="#">CHECK_SECURITY</a> , page 186	Checks a user or group's permissions level for one or more objects.
	<a href="#">GET_INBOX</a> , page 219	Returns items in user's Inbox.
	<a href="#">MARK_AS_ARCHIVED</a> , page 256	Sets the <code>i_is_archived</code> attribute of a <code>dm_audittrail</code> , <code>dm_audittrail_acl</code> , or <code>dm_audittrail_group</code> object to T.
	<a href="#">PURGE_AUDIT</a> , page 275	Deletes audit trail entries from the repository.
	<a href="#">RECOVER_AUTO_TASKS</a> , page 288	Recovers work items claimed by a workflow agent master session but not yet processed.
	<a href="#">ROLES_FOR_USER</a> , page 302	Returns the roles assigned to a user in a particular client domain.
Execute procedures	<a href="#">DO_METHOD</a> , page 197	Executes system-defined procedures such as <code>lpq</code> or <code>who</code> or user-defined procedures.

Category of Operation	Function	Description
Content storage management	<a href="#">HTTP_POST, page 228</a>	Directs the execution of a method to an application server.
	<a href="#">CAN_FETCH, page 177</a>	Determines whether content in a distributed storage area component can be fetched by the server.
	<a href="#">CHECK_RETENTION_EXPIRED, page 182</a>	Finds SysObjects in content-addressed storage that have an expired retention period or no retention period.
	<a href="#">CLEAN_LINKS, page 189</a>	Provides maintenance for linked store storage areas.  On Windows platforms, cleans up unneeded linkrecord objects and resets file storage object security.  On UNIX platforms, cleans up unneeded linkrecord objects, directories, and links associated with linked storage areas.
	<a href="#">DELETE_REPLICA, page 193</a>	Removes a replica from a distributed storage area.
	<a href="#">DESTROY_CONTENT, page 195</a>	Removes a content object and its associated file from the repository. (Do not use this for archiving; use PURGE_CONTENT instead.)
	<a href="#">GET_FILE_URL, page 217</a>	Returns the URL to a content file.
	<a href="#">GET_PATH, page 225</a>	Returns the path to a particular content file in a particular distributed storage area component.
	<a href="#">IMPORT_REPLICA, page 234</a>	Imports an external file as a replica of content already in the repository.

Category of Operation	Function	Description
	<a href="#">MIGRATE_CONTENT</a> , page 260	Moves content files from one storage area to another.
	<a href="#">PURGE_CONTENT</a> , page 283	Deletes a content file from a storage area. (Used as part of the archiving process.)
	<a href="#">PUSH_CONTENT_ATTRS</a> , page 285	Sets the content metadata in a content-addressed storage system for a document stored in that storage system.
	<a href="#">REGISTER_ASSET</a> , page 290	Queues a request for the creation of a thumbnail, proxies, and metadata for a rich media content file. The request is queued to Media Server.  This method is only available or useful if you have Documentum Media Transformation Services running.
	<a href="#">REPLICATE</a> , page 295	Copies content in one component of a distributed storage area to another area.
	<a href="#">RESTORE_CONTENT</a> , page 300	Moves a file or files from archived storage to the original storage location.
	<a href="#">SET_CONTENT_ATTRS</a> , page 307	Sets the <code>content_attr_name</code> and <code>content_attr_value</code> attributes in the content object associated with the content file.
	<a href="#">SET_STORAGE_STATE</a> , page 316	Sets the state of a storage area to off-line, on-line, or read-only.

Category of Operation	Function	Description
	<a href="#">SYNC_REPLICA_RECORDS</a> , page 321	Synchronizes the replica record objects with the current definition of a distributed storage area.  <b>Note:</b> This method is not available through Documentum Administrator.
	<a href="#">TRANSCODE_CONTENT</a> , page 325	Queues a request for a content transformation to Media Server.  This method is only available or useful if you have Documentum Media Transformation Services running.
Database Methods	<a href="#">DB_STATS</a> , page 191	Provides database operation statistics for a session.
	<a href="#">DROP_INDEX</a> , page 206	Drops an index.
	<a href="#">EXEC_SQL</a> , page 211	Executes SQL statements.
	<a href="#">EXPORT_TICKET_KEY</a> , page 213	Exports a repository's login ticket key.
	<a href="#">FINISH_INDEX_MOVES</a> , page 215	Completes an interrupted object type index move operation.
	<a href="#">IMPORT_TICKET_KEY</a> , page 236	Imports a login ticket to the repository.
	<a href="#">MAKE_INDEX</a> , page 253	Creates an object type index.
	<a href="#">MOVE_INDEX</a> , page 271	Moves an object type index from one tablespace to another.  <b>Note:</b> Not supported on DB2.
	<a href="#">REORGANIZE_TABLE</a> , page 292	Reorganizes a database table for query performance.
	<a href="#">RESET_TICKET_KEY</a> , page 298	Generates a login ticket key for the repository.

Category of Operation	Function	Description
Full-Text Methods	UPDATE_STATISTICS, page 329	Updates the statistics for a database table.
	ESTIMATE_SEARCH, page 208	Returns the number of results matching a particular SEARCH condition.
	MARK_FOR_RETRY, page 258	Finds all queue items representing objects that failed indexing and marks them as pending indexing.
Session Management	MODIFY_TRACE, page 269	Sets the tracing level for full-text querying operations.
	CHECK_CACHE_CONFIG, page 179	Requests a consistency check on a particular cache config object.
	GET_LAST_SQL, page 223	Returns the last SQL statement issued.
	GET_SESSION_DD_LOCALE, page 227	Returns the locale in use for the current session.
	LIST_AUTH_PLUGINS, page 238	Lists the authentication plugins loaded by Content Server.
	LIST_RESOURCES, page 240	Provides information about the server operating system environment.
	LIST_SESSIONS, page 244	Provides information about current, active sessions.
LIST_TARGETS, page 248	Lists the connection brokers defined as targets for the server.	
		The information is returned in a collection with one result object whose attributes list the connection brokers defined as targets for the server.

Category of Operation	Function	Description
	<a href="#">LOG_ON</a> , page 251 and <a href="#">LOG_OFF</a> , page 250	Turn server logging of information about RPC calls on or off.
	<a href="#">SET_APIDEADLOCK</a> , page 304	Sets a deadlock trigger on a particular API method or operation.
	<a href="#">SET_OPTIONS</a> , page 312	Turn various tracing options on or off.
	<a href="#">SHOW_SESSIONS</a> , page 319	Provides information about current, active sessions and a user-specified number of timed-out sessions.

Unlike the Apply method, the EXECUTE statement is not case sensitive. Also, you do not have to specify the datatype of the arguments for each function. The statement determines the datatype from the value you assign to the argument.

## Examples

Refer to the individual descriptions of the method in [Chapter 3, Administration Methods](#), for examples.

# Grant

**Purpose** Gives one or more user privileges to one or more users.

## Syntax

```
GRANT privilege {,privilege} TO users
```

## Syntax description

**Table 2-15. GRANT syntax description**

Variable	Description
<i>privilege</i>	<p>Identifies the privilege you want to grant. Valid privileges are</p> <p>SUPERUSER            SYSADMIN            CREATE TYPE            CREATE CABINET            CREATE GROUP            CONFIG AUDIT            PURGE AUDIT            VIEW AUDIT</p>
<i>users</i>	<p>Identifies the users to whom you want to a privilege or privileges. The user name must belong to an individual user. You can specify one or more users as a comma-separated list of user names or use a SELECT statement to identify the user names. (Refer to <a href="#">Examples, page 104</a> for the use of a SELECT statement.)</p> <p>The user name must be the value is found in the user_name attribute of the dm_user object associated with the user.</p>

## Permissions

To grant Superuser or Sysadmin user privileges, you must have Superuser privileges.

To grant Create Group, Create Type, or Create Cabinet user privileges, you must have Superuser or Sysadmin user privileges.

To grant Config Audit, Purge Audit, Or View Audit privileges, you must be the repository owner or a Superuser. Repository owners and Superusers cannot grant these privileges to themselves.

## General notes

Granting a privilege to a user who already has that particular privilege does not generate an error.

## Related statements

[Revoke, page 113](#)

## Examples

The following example grants the Create Type user privilege to the users donna and carol:

```
GRANT CREATE TYPE TO donna,carol
```

This example grants the Create Cabinet user privilege to all individual users (that is, to those user names that do not represent a group):

```
GRANT CREATE CABINET TO  
(SELECT "user_name" FROM "dm_user"  
 WHERE "r_is_group" = FALSE)
```

The final example grants two privileges to the users jim and mike:

```
GRANT SYSADMIN,SUPERUSER TO jim,mike
```



# Insert

**Purpose** Inserts a row into a registered table.

## Syntax

```
INSERT INTO table_name [(column_name {,column_name})]
VALUES (value {,value}) | dql_subselect
```

## Syntax description

**Table 2-16. INSERT syntax description**

Variable	Description
<i>table_name</i>	Identifies the registered table in which you want to insert new rows.
<i>column_name</i>	Identifies a column in the table to receive an assigned value.
<i>value</i>	Specifies the value assigned to a column. The number of values specified must equal the number of columns named or the number of columns in the table. Refer to <a href="#">Assigning values to columns, page 106</a> for more information.
<i>dql_subselect</i>	Specifies a SELECT statement. The returned values are inserted into the table.

## Return value

The INSERT statement returns a collection whose result object has one attribute, `rows_inserted`, that contains the number of rows inserted into the table.

## Permissions

To use the INSERT statement, the following conditions must be true:

- Your object-level permission for the `dm_registered` object representing the RDBMS table must be at least Browse.

- Your table permission for the `dm_registered` object representing the table must be `DM_TABLE_INSERT`.
- The user account under which Content Server is running must have the appropriate RDBMS permission to insert data into the specified table. (The actual name of this permission will depend on your RDBMS.)

For more information about security and object-level and table permissions, refer to [Chapter 4, Security Services](#), in the *Content Server Administrator's Guide*.

## Assigning values to columns

Which value is inserted in each column you specify is determined by position. That is, the first column you specify in the statement receives the first value. The second column receives the second value, and so forth. Consequently, if you are inserting a value in every column, it is not necessary to specify the columns; the server automatically inserts the values in each column in turn. However, if you omit the column names, the number of values must equal the number of columns in the table.

If you specify column names, you can insert values into a subset of the table's columns. Also, it is not necessary to specify the column names in the same order in which they appear in the table. Columns that are not specified in the `INSERT` statement receive default values. (Refer to the documentation for your underlying RDBMS for information about the default values assigned to columns in this situation. Different systems have different rules. For example, some database management systems only allow you to default nullable columns. In such systems, all non-nullable columns must be specified.)

## Defining values

There are two possible ways to define the values to assign to columns. You can use the `VALUES` clause or you can use a DQL subselect statement.

The `VALUES` clause has the syntax:

```
VALUES value {,value}
```

One value must be specified for each column named in the statement. Each value must be a literal value appropriate for the datatype of the column to which it is being assigned. For example, assume there is a registered table called `customer_accounts` and that you want to insert values into four of its columns: `customer_name`, `acct_number`, `open_date`, and `balance`. The `customer_name` and `acct_number` columns are character string datatypes. The `open_date` column is a date datatype, and the `balance` column is a floating point datatype. The following statement inserts values into these columns:

```
INSERT INTO "customer_accounts" ("customer_name","account_number","open_date,
balance") VALUES ('Henrietta Hornsby','01264',date('4/2/1993'),125.00)
```

---

The value Henrietta Hornsby is inserted into the `customer_name` column, the value 01264 is inserted into the `account_number` column, and so forth.

Similarly, when you use a subselect statement, the returned values must be equal in number to the number of columns named in the `INSERT` statement and the values must be appropriate for the column into which they will be inserted. (Refer to [Select, page 115](#) for the subselect statement's syntax.)

## Related statements

[Delete, page 87](#)

[Register, page 108](#)

[Select, page 115](#)

[Unregister, page 153](#)

[Update, page 155](#)

## Example

This example saves the object names and creation dates of the contents of the `flowers` folder into the `objects` table:

```
INSERT INTO "objects" ("o_name", "c_date")
SELECT "object_name", "r_creation_date"
FROM "dm_document"
WHERE FOLDER ('/public/subject/flowers')
```

# Register

**Purpose** Registers a table from the underlying RDBMS with the repository.

## Syntax

```
REGISTER TABLE [owner_name.]table_name
(column_def {,column_def})
[[WITH] KEY (column_list)]
[SYNONYM [FOR] 'table_identification']
```

## Syntax description

Table 2-17. REGISTER syntax description

Variable	Description
<i>owner_name</i>	<p>Identifies the table's owner.</p> <p>Use the owner's RDBMS user name. If the owner is a repository user, this value may be found in the user's <code>user_db_name</code> attribute.</p> <p>If the RDBMS is Oracle or DB2 and the owner is the DBA, you can use the alias <code>dm_dbo</code>.</p> <p>If the RDBMS is MS SQL Server or Sybase and the owner is the DBA, you can use either <code>dbo</code> or <code>dm_dbo</code> as an alias.</p> <p><i>owner_name</i> is optional if the current user is the table's owner. The default value is the current user.</p>

Variable	Description
<i>table_name</i>	<p>Identifies the RDBMS table to register with the repository.</p> <p>You can specify the table's actual name or a synonym. The name must consist of ASCII characters.</p> <p>If you specify the table's actual name, do not include the SYNONYM clause in the statement.</p> <p>For Oracle or DB2, if you specify a synonym, that synonym must be previously defined through the RDBMS. Including the SYNONYM clause in the REGISTER statement is optional.</p> <p>For MS SQL Server or Sybase, if you specify a synonym, you must include the SYNONYM clause in the statement.</p> <p>For MS SQL Server, you cannot register a remote table.</p> <p>(Refer to <a href="#">The SYNONYM clause, page 111</a> for details.)</p> <p>If you do not have Superuser user privileges, you must be the owner of the table.</p>
<i>column_def</i>	<p>Describes columns in the table. The format for a column definition is:</p> <pre>column_name datatype [(length)]</pre> <p>where <i>column_name</i> is the name of the column in the table and <i>datatype</i> is the DQL datatype that corresponds to the column's RDBMS datatype. The column name must consist of ASCII characters. Valid datatypes in a column definition are:</p> <p>float, double integer, int char, character, string date, time</p> <p>You must also specify a length for columns that have a character, char, or string datatype (for example, char(24)).</p>

Variable	Description
<i>column_list</i>	Identifies the columns in the table on which indexes have been built. Use a comma-separated list of column names.
<i>table_identification</i>	<p>The name of the table in the underlying RDBMS.</p> <p>You must include the SYNONYM clause if you run against MS SQL Server or Sybase and you specify a synonym as the table name in the statement.</p> <p>Including the SYNONYM clause is optional if you are running against Oracle or DB2 and specify a synonym as the table name in the statement.</p> <p>(Refer to <a href="#">The SYNONYM clause, page 111</a> for details.)</p>

## Return value

When issued through IDQL, the REGISTER statement returns the object ID of the dm\_registered object for the table. If you issue the statement through IAPI (using the Query method), the statement returns a collection whose result object has one attribute, called new\_object\_id, that contains the object ID of the dm\_registered object for the table.

## Permissions

To register a table, you must own the table or have Superuser user privileges. If folder security is enforced in the repository, you must also have at least Write permission on the System cabinet.

## General notes

When you execute the REGISTER statement, the server creates an object of type dm\_registered (a SysObject subtype) that represents the RDBMS table in the repository. This object is automatically linked to the system (/system) cabinet.

A dm\_registered object can be manipulated like any other SysObject or SysObject subtype with one exception: you cannot version a dm\_registered object.

You do not have to include all of the columns in the underlying RDBMS table in the statement. You can specify a subset of the underlying table's columns. The column definitions you provide need not match the column definitions for the underlying table.

When the server creates the `dm_registered` object for the table, it uses your column definitions.

The `REGISTER` statement automatically assigns a default ACL to the `dm_registered` object. (The default is determined by the value in the `default_acl` attribute of the server's server config object.)

The statement also sets default table permissions. The table permissions are set to `SELECT` for the owner. The group and world are not given any default table permissions. You can change these by setting the attributes directly. Note that table permissions for registered tables are not hierarchical. (For a complete description of registered table permits, refer to [Table permits, page 404](#), in the *Content Server Administrator's Guide*.)

Changes to the definition of an underlying table are not automatically reflected in its corresponding `dm_registered` object. If the underlying table is modified and you want the `dm_registered` object to match the underlying table definition, you must unregister the table and then register it with new column definitions.

## The `SYNONYM` clause

Use the `SYNONYM` clause to register RDBMS tables that have synonyms in the underlying tables. (A synonym for an RDBMS table must be created independently in the RDBMS before you register the table. The Register statement does not create a synonym.)

The `SYNONYM` clause records the actual name of the table that corresponds to the table's synonym. The name is stored in the `synonym_for` attribute of the table's `dm_registered` object.

When you run against Oracle or DB2, Content Server passes the synonym directly to the database server. The actual table name, as specified in the `SYNONYM` clause and recorded in the `synonym_for` attribute is purely informational. For example, suppose johndoe has a table called `myremotetable` in the remote Oracle database `londonremote`, and that this table has the synonym `remotel`. To register this table, he uses:

```
REGISTER TABLE johndoe."remotel" ("columnA" int)
SYNONYM FOR johndoe.myremotetable@londonremote
```

After he registers the table, he can use the synonym in DQL statements and it is passed directly to the database server. For example, issuing the following DQL:

```
SELECT * FROM johndoe."remotel"
```

generates the following SQL:

```
SELECT * FROM johndoe."remotel"
```

When you run against MS SQL Server or Sybase, Content Server substitutes the value you specified in the `SYNONYM` clause (recorded in the `synonym_for` attribute) for the table name in the `SELECT`, `INSERT`, `UPDATE` and `DELETE` statements. For example, suppose johndoe issues the following `REGISTER` statement:

```
REGISTER TABLE johndoe."remotel" ("columnA" int)
SYNONYM FOR londonserver.londonremote.johndoe.myremotetable
```

After he registers the table, he issues the following SELECT statement:

```
SELECT * FROM johndoe."remotel"
```

Content Server substitutes the actual table name for the synonym and the following SQL is generated:

```
SELECT * FROM londonserver.londonremote.johndoe.myremotetable
```

### Special note for Oracle platforms

If you create a database link between the database on which the Documentum repository is installed and another database, and then use the SYNONYM feature to obtain information from that second database, both databases must be in the same codepage.

### Related statements

[Delete, page 87](#)

[Insert, page 105](#)

[Unregister, page 153](#)

[Update, page 155](#)

### Example

The following statement registers the RDBMS table named departments:

```
REGISTER TABLE "departments"
("dept_name" CHAR(30), "dept_code" INT)
KEY ("dept_code")
```



# Revoke

**Purpose** Removes one or more user privileges from one or more users.

## Syntax

```
REVOKE privilege {,privilege} FROM users
```

## Syntax description

Table 2-18. REVOKE syntax description

Variable	Description
<i>privilege</i>	<p>Specifies the privilege to revoke. Valid privileges are:</p> <p>SUPERUSER            SYSADMIN            CREATE TYPE            CREATE CABINET            CREATE GROUP            CONFIG AUDIT            PURGE AUDIT            VIEW AUDIT</p>
<i>users</i>	<p>Specifies the users from whom you are revoking the specified privilege or privileges. The user name must belong to an individual user. You can specify one or more users as a comma-separated list of user names or use a SELECT statement to identify the user names. (Refer to <a href="#">Examples, page 114</a> for the use of a SELECT statement.)</p> <p>The user name must be the value is found in the user_name attribute of the dm_user object associated with the user.</p>

## Permissions

To revoke the Sysadmin or Superuser user privilege, you must have Superuser user privileges.

To revoke the Create Group, Create Type, or Create Cabinet user privilege, you must have either Superuser or Sysadmin user privileges.

To revoke Config Audit, Purge Audit, Or View Audit privileges, you must be the repository owner or a Supersuser. Repository owners and Superusers cannot revoke these privileges from themselves.

## General notes

The statement succeeds even if a user does not have the privilege being revoked.

## Related sStatements

[Grant, page 103](#)

## Examples

The following example revokes the Create Cabinet and Create Type privileges from the users john and howard:

```
REVOKE CREATE CABINET, CREATE TYPE FROM john, howard
```

The next example demonstrates the use of a subselect statement to identify the users. This example revokes the Superuser user privilege from all users except haroldp:

```
REVOKE SUPERUSER FROM  
(SELECT "user_name" FROM "dm_user"  
WHERE "user_privilege" >= 16 AND "user_name" != 'haroldp')
```

# Select

**Purpose** Retrieves information from the repository, the database, or both.

## Syntax

A SELECT statement may be either a standard SELECT or an FTDQL SELECT statement. The syntax for an FTDQL SELECT statement is a subset of the standard syntax. For a brief description of an FTDQL SELECT statement, refer to [FTDQL SELECT statements](#), page 120.

The syntax for a standard SELECT statement is:

```
SELECT [FOR base_permit_level] [ALL|DISTINCT] value [AS
name] {,value [AS name]}
FROM [PUBLIC] source_list
[IN DOCUMENT clause | IN ASSEMBLY clause]
[SEARCH [FIRST|LAST] fulltext_search_condition]
[IN FTINDEX index_name{,index_name}]
[WHERE qualification]
[GROUP BY value_list]
[HAVING qualification]
[UNION dql_subselect]
[ORDER BY value_list]
[ENABLE (hint_list)]
```

The syntax for an FTDQL SELECT statement is:

```
SELECT [FOR base_permit_level] [ALL|DISTINCT] value [AS
name] {,value [AS name]}
FROM [PUBLIC] source_list
[SEARCH fulltext_search_condition]
[IN FTINDEX index_name{,index_name}]
[WHERE qualification]
[ORDER BY SCORE]
[ENABLE (hint_list)]
```

## Syntax description

**Table 2-19. SELECT syntax description**

Variable	Description
<i>base_permit_level</i>	Defines a minimum permission level for the returned objects. If specified, the query returns only those objects for which the user has at least the specified permit level.

Variable	Description
<i>value</i>	<p>Valid values are: NONE, BROWSE, READ, NOTE, VERSION, WRITE, and DELETE.</p> <p>If unspecified, the default is BROWSE.</p> <p>There are no FTDQL constraints on this value.</p> <p>Identifies the information to retrieve. Valid values are:</p> <ul style="list-style-type: none"> <li>• Attribute and column names</li> <li>• Scalar and aggregate functions</li> <li>• MFILE_URL function</li> <li>• Arithmetic expressions</li> <li>• The keywords CONTAIN_ID, CONTENTID, DEPTH, ISCURRENT, ISPUBLIC, ISREPLICA, OBJTYPE, PARENT, SCORE, SUMMARY, SYSOBJ_ID, TEXT, THUMBNAIL_URL, USER</li> <li>• An asterisk (*)</li> </ul> <p>Multiple values can appear in any order. If the SELECT statement is a subselect or a subquery, you can only include one value.</p>
<i>AS name</i>	<p>There are constraints on the values in the selected values list for FTDQL queries. For details, refer to the Usage Notes sections describing each kind of selected value.</p> <p>Defines a name for the result object attribute that will contain the retrieved value. This argument is optional. If you omit it, the system provides a default name. (Refer to <a href="#">The AS clause, page 130</a> for information about the default name.)</p> <p><i>name</i> must consist of ASCII characters</p> <p>You may include the AS <i>name</i> option in an FTDQL SELECT statement.</p>

Variable	Description
<i>source_list</i>	<p>Identifies the object types and RDBMS tables to search. You can specify any combination of object types and RDBMS tables in a standard SELECT statement.</p> <p>Types are specified using the following syntax:</p> <pre>type_name [(ALL)   (DELETED)] [correlation_variable] [WITH passthrough_hint_list]</pre> <p>Tables are specified using the following syntax:</p> <pre>[owner_name.]table_name [correlation_variable] [WITH passthrough_hint_list]</pre> <p><i>passthrough_hint_list</i> is a list of hints for one or more databases. The hint list for an individual database has the following format:</p> <pre>db_keyword('hint'{'hint'})</pre> <p>Valid <i>db_keywords</i> are: ORACLE, SQL_SERVER, SYBASE, and DB2. <i>hint</i> is any valid hint accepted by the particular RDBMS.</p> <p>If you include hints for multiple databases, the hint lists for each must be separated by commas and the entire set (for all databases) enclosed in parentheses. <a href="#">Passthrough hints, page 365</a> describes using passthrough hints fully.</p> <p>On Oracle and DB2, a passthrough hint in the source list in a subquery is applied to the entire SELECT statement.</p> <p>The RDBMS table must be a registered table unless the user issuing the SELECT is a superuser. Only superusers can query unregistered RDBMS tables. Additionally, there are several constraints on specifying more than one type in the source list.</p> <p>For complete information about specifying type and table names refer to <a href="#">Source list, page 132</a>. For a description of the constraints imposed on the source list by an FTDQL query, refer to <a href="#">FTDQL constraints on the source list, page 134</a>.</p>

Variable	Description
IN DOCUMENT clause	<p>Identifies the target of the SELECT statement as a particular virtual document. This clause can assemble a virtual document or identify the components of a virtual document. If you include this clause, you cannot specify the IN ASSEMBLY clause. Refer to <a href="#">The IN DOCUMENT clause, page 134</a> for the syntax and usage of the IN DOCUMENT clause.</p> <p>This clause may not be included in an FTDQL SELECT statement.</p>
IN ASSEMBLY clause	<p>Identifies the target of the SELECT statement as an assembly of a virtual document. The server uses the components of the assembly as the definition of the virtual document and applies any other specified conditions to those components, rather than searching the document's hierarchy for components.</p> <p>This clause may not be included in an FTDQL SELECT statement.</p>
<i>fulltext_search_condition</i>	<p>Defines criteria for searching of the full-text index. Refer to <a href="#">The SEARCH clause, page 138</a> for a description of its syntax and use.</p> <p>You can include a fulltext search condition in an FTDQL SELECT statement.</p>
<i>index_name</i>	<p>Identifies the index to be searched. Use the index name as defined in the associated fulltext index object. Currently, only one index for a repository is supported.</p> <p>You can include the IN FTINDEX clause only if the SELECT statement includes a SEARCH clause.</p>
<i>qualification</i>	<p>Restricts returned results to objects meeting the conditions in the qualification. Refer to <a href="#">The WHERE clause, page 142</a> for a full description of the valid forms of a qualification for a WHERE clause or a HAVING clause.</p>

Variable	Description
<i>dql_subselect</i>	<p>Defines an additional DQL SELECT statement. The columns returned by the statement must correspond to those returned by the outermost SELECT statement in number and datatype.</p> <p>A <i>dql_subselect</i> may not be included in an FTDQL SELECT statement.</p>
<i>value_list</i>	<p>Defines an order in which to group or sort the returned results. The values in this list must also be selected values. Refer to <a href="#">The GROUP BY clause, page 146</a> for the specific use of this in each clause.</p> <p>A GROUP BY clause may not be included in an FTDQL SELECT statement.</p>
<i>hint_list</i>	<p>One or more standard or passthrough DQL hints.</p> <p>Refer to the description of the source list for the format of a passthrough hint.</p> <p>Standard hints are the following:</p> <pre>SQL_DEF_RESULT_SET N FORCE_ORDER RETURN_TOP N OPTIMIZE_TOP N FETCH_ALL_RESULTS N OPTIMIZATION_LEVEL level_1 level_2 UNCOMMITTED_READ FTDQL NOFTDQL ROW_BASED</pre> <p>N is an integer value and level_1 and level_2 are optimization levels.</p> <p>For a brief description of the standard hints, refer to <a href="#">Table 2–20, page 149</a>. The ROW_BASED hint affects the selected values list and the WHERE clause qualification. For information about its effects, refer to the descriptions of those portions of the syntax. Information about ROW_BASED is also found in <a href="#">Appendix A, Using DQL Hints</a>, which contains full information about all the standard hints.</p>

## General notes

The SELECT statement retrieves information from object types and RDBMS tables. By default, the statement returns only objects for which the user has at least Browse permission. If you want to enforce a higher level of permission on the returned objects, you can include the FOR *base\_permit\_level* clause. For example, suppose you issue the following SELECT statement:

```
SELECT FOR VERSION "r_object_id","object_name","owner_name"  
FROM "dm_document" WHERE "subject"='budget_proposal'
```

The query returns all budget proposal documents for which the currently logged-in user has at least Version permission.

You can execute the statement as a stand alone statement and indirectly, as part of a variety of other DQL statements. For example, the following CREATE GROUP statement uses a SELECT statement to select the users that will populate the new group:

```
CREATE GROUP supers MEMBERS  
(SELECT "user_name" FROM "dm_users"  
WHERE "user_privilege" >= 16)
```

(For more information about CREATE GROUP, refer to [Create Group, page 65](#).)

## FTDQL SELECT statements

An FTDQL SELECT statement is a SELECT statement whose syntax conforms to a particular set of rules that allow the query to be executed directly against the fulltext index. If the statement conforms to FTDQL syntax, the statement is run solely against the fulltext index; the repository is not queried.

This feature provides enhanced performance for the query. The general syntax accepted for an FTDQL query is shown in the formal syntax description. The syntax constraints on particular clauses enforced for an FTDQL query are listed in the Usage Notes sections describing each clause. (A summary of the rules is found in [Table C-1, page 387](#).)

An FTDQL SELECT statement must be a standalone statement. That is, SELECT statements embedded as a subselect in another DQL statement are never executed as FTDQL SELECT statements. Nor are unioned SELECT statements executed as FTDQL SELECT statements. All FTDQL SELECT statements must include one of the following:

- A SEARCH clause
- The keyword SCORE or SUMMMARY in the list of selected values
- The DQL hint ENABLE(FTDQL)

A query that conforms to the FTDQL syntax rules is automatically executed as an FTDQL query. If you are wondering whether a particular query conforms to the rules, you can include the ENABLE(FTDQL) hint in the query. If the query conforms, the hint has no effect and the query is executed as an FTDQL query. If the query does not conform, Content Server returns an error.



If you do not want a particular query to execute as an FTDQL query even though it conforms to the FTDQL syntax, include the `ENABLE(NOFTDQL)` in the query.

**Note:** A query that does not conform to the FTDQL syntax and does not include the `ENABLE(FTDQL)` hint is executed as a standard `SELECT` query. It does not return an error.

FTDQL metadata queries that contain `LIKE` predicates with pattern matching or on metadata that contains underscores in the values may return different results than non-FTDQL queries on the same metadata. This occurs because the index server processes FTDQL queries against metadata and the database server processes non-FTDQL queries against metadata.

FTDQL queries are not affected by the `distinct_query_results` key in the `server.ini`. Setting this key to `T` (`TRUE`) does not affect how an FTDQL query is processed.

## The ALL and DISTINCT keywords

The optional `ALL` and `DISTINCT` keywords determine whether the `SELECT` statement returns duplicate rows. `ALL` returns all rows, including any duplicates. `DISTINCT` does not return duplicate rows. If neither keyword is included, the default is determined by the server's default behavior. (The server's default behavior is determined by the `distinct_query_results` flag in the server's `server.ini` start-up file.)

The `DISTINCT` keyword is ignored if the `SELECT` statement also includes an `IN DOCUMENT` or `IN ASSEMBLY` clause. The server issues a warning if the statement includes both the `DISTINCT` keyword and an `IN DOCUMENT` or `IN ASSEMBLY` clause. The warning is also issued if the server's default setting is not to return duplicates and the query contains an `IN DOCUMENT` or `IN ASSEMBLY` clause.

## Selecting attribute values

Select object attribute values by specifying the attribute's name as a selected value. The attributes you specify must belong to an object type identified in the `FROM` clause. The attributes can be either single-valued or repeating attributes. [Repeating attributes, page 122](#), contains guidelines for selecting repeating attribute values.

You cannot select any of the following attributes if you are querying audit trail entries unless you have `Superuser` or `View_audit` privileges:

- `acl_name`
- `acl_domain`
- `attribute_list`
- `object_name`
- `object_type`
- `owner_name`

- attribute\_list\_id
- chronicle\_id
- session\_id
- version\_label

The following statement returns the value of title, a single-valued attribute, for all documents in the repository:

```
SELECT "title" FROM "dm_document"
```

This next example returns the value of authors, a repeating attribute, from all documents that have bread as their subject:

```
SELECT "authors" FROM "dm_document"
WHERE "subject"='bread'
```

You can use an asterisk (\*) to select a predefined set of system-defined attributes for the object. Refer to [The Asterisk \(\\*\) as a selected value, page 130](#) for details.

### Repeating attributes

Including a repeating attribute in the selected values list is subject to the following constraints:

- You cannot select a repeating attribute and the name of a column from a registered table unless you also include the DQL hint, ROW\_BASED.
- You cannot select a repeating attribute if there are multiple object types identified in the FROM clause unless you also include the DQL hint, ROW\_BASED.

You can select both repeating and single-valued attributes in the same statement. If you do, the format of the returned results varies depending on how the query is written. The results can be returned with a separate result row for each value returned for a selected repeating attribute or the result rows can be returned in a master-detail format, with each row representing the values, including all repeating values, for one object.

#### To return results by object ID

To obtain query result objects that include all selected repeating attribute values for one object in one query result object, include the r\_object\_id attribute in the selected values list. Additionally, the source list in the FROM clause may contain only item.

To ensure that the results are ordered in either ascending or descending order, you can include an ORDER BY clause also. Typically, results are returned in ascending order by default.

**Note:** Refer to [The ORDER BY clause, page 148](#) for more information about using the ORDER BY clause.

For example, the following statement returns one result object that contains the names of all three authors in the authors attribute:

```
SELECT "r_object_id", "object_name", "authors"
FROM "dm_document"
WHERE "title" = 'Breads of France'
ORDER BY "r_object_id"
```

This statement returns the following object:

r_object_id	object_name	authors
<i>object_id</i>	French_breads	Jean Connie Corwin

### To return each repeating value in an individual result object

To obtain result objects that contain one repeating attribute value in each object, include the DQL hint, `ROW_BASED`, in the query.

For example, suppose you want to obtain the object ID and authors of a document whose `object_name` is French Bread, and you want the results in a separate row for each author. The following query that includes the `ROW_BASED` hint returns a separate row for each returned value for the authors attribute:

```
SELECT "r_object_id","authors" FROM dm_document
WHERE "title" = 'Breads of France'
ENABLE (ROW_BASED)
```

Assuming that there are three authors, named Jean, Connie, and Corwin, the returned rows are:

r_object_id	authors
0900000334fa21c3	Jean
0900000334fa21c3	Connie
0900000334fa21c3	Corwin

Additionally, if the query has a `WHERE` clause qualification that references a value in the selected repeating attribute, then the query will return only the row that contains that value. For example, suppose you execute the following statement:

```
SELECT "r_object_id","title","authors" FROM dm_document
WHERE "authors" = 'JPierpont'
ENABLE (ROW_BASED)
```

The query returns one row:

r_object_id	title	authors
0900000334fa21c3	English Muffins	JPierpont

### FTDQL requirements

The requirements for single or repeating attributes in a selected values list for an FTDQL query are the same as those for a standard query.

## Selecting column values

Select column values by specifying column names as values. The columns must belong to a RDBMS table that you identify in the FROM clause. You can use column names to select some or all of the column values, or you can use an asterisk to retrieve all column values.

**Note:** The account from which Content Server was installed must have SELECT privileges on the underlying table in the RDBMS before you can use DQL to select from an RDBMS table. Additionally, to query a registered table, you must have DM\_TABLE\_SELECT table permission and at least Browse object-level permission for the dm\_registered object representing the table. If the SELECT statement includes a SEARCH clause, your object-level permission must be at least Read.

### FTDQL requirements

The requirements for column values in a selected values list for an FTDQL query are the same as those for a standard query.

## Scalar and aggregate functions as selected values

You can include scalar and aggregate functions to manipulate returned values.

A scalar function operates on one value. The DQL SELECT statement accepts three scalar functions:

- UPPER, which returns the uppercase form of a character string
- LOWER, which returns the lowercase form of a character string
- SUBSTR, which returns a subset of a specified character string

An aggregate function operates on a set of values and returns one value. The DQL SELECT statement accepts five aggregate functions:

- COUNT, which returns the number of values in a group of values
- MIN, which returns the minimum value in a group of values
- MAX, which returns the maximum value in a group of values
- AVG, which returns the average value of a group of values
- SUM, which returns the total of all values in a group of values

For example, the following statement returns the number of documents owned by the user horace:

```
SELECT COUNT(*) FROM "dm_document"  
WHERE "owner_name"='horace'
```

You cannot include an aggregate function if your statement includes an IN DOCUMENT clause unless the statement also includes the USING ASSEMBLIES clause and an assembly exists for the virtual document itself. (Refer to [Snapshots](#), page 175 in *Content Server Fundamentals* for information about assemblies.) However, in this case, you may

prefer to use the IN ASSEMBLY clause instead of the IN DOCUMENT clause with USING ASSEMBLIES. Refer to [The IN ASSEMBLY clause, page 137](#) for more information.

### FTDQL requirements

You can include scalar and aggregate functions in the selected values list of an FTDQL query.

### MFILE\_URL function as a selected value

The MFILE\_URL function returns URLs for the content files and renditions associated with the objects returned by the SELECT statement. The function has three arguments that are used to control which URLs are returned. For a description of the function and its arguments, refer to [The MFILE\\_URL function, page 28](#).

You can include the MFILE\_URL function in the selected values list of an FTDQL query.

### Arithmetic expressions as selected values

You can include arithmetic expressions to perform calculations on the returned values. Arithmetic expressions are expressions that use arithmetic operators to form the expression, such as:

```
attribute A + attribute B
column C * column D
(2 * column B) - attribute F
attribute X + 4
```

The following statement returns the total of all rental charges paid in June, including regular rents and late charges:

```
SELECT "rent_received" + "late_charge_total" AS month_total
FROM "yearly_rent_records"
WHERE "mon" = 'june'
```

### FTDQL requirements

Arithmetic expressions are not allowed in the selected values list in an FTDQL query.

### Keywords as selected values

Keywords are words that have a special meaning for the server. The majority are full-text keywords. That is, the values they return are attribute values stored in the full-text index or values computed during a full-text search. Three keywords return information about an object's relationships to other objects in a virtual document. The remaining keywords determine whether an object is a replica, return the current user name, and return the URL to an object's thumbnail file.

It isn't necessary to include a `SEARCH` clause in the query to include a full-text key in the selected values list. However, for some keywords such as `TEXT`, the returned values are not useful unless a `SEARCH` clause is included.

### Full-Text keywords

The following keywords return information about full-text indexes.

- `CONTENTID`

The `CONTENTID` keyword returns the object ID of the content object representing the content file that matches the select criteria. (A content object links a content file to all documents that contain the file.)

- `ISCURRENT`

The `ISCURRENT` keyword is a Boolean keyword that returns 1 (`TRUE`) if the object is the current version and 0 (`FALSE`) if the object is not the current version.

- `ISPUBLIC`

The `ISPUBLIC` keyword is a Boolean keyword that returns 1 (`TRUE`) if the object is a public object (its `r_is_public` attribute is `TRUE`) and 0 (`FALSE`) if the object is not a public object.

- `OBJTYPE`

The `OBJTYPE` keyword returns the `r_object_type` of the selected object.

- `SCORE`

The `SCORE` keyword returns a document's relevance ranking as determined by the index server. If an `ORDER BY` clause is not included in the query, the returned document order is by descending score.

- `SYSOBJ_ID`

The `SYSOBJ_ID` keyword returns the object ID of the objects returned with a fulltext search. It is the object ID stored in the fulltext index for the object.

- `SUMMARY`

The `SUMMARY` keyword returns a summary of each document returned by the `SELECT` statement. The summary is up to 4096 bytes from the document, chosen by the Full-Text Engine. For example, the following statement returns a summary of the document whose name is Company History:

```
SELECT SUMMARY FROM "dm_document"  
WHERE "object_name"='Company History'
```

- `TEXT`

The `TEXT` keyword returns the variant text forms to which a specified term was match, to return a particular document. The values are returned as a repeating attribute. For example, if the search string specifies "talk", the index server may return documents that contain the words "talking" or "talked". If so, then `TEXT` would return those forms of the word "talk".

## FTDQL requirements

All the keywords classified as fulltext keywords can be included in an FTDQL query.

### Virtual document keywords

The following keywords return information about relationships in a virtual document.

- **CONTAIN\_ID**

The **CONTAIN\_ID** keyword returns the object IDs of the containment objects associated with the components of a virtual document. The **CONTAIN\_ID** keyword cannot return containment IDs for components that are found by searching an assembly. Consequently, if the **SELECT** statement includes an **IN DOCUMENT** clause with the **USING ASSEMBLIES** option, the containment IDs of any components found using an assembly will be zeros ('0000000000000000').

For example, suppose you have a virtual document, A, with one component, B. B is also a virtual document with four components and an assembly. If you execute a **SELECT** statement to retrieve the containment IDs of the components of A and include the **DESCEND** keyword, you receive the containment IDs for A's components at all levels. If the **SELECT** statement includes **DESCEND** and **USING ASSEMBLIES**, you receive only the containment ID for B. (B's containment object links it to A.) Because the direct components of B are found using an assembly, their containment IDs are not returned.


- **DEPTH**

The **DEPTH** keyword returns a component's level within a virtual document. You can only use the **DEPTH** keyword if the **SELECT** statement includes the **IN DOCUMENT** clause and the clause includes the **DESCEND** keyword.

To illustrate its use, suppose the following query is executed against the virtual document shown in [Figure 2-1, page 127](#) :

```
SELECT "r_object_id", DEPTH FROM "dm_document"
IN DOCUMENT ID('object_id_A') DESCEND
```

**Figure 2-1. Sample virtual document**



The statement returns

Object ID	DEPTH
A	0
B	1
D	2
C	1

- **PARENT**

The PARENT keyword returns the object ID of the object that directly contains a direct or indirect component of a virtual document. You can only include the PARENT keyword in the value list if the SELECT statement contains an IN DOCUMENT clause.

For example, suppose the following query is executed against the virtual document shown in [Figure 2-1, page 127](#):

```
SELECT "r_object_id", PARENT FROM "dm_document"
IN DOCUMENT ID('object_id_A') DESCEND
```

The statement returns the object IDs and parents of each component of the virtual document:

Object ID	PARENTS
A	A
B	A
D	B
C	A

Note that a virtual document is always considered to be its own parent.

### **FTDQL requirements**

You may not use any of the keywords classified as virtual document keywords in an FTDQL query.

### **Miscellaneous keywords**

The following keywords return miscellaneous information.

- **ISREPLICA**

The ISREPLICA keyword is a Boolean keyword that returns 1 (TRUE) if the returned object is a replica and 0 (FALSE) if the object is not a replica. (*Replica* is the term used to describe any object in the repository that has been placed in the repository as a copy of an object from another repository. Only environments using object replication have replicas.)



You can only include the ISREPLICA keyword in the value list of the outermost SELECT statement. You cannot use ISREPLICA in a subquery.

- THUMBNAIL\_URL

The THUMBNAIL\_URL keyword returns the URL to a thumbnail file. Use THUMBNAIL\_URL in a SELECT statement's selected values to return the URL to the thumbnail associated with the returned objects. For example:

```
SELECT object_name, THUMBNAIL_URL FROM dm_document
WHERE FOLDER('/Corporate Objectives')
```

The statement returns the documents in the folder Corporate Objectives. If the content of any of returned document has an associated thumbnail, the URL to that thumbnail is returned also. If a document has multiple content pages with thumbnails, the keyword returns only the thumbnail associated with the first content page that has a thumbnail.

The URL contains the following information:

- The base URL defined for the thumbnail storage area
- A path relative to the storage area identified in the store attribute that points to the thumbnail file
- A store attribute that identifies the storage area containing the thumbnail

If the storage area's require\_ticket attribute is TRUE, the URL also contains a ticket that contains an encryption of the path plus a time stamp. The ticket is valid for five minutes, which means that the URL is good only for five minutes.

For example:

```
http://myserver.documentum.com:8080/getThumbnail?
path=00232803/80/00/01/Ob.jpg&store=thumbnail_store_01&ticket=8002DWR670X
```

http://myserver.documentum.com:8080/getThumbnail? is the base URL.

path=00232803/80/00/01/Ob.jpg specifies the path to the file.

store=thumbnail\_store\_01 identifies the storage area.

ticket=8002DWR670X is the optional ticket.

- USER

The USER keyword returns the current user's user name.

### FTDQL requirements

The THUMBNAIL\_URL keyword is the only keyword in the miscellaneous category that is acceptable in an FTDQL query. You may not use the IS\_REPLICA or USER keywords in an FTDQL query.

## The Asterisk (\*) as a selected value

If the FROM clause references only registered tables or only object types, you can use an asterisk (\*) in the values list.

For registered tables, the asterisk returns all columns in the table.

For object types, what is returned by an asterisk depends on whether or not the ROW\_BASED hint is included in the query. If the hint is not included, an asterisk returns values only for the first object type listed in the FROM clause. The values returned depend on the object type:

- For objects that are subtypes of persistent object type, the asterisk returns:
  - All read/write single-valued attributes
  - The r\_object\_id attribute
- If the object type is SysObject or a SysObject subtype, in addition to the attributes returned for a persistent object, the asterisk also returns:
  - r\_object\_type
  - r\_creation\_date
  - r\_modify\_date
  - a\_content\_type

If the query includes the ROW\_BASED hint, the asterisk returns the values of all attributes of all object types in the FROM clause. For example, the following query returns all attribute values from both object types for returned objects:

```
SELECT * FROM dm_user, dm_group
```

If you want to limit the returned values to only the attribute values for a particular type in the FROM clause, qualify the asterisk with the type. For example:

```
SELECT a.*, b.r_object_id FROM dm_document a, mydoc b
WHERE a.subject = b.title ENABLE(ROW_BASED)
```

In the above example, the asterisk returns the values only for attributes from dm\_document objects. You can use an asterisk in this manner only when the ROW\_BASED hint is included in the query.

### FTDQL requirements

You may not use an asterisk as a selected value in an FTDQL query.

## The AS clause

The AS clause lets you name the attributes of the query result objects that represent the returned results. For example, suppose you issue the following SELECT statement:

```
SELECT "a"+"b" AS total FROM "accts"
```

Total is assigned as the name of the query result object attribute that contains the returned a+b values.

When the statement returns the value of an attribute or column name, providing a name for the query result attribute is optional. The system assigns the attribute or column name as the default name. If you provide a name, the name must consist of ASCII characters.

When the statement includes a function, arithmetic expression, or a keyword, you must specify a name if you are running against MS SQL Server or Sybase, and we recommend that you also do so if you are running against Oracle or DB2.

MS SQL Server and Sybase do not provide default names when values are functions, arithmetic expressions, or keywords. For example, suppose you issue the following statement:

```
SELECT COUNT(*) FROM "dm_document"
WHERE "subject" = 'cake'
```

Your return value would be:

\_\_\_\_\_

*integer*

Notice that there is no name above the integer return value.

For DB2, if you don't provide a name, the returned column is named with its position in the selected values list. For example, in the above SELECT example, the count would be returned in a column named 1.

Because the server does not allow duplicate attribute names in an object, if you assign the same name to more than one attribute, the system ignores all subsequent occurrences of the name and assigns those attributes different names.

Oracle provides a default name, but it is difficult to predict what forms the names for selected functions, arithmetic expressions, or keywords will have. Consequently, explicitly assigning a name is recommended.

### **FTDQL requirements**

You may use the AS clause in an FTDQL query.

## **The FROM clause**

The FROM clause defines which items are searched. Both the PUBLIC keyword and the source list restrict the search.

### **PUBLIC keyword**

The PUBLIC keyword restricts the search to objects for which the r\_is\_public attribute is set to TRUE. If the query contains a SEARCH clause, the full-text search is restricted to those documents for which ISPUBLIC is TRUE. When the server searches public objects only, it uses the setting of r\_is\_public for security checks.

You may include PUBLIC in an FTDQL query.

Including PUBLIC increases the search performance.

## Source list

A source list defines which object types and which RDBMS tables to search. The source list for a standard query can include:

- One or more object types, to a maximum of 10 types
- One or more RDBMS table names

The total number of object types and tables that you can include in the FROM clause is constrained by the RDBMS. Each relational database has a limit on the number of tables that can be joined in one query. In the underlying RDBMS query, for each object type in the FROM clause, all the *\_s* tables in the type's hierarchy plus any needed indexes are candidates for inclusion in the underlying query. If the query references repeating attributes, the *\_r* tables are included also.

For registered tables, all object types included in the view's definition are included in the query.

DQL passthrough hints included in the source list are applied only to the type or table after which the hint appears. For portability, you can include hints for multiple databases. If you include hints for multiple databases, only the hints that are applicable to the database receiving the request are used.

For more information about using passthrough hints, refer to [Passthrough hints, page 365](#).

### Including type names in the source list

Use the following syntax to specify a type name:

```
type_name [(ALL) | (DELETED)] [correlation_variable]
```

The *type name* argument identifies the object type to search. The server searches objects of the specified type and all subtypes. For example, to search all SysObjects, specify *dm\_sysobject*. The server searches *dm\_sysobject* and all of its subtypes, such as *dm\_document*, *dm\_process*, and so forth.

The keyword ALL directs the server to search all versions of each object. The keyword DELETED directs the server to search all versions of each object including any versions for which *i\_is\_deleted* is set to TRUE. (This attribute is set to TRUE if you delete the object and it is the root version of a version tree.) You must enclose ALL and DELETED in parentheses.

If the FROM clause includes neither ALL nor DELETED, the server searches only the CURRENT version.

The value of *correlation\_variable* is a qualifier for the type name that is used to clarify attribute references in the query.

Any user can specify any object type in a query with one exception. The exception is the type `dmi_audittrail_attrs`. To specify that type, you must have either superuser or View Audit privileges.

The server searches only objects within that specified type that are owned by the user or objects to which the user has access. Whether a user has access depends on the object-level permissions of the objects. The user must have at least Browse permission on an object for the object to be included in a query run by the user. If the query includes a `SEARCH` clause, the user must have at least Read permission on the object. (For a full description of server security, refer to [Chapter 4, Security Services](#), in the *Content Server Administrator's Guide*.)

Including multiple object types is subject to the following constraints:

- The selected values list cannot contain repeating attributes unless the `ROW_BASED` hint is included in the query.
- If the selected values list contains an asterisk (\*), only the first type's attributes will be expanded unless the `ROW_BASED` hint is included in the query.

[The Asterisk \(\\*\) as a selected value, page 130](#), contains details about including an asterisk in the selected values list.

- The query cannot include a `SEARCH` clause.
- Unqualified attribute names in the query are disambiguated from left to right, as the object types appear in the `FROM` clause.
- If the query includes an `IN DOCUMENT` clause, the object types must join in one-to-one relationship.
- If the query includes an `IN DOCUMENT` or `IN ASSEMBLY` clause or a `TYPE` or `FOLDER` predicate, the clause or predicate is applied to only one type in the join, and that type will be the first object type identified in the `FROM` clause.

### Including table names in the source list

You can include one or more RDBMS table names in the `FROM` clause of a standard query. The syntax is:

```
[owner_name.]table_name [correlation_variable]
```

The *owner\_name* argument identifies the owner of the table. You must include the owner name if you are not the owner of the table. If the owner is the DBA of the database, you can use the alias `dm_dbo` as the owner name. (Using `dm_dbo` when appropriate makes your statement portable across databases.) If the owner name is unspecified, the current user is assumed.

You must be a superuser to include an unregistered RDBMS table in the list. Any user with appropriate permissions can include a registered table.

*table\_name* is the name of the table that you want to search.

*correlation\_variable* is a qualifier for the table name used to clarify column references in the query.

### Clarifying type and table names

Type and table names can be the same. In such instances, the server must distinguish which name identifies a type and which identifies a table. The server uses the following rules for distinguishing between type and table names:

- If (ALL) or (DELETED) is present, the name is a type name.
- If *ownername* is present, the name is a table name.
- All other cases are ambiguous references. In such cases, the server checks to see whether the name is a type name. If it is, the name is assumed to be a type.

The values in the *source\_list* argument are processed from left to right. Consequently, if you include tables and types with the same name in the FROM clause, be sure to qualify the table name with its owner's name. This will ensure that the server does not mistakenly assume that the table name is the type name.

To illustrate, suppose you have an object type named `legal_docs` and a registered table of the same name. The following statement shows one correct way to include them both in one SELECT statement:

```
SELECT a."r_object_id", b."client"
FROM "legal_docs" a, jolenef."legal_docs" b
```

The first `legal_docs` (with correlation name `a`) is a type name. The server knows the second is a registered table because an owner name is specified. Note that correlation variables should be used in this case to clarify column references.

The following statement is valid because `billing_docs` is not found to be a type and so is assumed to be a registered table:

```
SELECT a."r_object_id", b."client"
FROM "dm_document" a, "billing_docs" b
```

In the following statement, the server assumes that the query references the `legal_docs` type, not the registered table, because no owner name is specified and there is a type called `legal_docs`:

```
SELECT a."object_id", b."client"
FROM "legal_docs" b, "dm_document" a
```

### FTDQL constraints on the source list

Observe the following rules if you want to execute the query as an FTDQL query:

- You can only reference one object type and that type must be `dm_sysobject` or a `dm_sysobject` subtype.
- You cannot reference a registered table in the source list.

### The IN DOCUMENT clause

**Note:** You may not include an IN DOCUMENT clause in an FTDQL SELECT statement.

The IN DOCUMENT clause lets you use the SELECT statement to perform the following operations:

- Identify only the directly contained components of a virtual document
- Identify the indirectly contained components that reside in the current repository with the virtual document

If a component is a folder that doesn't reside in the current repository, the query returns the object ID of the folder's mirror object in the repository, but not any folders or documents that the remote folder contains.

- Assemble a virtual document
- Identify the virtual documents that directly contain a particular object

If the value list for the SELECT statement contains a repeating attribute and the statement contains an IN DOCUMENT clause, the server automatically removes any duplicates of the repeating attribute.

You cannot include an ORDER BY clause when you include an IN DOCUMENT clause.

The syntax of the IN DOCUMENT clause is:

```
IN DOCUMENT object_id [VERSION version_label]
[DESCEND] [USING ASSEMBLIES]
[WITH binding condition]
[NODESORT BY attribute {,attribute} [ASC|DESC]]
```

### Specifying the virtual document

To identify which virtual document to search, use either the document's object ID or the combination of an object ID and a version label. When you use only the object ID, the server searches the version identified by that object ID. When you use a version label in addition to an object ID, the server searches that version of the object. Note that this version may not be the same as the version identified by the object ID.

If the SELECT statement is not a subquery, specify the *object\_id* using the special ID function and the actual object ID of the virtual document. For example:

```
IN DOCUMENT ID('object_id')...
```

If the SELECT statement is a subquery, you can also specify the *object\_id* using a correlated *r\_object\_id* attribute reference. For example:

```
SELECT "r_object_id" FROM "dm_document" x
WHERE EXISTS
(SELECT * FROM "dm_document"
IN DOCUMENT x."r_object_id"
WHERE "object_name"='Chapt1')
```

The VERSION option lets you specify a version label. This label can be assigned to any version in the version tree that contains the specified virtual document. If you do specify a label, the server searches the specified version rather than the version identified by the object ID.

### The DESCEND option

The keyword DESCEND directs the server to search the virtual document's hierarchy, including all components directly or indirectly contained in the virtual document.

Otherwise, the server searches only those components that are directly contained. You cannot use the DESCEND keyword in the IN DOCUMENT clause if the SELECT statement is a subquery.

### The USING ASSEMBLIES option

An assembly defines the components of a particular version of a virtual document at a particular time (the time the assembly was created). Creating an assembly ensures that a particular version of a virtual document always contains the same material. (Refer to [Snapshots, page 175](#) in *Content Server Fundamentals* for information about creating assemblies.)

If a component selected for inclusion is a virtual document and the USING ASSEMBLIES clause is included, the server determines whether an assembly is defined for that component. If so, the server uses the information in the assembly as the definition of the component's composition.

Refer to [Snapshots, page 175](#) in *Content Server Fundamentals* for a full discussion of assemblies.

### The WITH option

The WITH option lets you identify which version of a virtual document's component you want to include in the document. The syntax of the WITH option is:

```
WITH binding condition
```

where *binding condition* is one or more expressions that resolve to a single Boolean TRUE or FALSE. The expressions can include comparison operators, literals, attribute or column names, scalar and date functions, arithmetic expressions, and any predicates except SysObject predicates. Multiple expressions in the condition are joined using logical operators.

The condition defined by the WITH option is applied to any component of the virtual document that was added without a version being specified at the time of its addition. For example, suppose you add component C to virtual document A without specifying a particular version of component C. Later, when you assemble virtual document A, you can use the WITH clause condition to determine which version of component C to include in the document. You may want to include the version that has the word accounting as a keyword or the version that carries the symbolic label CURRENT. Each time you assemble the document, you can select a different version of component C. Choosing a version for inclusion at the time of actual assembly is called late binding. (For more information about late binding, refer to [Late binding, page 167](#) in *Content Server Fundamentals*.)

More than one version of a component may qualify for inclusion in the virtual document. To resolve this, you can include the NODESORT BY option in the IN DOCUMENT clause or you can allow the server to make a default choice. If you allow the server to make the choice, the server includes the version with the smallest object ID; that is, the earliest version of the component.



## The NODESORT BY option

The NODESORT BY option works in conjunction with the WITH option to identify one version of a particular component for inclusion in a virtual document. (In some instances, more than one version of a component may fulfill the condition imposed by the WITH option.) The NODESORT BY option sorts the versions by the value of one or more attributes. The server includes the first version.

The syntax of the NODESORT BY option is:

```
NODESORT BY attribute {,attribute} [ASC|DESC]
```

where *attribute* is any attribute belonging to the component. You can sort in ascending (ASC) or descending (DESC) order. If you do not specify a sort order, the default is ascending (ASC).

If you use the NODESORT BY option and two or more versions still qualify for inclusion, the server picks the version having the smallest object ID. This situation could occur if two or more versions have exactly the same values in the attributes specified in the NODESORT BY option.

## The IN ASSEMBLY clause

**Note:** You may not include the IN ASSEMBLY clause in an FTDQL SELECT statement.

The IN ASSEMBLY clause allows you to identify an assembly associated with a particular document. An assembly is a snapshot of a virtual document at a particular point in time and under conditions defined when the assembly was created. Using the IN ASSEMBLY clause means that the SELECT statement queries the virtual document represented by the assembly; however, the server uses only the components found in the assembly, rather than recursively searching the virtual document's hierarchy.

You can also use this clause to identify a single component of a virtual document and any components contained by that component. The component must be contained in the assembly.

The IN ASSEMBLY clause is more flexible than the IN DOCUMENT clause. For example, you can include aggregate functions such as COUNT or MAX in the value list when you use the IN ASSEMBLY clause. You cannot include aggregate functions when a SELECT statement includes an IN DOCUMENT clause.

The syntax of the IN ASSEMBLY clause is:

```
IN ASSEMBLY [FOR] document_id [VERSION version_label]  
[NODE component_id] [DESCEND]
```

where *document\_id* identifies the document with which the assembly is associated. Use the document's object ID for this argument, specified as an ID literal:

```
ID('object_id')
```

**Note:** You can create an assembly for a virtual document and assign that assembly to another document. In such cases, the value of the *document\_id* argument identifies the document to which the assembly is assigned, *not* the virtual document that the assembly represents.

If the specified document does not have an associated assembly, the statement fails with an error.

The VERSION option lets you specify a particular version of a document. If you include a version label, the server finds the version of the document carrying that label and uses the assembly associated with that version. You can specify either a symbolic label or an implicit label.

The NODE option allows you to identify a particular component of a virtual document. Use the component's object ID, specified as an ID literal:

```
ID('component_id')
```

The component must be contained in the assembly.

You cannot include the NODE option if the SELECT statement contains an aggregate function in the value list or if the SELECT statement is a subselect or subquery.

The DESCEND keyword directs the server to return not only directly contained components but also any indirectly contained components that meet the criteria in the statement. If you do not include this keyword, the server returns only directly contained components of the document or component.

## The SEARCH clause

The SEARCH clause identifies a subset of the full-text indexed documents. When you include a SEARCH clause, Content Server submits the query to the index server, which returns all objects that meet the SEARCH clause condition. The index server searches both content files and indexed metadata (attribute values) for matches with the specified condition.

The syntax of the SEARCH clause for a standard query is:

```
SEARCH [FIRST|LAST] fulltext_search_condition
[IN FTINDEX index_name {,index_name}]
```

The syntax of the SEARCH clause for an FTDQL query is:

```
SEARCH fulltext_search_condition
[IN FTINDEX index_name {,index_name}]
```

The FTDQL SEARCH clause may not contain either the FIRST or LAST keywords.

The *fulltext\_search\_condition* has one preferred syntax variant and one deprecated syntax variant. The preferred syntax is described in [SEARCH DOCUMENT CONTAINS, page 139](#). The deprecated syntax is briefly described in [SEARCH TOPIC, page 141](#).

## FIRST and LAST keywords

The FIRST and LAST keywords determine when the search is conducted. These keywords are only valid in standard SELECT queries. You may not include them if the query is an FTDQL query.

If you include FIRST, the server runs the full-text search query first and then applies the SELECT statement. For example, suppose legal\_documents is a subtype of dm\_document. The following statement first searches the full-text index for all documents that have the word fiduciary in their content and then finds the members of that group that are also legal\_documents:

```
SELECT "object_name" FROM "legal_documents"
SEARCH FIRST DOCUMENT CONTAINS 'fiduciary'
```

If you include LAST, the server first searches the repository for all objects that meet the SELECT criteria and then applies the criteria in the SEARCH clause. For example, the following statement first finds all legal\_documents authored by G. Oliphant and then finds the members of that group that contain the word fiduciary:

```
SELECT "object_name" FROM "legal_documents"
SEARCH LAST DOCUMENT CONTAINS 'fiduciary'
WHERE ANY "authors"='G.Oliphant'
```

The default behavior is to search the full-text index first. If you include a full-text keyword in the selected values list, the full-text search is conducted first even if LAST is specified in the SEARCH clause.

## SEARCH DOCUMENT CONTAINS

The SEARCH DOCUMENT CONTAINS syntax is the preferred way to specify a fulltext search condition in a SEARCH clause.

The syntax is:

```
SEARCH [FIRST|LAST]DOCUMENT CONTAINS 'search_string'
```

*search\_string* may be a list of words, phrases, or both, separated by spaces and quoted appropriately. For information about this syntax, refer to [Specifying words and phrases in SEARCH DOCUMENT CONTAINS](#), page 139.

The *search\_string* may be also specified using the following syntax:

```
[NOT] word {AND|OR [NOT] word}
```

The more complex syntax using logical operators is not recommended as future support may be dropped for this syntax.

### Specifying words and phrases in SEARCH DOCUMENT CONTAINS

Including a list of words or phrases in a SEARCH DOCUMENT CONTAINS clause returns those documents whose content contains the one or more of the specified words or phrases. The documents are returned in order of importance.

A *word* can be any character string that does not include spaces or punctuation. For example, the following statement finds all indexed documents that contain the word yeast:

```
SELECT * FROM "dm_documents"  
SEARCH DOCUMENT CONTAINS 'yeast'
```

To search for multiple words, separate each word in the list with a space. For example, the following query returns documents that have yeast or cake or both words in their content or metadata:

```
SELECT * FROM "dm_documents"  
SEARCH DOCUMENT CONTAINS 'yeast cake'
```

If you want to include a single quote in the search string, you must escape the quote with another single quote. For example:

```
SELECT * FROM "dm_documents"  
SEARCH DOCUMENT CONTAINS 'O' 'Malley'
```

To search on a phrase, enclose the phrase in double quotes inside the quoted search string. For example, the following query returns all documents that include the phrase “blend butter and sugar”:

```
SELECT * FROM "dm_documents"  
SEARCH DOCUMENT CONTAINS '"blend butter and sugar"'
```

**Note:** You can use single quotes to enclose a phrase, but if you do so, you must escape the quotes with single quotes.

You can combine words and phrases in the same search string. For example, the following query returns documents that contain both the word “yeast” and the phrase “blend butter and sugar”:

```
SELECT * FROM "dm_documents"  
SEARCH DOCUMENT CONTAINS 'yeast "blend butter and sugar"'
```

### Case sensitivity

The index server conducts searches using the following rules for case sensitivity:

- If a term or phrase is written in all lowercase characters, the search for that term or phrase is case insensitive.
- If a term or phrase is written in mixed case (some lowercase and some uppercase), the search is case sensitive.

For example, suppose you issue the following query:

```
SELECT owner_name,r_creation_date FROM dm_document  
SEARCH DOCUMENT CONTAINS 'budget'
```

That query is run as a case-insensitive query and returns all documents that contain, either in content or metadata, the word budget in lowercase or uppercase or in any combination (Budget, buDget, budgeT, and so forth). Now suppose you issue the following statement, specifying Budget:

```
SELECT owner_name,r_creation_date FROM dm_document
```

```
SEARCH DOCUMENT CONTAINS 'Budget'
```

That query is run as a case-sensitive query and returns only documents that contain the word Budget, capitalized as specified in the query.

### Accent and diacritical mark sensitivity

Some languages use accents and diacritical marks on some characters or syllables in words. Depending on how you specify the terms in the search string, the search may or may not be sensitive to such marks:

- If you do not include such a mark on a term in the search string, the index server will match the term with and without the marks.
- If you include such marks on a term in the search string, the index server will match the term exactly as it is specified in the search string.

For example, suppose you issue the following query:

```
SELECT owner_name,r_creation_date FROM dm_document
SEARCH DOCUMENT CONTAINS 'cote'
```

That query is run as insensitive to accent and diacritical marks. It returns all documents that contain, in metadata or content, the word cote, including those with instances of the word with accents or diacritical marks (côte, côté, and so forth).

Now, suppose you issue the following query that specifies a search term that includes an accent:

```
SELECT owner_name,r_creation_date FROM dm_document
SEARCH DOCUMENT CONTAINS 'coté'
```

That query is sensitive to accent and diacritical marks and returns only documents that contain the word “coté” exactly as specified in the query, with the accent mark.

### Asterisk as wildcard character

You can use the asterisk (\*) as a wildcard character in a search string. The asterisk matches any string of letters. You can use it in any position in the string. For example, foo\* matches any word that begins with “foo”. The string \*foo\* matches any word that contains “foo” in its sequence of letters. The string faa\*foo matches any word that begins with faa and ends with foo.

You may also use the asterisk as a wildcard in a phrase search. For example, suppose you provide the following search string:

```
'"n* is th* time*"'
```

That string matches phrases such as “now is the time”, “nothing is thoroughly timed”, and “November is thinking time”.

## SEARCH TOPIC

SEARCH TOPIC is a deprecated syntax. Although still supported for backwards compatibility, queries are not guaranteed to behave exactly as in previous releases because the VQL in the *topic\_query\_string* is translated by the new (5.3) index server.

Additionally, SEARCH TOPIC may not be supported in future releases. It is strongly suggested that new applications use the SEARCH DOCUMENT CONTAINS syntax.

The syntax of the clause is:

```
SEARCH [FIRST|LAST]TOPIC 'topic_query_string'
```

The query string is passed to the index server, which translates and executes the query.

### The IN FTINDEX option

The IN FTINDEX option is included for future use. It allows you to identify which index to search. However, Content Server only supports one index for a repository. If you include this clause, specify the name of the index associated with the repository. The index name is the name of its fulltext index object. If the index name begins with a digit, enclose the name in single quotes.

### The WHERE clause

Use the WHERE clause to restrict which objects are returned. For example, the following statement selects the title of every document but only returns the titles of documents owned by the current user:

```
SELECT "title" FROM "dm_document"
WHERE "user_name" = USER
```

The syntax of the WHERE clause is:

```
WHERE qualification
```

The *qualification* argument consists of one or more expressions that resolve to a single Boolean TRUE or FALSE. The expressions can include comparison operators, literals, attribute or column names, scalar and date functions, arithmetic expressions, predicates, and full-text keywords. Multiple expressions are joined together using the logical operators. ([Chapter 1, DQL Language Elements](#), contains descriptions of all accepted comparison and arithmetic operators, predicates, functions, and logical operators. Full-text keywords are described in [Full-Text keywords, page 126](#).)

A qualification can also include a subquery:

```
SELECT "r_object_id" FROM "dm_document"
WHERE ANY "authors" IN
(SELECT "user_name" FROM "dm_user"
WHERE "r_is_group" = TRUE)
```

The qualification can be simple or as complex as necessary. For example, the following WHERE clause is simple:

```
SELECT * FROM "dm_workflow"
WHERE "supervisor_name" = 'horace'
```

This next example contains a more complex qualification:

```
SELECT "r_object_id", "title," FROM "dm_document"
```

```
WHERE "r_creation_date" >= DATE('01/01/99', 'mm/dd/yy')
AND "r_modify_date" <= NOW
```

### General constraint on the qualification

A qualification cannot reference any of the following attributes of audit trail objects unless you have Superuser or View\_audit privileges:

- acl\_name
- acl\_domain
- attribute\_list
- attribute\_list\_id
- chronicle\_id
- object\_name
- object\_type
- owner\_name
- session\_id
- version\_label

### Using repeating attributes in qualifications

Referencing repeating attributes in a qualification is supported. However, there are some rules and guidelines. This section discusses those rules and guidelines.

#### ANY keyword use

If an expression references a repeating attribute, you must include the ANY keyword in the expression unless the ROW\_BASED hint is included in the query. For example, the following query includes the ANY keyword because it does not include ROW\_BASED:

```
SELECT "r_object_id", "title", "subject" FROM "dm_document"
WHERE ANY "authors" IN ('gillian')
```

If ROW\_BASED is included, the query may be written without the ANY predicate:

```
SELECT "r_object_id", "title", "subject" FROM "dm_document"
WHERE "authors" IN ('gillian') ENABLE(ROW_BASED)
```

If you are referencing a repeating attribute and including a subquery, you can use the IN or EXISTS keyword after the ANY keyword to control the generated SQL query. The syntax is:

```
WHERE ANY [IN|EXISTS] attr_name IN subquery
```

Including IN or EXISTS may change the generated SQL query and consequently, enhance performance. [IN and EXISTS, page 362](#) contains an example that shows the differences between the options in the generated query.

#### Referencing repeating attributes in compound expressions

A compound expression is multiple expressions linked by logical operators. For example, the following is a compound expression:

```
object_name = 'book_proposal' AND owner_name = 'john doe'
```

If two or more expressions in a compound expression reference repeating attributes, the attributes must be from the same object type unless the ROW\_BASED hint is included in

the query. For example, the following query contains a compound expression that has two expressions referencing repeating attributes from `dm_document`:

```
SELECT r_object_id, object_name FROM dm_document
WHERE ANY authors IN ('joe','john') AND ANY keywords LIKE 'eng%'
```

However, suppose you want to select from two object types and use a compound expression that referenced repeating attributes from both types. To do that, you must include the `ROW_BASED` hint:

```
SELECT a.r_object_id, b.i_acceptance_date
FROM dm_document a, dmi_package b
WHERE a.authors IN ('john doe') AND b.r_package_label = 'APPROVED' ENABLE(ROW_BASED)
```

Without the `ROW_BASED` hint, the query would fail with an error stating that the two repeating attributes it references, `authors` and `r_package_label`, are not from the same type.

Additionally, when `ROW_BASED` is included, it is not necessary to include the `ANY` predicate. Including `ANY` does not generate an error in such cases, but it is not necessary.

### Referencing repeating attributes in comparison expressions

A comparison expression is an expression that uses a comparison operator, such as `=` or `>`, to compare the values of two attributes. The following rules apply when comparing a repeating attribute to another attribute in a qualification:

- You can compare a repeating attribute to a single-valued attribute from the same or a different object type.

For example, the following queries are legal:

```
SELECT r_object_id, title FROM dm_document
WHERE ANY authors=object_owner
```

```
SELECT a.r_object_id, b.user_name
FROM dm_document a, dm_user b
WHERE ANY a.authors=b.user_name
```

- You cannot compare a repeating attribute to another repeating attribute unless you include the `ROW_BASED` hint.

For example, the following query is illegal:

```
SELECT a.r_object_id,b.r_object_id,a.title
FROM dm_document a,dm_sysobject b
WHERE ANY a.authors = b.keywords
```

To make that example a legal statement, you must include the `ROW_BASED` hint:

```
SELECT a.r_object_id,b.r_object_id,a.title
FROM dm_document a,dm_sysobject b
WHERE ANY a.authors = b.keywords
ENABLE(ROW_BASED)
```

### FTDQL requirements for where clauses

A `WHERE` clause in an FTDQL query is subject to the following rules:



- Any repeating attributes referenced in the clause must be of type string or ID.
- Only the DQL UPPER and LOWER functions are allowed. The functions SUBSTR and MFILE\_URL and all aggregate functions are not acceptable.
- The following predicates are not allowed:
  - BETWEEN
  - NOT LIKE
  - NOT FOLDER
  - [NOT] CABINET
  - ONLY
  - TYPE
- The IN and EXISTS keywords are not allowed.
- Any valid form of the FOLDER predicate (except NOT FOLDER) may be used. The DESCEND option is also allowed.
- The following rules apply to the LIKE predicate:
  - The LIKE predicate may be used with pattern matching characters.
  - The LIKE predicate may not be used with an escape clause.
- The keywords TODAY, YESTERDAY, TOMORROW may not be used in the DATE function.
- The following rules apply to all expressions in the WHERE clause:
  - Expressions may not contain the ISREPLICA or USER keywords.
  - Expressions may use any comparison operator.
  - Expressions may use an arithmetic operator, but they may not be used to form a compound expression.
  - Expressions that compare one attribute to another are not allowed. For example, subject=title is invalid.
  - Expressions in the following format are not allowed:  
*attribute\_name operator('literal' operator 'literal')*
  - Expressions that force index correspondence between repeating attributes are not allowed. Such expressions AND together expressions that reference repeating attributes in the format:  
*predicate(repeating\_attr\_expr AND repeating\_attr\_expr)*  
(For more information about forcing index correspondence, refer to [Forcing index correspondence in query results, page 346.](#))
- The following additional rules apply to expressions in the format *first\_expression operator second\_expression*

- The *first\_expression* is limited to one of the following:
  - attribute name*
  - `upper(attribute name)`
  - `lower(attribute name)`
- The *second\_expression* is limited to one of the following:
  - literal value*
  - `upper(literal value)`
  - `lower(literal value)`
  - the `DATE()` function
- The operator may be any valid operator.

## The GROUP BY clause

**Note:** You may not include a GROUP BY clause in an FTDQL SELECT statement.

The GROUP BY clause groups results. Use it when you include a function in the value list for the SELECT statement. The *value\_list* argument for the GROUP BY clause must contain all of the values in the value list for the SELECT statement that are not aggregate function values.

For example, the following statement returns the names of all documents, their owners, and a count of the documents that each owner owns. The results are grouped by owner name:

```
SELECT "owner_name", count(*)
FROM "dm_document"
GROUP BY "owner_name"
```

This next example selects the names of all documents, their owners, and their subjects, and groups them first by owner and then, within each owner group, by subject:

```
SELECT "owner_name", "subject", count (*)
FROM "dm_document"
GROUP BY "owner_name", "subject"
```

## The HAVING clause

**Note:** You may not include a HAVING clause in an FTDQL SELECT statement.

The HAVING clause restricts which groups are returned. It is most commonly used in conjunction with the GROUP BY clause.

For example, the following statement returns owner names and the number of documents each owner owns, grouped by owner names, for all owners having more than 100 documents:

```
SELECT "owner_name", count(*)
FROM "dm_document"
```

```
GROUP BY "owner_name"
HAVING count(*) > 100
```

This statement first finds and counts all documents for each owner name. It returns only those groups (owner's name and count) for which the count is greater than 100.

You can also use the HAVING clause in a SELECT statement that does not include the GROUP BY clause when a value in the value list is an aggregate function. For example, the following statement returns a count of the workflows supervised by haroldk if that count is greater than or equal to 25:

```
SELECT COUNT(*) FROM "dm_workflow"
WHERE "supervisor_name" = 'haroldk'
HAVING COUNT (*) >=25
```

If the number of workflows supervised by haroldk is less than 25, the statement returns nothing.

Note that if you do not include a GROUP BY clause when you use a HAVING clause, the only value permitted in the value list for the SELECT statement is one aggregate function.

## The UNION clause

**Note:** You may not include a UNION BY clause in an FTDQL SELECT statement.

The UNION clause lets you obtain results from more than one SELECT statement. The syntax is:

```
UNION dql_subselect
```

The *dql\_subselect* argument must return the same number of attributes or columns as the first SELECT statement, and each attribute or column must have the same datatype as its corresponding attribute or column in that SELECT statement. For example, if the first SELECT statement returns three attributes, then the subselect argument in the UNION clause must also return three attributes. The datatypes of the first attributes returned by each must be the same, as must the datatypes of the second and third attributes.

Neither the first SELECT statement nor any subsequent UNION SELECT statements can contain an IN DOCUMENT clause. The IN ASSEMBLY clause can be included only if the clause is in the first SELECT statement and all unioned subselect statements.

For all databases except DB2, when you union two or more SELECT statements, the attribute names that are returned are derived from the first SELECT statement. For example, suppose you issue the following statement:

```
SELECT "name", "address" FROM "current_emp"
UNION SELECT "ret_name", "ret_address" FROM "retired_emp"
```

The query result objects for this statement have two attributes, name and address, taken from the value list of the first SELECT in the statement.

For DB2, if corresponding selected values in the unioned SELECT statements don't have the same name or alias, the returned attribute name is an number representing the attribute's position in the selected values list. For example, suppose you issue the following query:

```
SELECT "name", "address" FROM "current_emp"  
UNION SELECT "ret_name", "ret_address" FROM "retired_emp"
```

The query result objects for this statement have two attributes. The first attribute is named 1, representing selected values from name and ret\_name. The second attribute is named 2, representing selected values from address and ret\_address.

The server does not return duplicate rows when you union two or more SELECT statements. This is true even if the ALL keyword is included in the first SELECT statement.

## The ORDER BY clause

Use the ORDER BY clause to sort the results of the SELECT statements. The syntax of this clause is:

```
ORDER BY value [ASC|DESC] {,value [ASC|DESC]}
```

The *value* argument must be an attribute or column name that appears in the value list. You can sort in ascending (ASC) or descending (DESC) order. If you do not specify a sort order, the default is ascending.

The primary sort is on the first value specified, the secondary sort is on the second value specified, and so forth. To illustrate, the following statement returns the owner name, subject, and title of all documents in the system in ascending order by owner name and the subject:

```
SELECT "owner_name", "subject", "title"  
FROM "dm_document"  
ORDER BY "owner_name", "subject"
```

You can specify a value either explicitly, by name, or by position. For example, the following statement returns the same results as the previous example, but notice that the ORDER BY clause specifies values by their position in the SELECT value list:

```
SELECT "owner_name", "subject", "title"  
FROM "dm_document"  
ORDER BY 1, 2
```

The only exception occurs when you union two or more SELECT statements. In such cases, you can only specify a value by its position in the value list.

### FTDQL requirements for ORDER BY

If you include an ORDER BY clause in an FTDQL query, the sort order must be specified by SCORE:

```
ORDER BY SCORE
```

You cannot reference SCORE by its position in the selected values list. You must reference it by name. You may sort in either ascending or descending order.

## Including DQL hints

You can include processing hints for DQL and the RDBMS servers in SELECT statements. The hints for DQL are called standard hints and can be added at the end of the SELECT statement, using the ENABLE keyword. Hints to the RDBMS server are called passthrough hints and can be added in the FROM clause, after a table or type name, and at the end of the statement, using the keyword ENABLE.

Table 2–20, page 149 lists the standard hints and provides brief guidelines for their use. For detailed information about their implementation, refer to [Appendix A, Using DQL Hints](#).

**Table 2-20. DQL standard hints**

Standard hint	Description
SQL_DEF_RESULT_SET <i>N</i>	<p>Defines a maximum number (<i>N</i>) of rows to return. Set <i>N</i> to a positive integer number. If set to 0, all rows are returned.</p> <p><b>Guidelines</b></p> <ul style="list-style-type: none"> <li>• On a SQL Server database, this hint forces the use of default result sets instead of a cursor to return the rows. On all other databases, the hint only sets the number of rows to return.</li> <li>• It is recommended that you set <i>N</i> to a maximum limit, rather than specifying it as 0.</li> <li>• If you include SQL_DEF_RESULT_SET and other hints that control the number of rows returned, the last one listed in the hints is used.</li> </ul>
FORCE_ORDER	<p>Controls the join order for the tables and types listed in the source list. If this hint is included, the tables are joined in the order they are listed.</p> <p><b>Guideline</b></p> <ul style="list-style-type: none"> <li>• This hint is ignored on DB2.</li> </ul>

Standard hint	Description
RETURN_TOP <i>N</i>	<p>Limits the number of results returned by a query. Set <i>N</i> to a positive integer number.</p> <p><b>Guidelines</b></p> <ul style="list-style-type: none"> <li>• If you include RETURN_TOP and other hints that control the number of rows returned, the last one listed in the hints is used.</li> <li>• This hint is useful because it can limit the effect of a bad (unbounded) query on the database.</li> <li>• Sorting the results or including the DISTINCT keyword reduces or eliminates the benefits of using RETURN_TOP.</li> <li>• On a SQL Server database, this hint reduces the number of rows touched by the query.</li> <li>• On a DB2 database, using OPTIMIZE_TOP with RETURN_TOP is recommended.</li> </ul>
OPTIMIZE_TOP <i>N</i>	<p>Returns <i>N</i> rows very quickly, then continues with the remainder of the rows. Set <i>N</i> to positive integer.</p> <p><b>Guidelines</b></p> <ul style="list-style-type: none"> <li>• This hint is most useful in conjunction with RETURN_TOP.</li> <li>• On an Oracle database, <i>N</i> is ignored.</li> <li>• This hint is ignored on a Sybase database.</li> </ul>
FETCH_ALL_RESULTS <i>N</i>	<p>Fetches all results from the cursor immediately and then closes the cursor. Set <i>N</i> to a positive integer number or 0 to return all rows.</p> <p><b>Guideline</b></p> <ul style="list-style-type: none"> <li>• If you include FETCH_ALL_RESULTS and other hints that control the number of rows returned, the last one listed in the hints is used.</li> </ul>

Standard hint	Description
OPTIMIZATION_LEVEL level_1 level_2	<p>Allows you to change the query optimization level. Set level_1 to the optimization level you want for the current query and level_2 to the optimization level desired for the connection after the query completes.</p> <p><b>Guideline</b></p> <ul style="list-style-type: none"> <li>This hint is only useful against DB2 databases. It is ignored for other databases. For information about valid values for level_1 and level_2, refer to your DB2 documentation.</li> </ul>
UNCOMMITTED_READ	<p>Ensures that a read only query returns quickly even if another session is holding locks on the tables queried by the read only query.</p> <p>This hint is useful only on SQL Server, DB2, and Sybase databases.</p>
ROW_BASED	<p>Directs Content Server to return query results that contain repeating attribute values in a row-based format—one row for each returned repeating attribute value. In addition, the hint also affects the rules governing:</p> <ul style="list-style-type: none"> <li>Selecting repeating attribute values</li> <li>Using an asterisk as a selected value</li> <li>Referencing repeating attributes in WHERE clause qualifications</li> </ul> <p>More information about these effects is found in <a href="#">Repeating attributes, page 122</a>, <a href="#">The Asterisk (*) as a selected value, page 130</a>, and <a href="#">Using repeating attributes in qualifications, page 143</a>.</p>
[NO]FTDQL	<p>FTDQL directs Content Server to execute the query as an FTDQL query. If the syntax of the query violates the rules for FTDQL syntax, an error is returned.</p> <p>NOFTDQL directs Content Server to execute the query as a standard query.</p>

## Examples

The following example returns the names and titles of all documents owned by the current user:

```
SELECT "object_name", "title" FROM "dm_document"  
WHERE "owner_name" = USER
```

The next example selects all documents and their authors with the subject `employee_benefits`:

```
SELECT "r_object_id", "authors" FROM dm_document  
WHERE "subject" = 'employee_benefits'  
ORDER BY 1
```

The following example returns the names of all objects in the New Books cabinet:

```
SELECT "object_name" FROM "dm_sysobject"  
WHERE CABINET ('/New Books')  
ORDER BY "object_name"
```

The following FTDQL example returns those documents that contain the phrase “for publication:” for which the user has Write permission:

```
SELECT FOR WRITE object_name,title,owner_name FROM dm_document  
SEARCH DOCUMENT CONTAINS 'for publication'
```

There are additional examples demonstrating the use of specific clauses in the descriptions of the clauses.



# Unregister

**Purpose** Removes a registered table from the repository.

## Syntax

```
UNREGISTER [TABLE] [owner_name.]table_name
```

## Syntax description

**Table 2-21. UNREGISTER syntax description**

Variable	Description
<i>owner_name</i>	<p>Identifies the table's owner. The default value is the current user.</p> <p>Use the owner's RDBMS user name. If the owner is a repository user, this value is found in the user's <code>user_db_name</code> attribute.</p> <p>If the RDBMS is Oracle and the owner is the DBA, you can use the alias <code>dm_dbo</code>.</p> <p>If the RDBMS is MS SQL Server or Sybase and if the owner is the DBA, you can use the alias <code>dbo</code>.</p>
<i>table_name</i>	<p>Identifies a registered table to remove from the repository. Use the table name as it appears in the RDBMS.</p>

## Permissions

You must be the owner of the registered table or have Superuser privileges to unregister a table.

## General notes

Unregistering a table removes the object of type `dm_registered` that represents the table in the repository. It does not remove the table from the underlying RDBMS.

If you attempt to unregister a table that is not registered, you receive an error message.

## Related statements

[Delete, page 87](#)

[Insert, page 105](#)

[Register, page 108](#)

## Example

The following example unregisters the table called departments:

```
UNREGISTER TABLE "departments"
```

# Update

**Purpose** Updates the rows of a registered table.

## Syntax

```
UPDATE table_name SET column_assignments
[WHERE qualification]
```

## Syntax description

Table 2-22. UPDATE syntax description

Variable	Description
<i>table_name</i>	Identifies the registered table to update. (Refer to <a href="#">Register, page 108</a> for information about registering tables in the repository.)
SET clause	Specifies the columns to update and the new values to assign to them. The syntax for <i>column_assignments</i> is:  <i>column_name</i> = <i>value expression</i>  Refer to <a href="#">Identifying columns to update, page 156</a> for a full description of the syntax.
WHERE clause	Restricts the rows that are updated to those that meet the criteria in the <i>qualification</i> . Refer to <a href="#">The WHERE clause, page 142</a> for a full description of WHERE clauses and qualifications.

## Return value

The UPDATE statement returns a collection whose result object has one attribute, `rows_updated`, that contains the number of updated rows.

## Permissions

To use the UPDATE statement, the following conditions must be true:

- Your object-level permission for the `dm_registered` object in the repository that represents the RDBMS table must be at least Browse
- Your table permission for the `dm_registered` object representing the table must be `DM_TABLE_UPDATE`
- The user account under which Content Server is running must have the appropriate RDBMS permission to update the specified table. (The actual name of this permission depends on your RDBMS.)

(For more information about security, object-level and table permissions, refer to [Chapter 11, Protecting Repository Objects](#), in the *Content Server Administrator's Guide*.)

## General notes

The SET clause identifies which columns are updated and the WHERE clause determines which rows are updated. The server searches for all rows that meet the qualification and updates the specified columns. If a WHERE clause is not included, then all rows are updated.

## Identifying columns to update

In the SET clause, *column\_assignment* has the format:

```
column_name = value_expression
```

where *column\_name* is the name of a column in the table.

The *value\_expression* can be a simple or complex expression but must resolve to a single value. It can include literal values, other column names, special keywords, date and scalar functions, and arithmetic expressions. The resulting value must be appropriate for the datatype of the specified column.

## Identifying rows to update

The WHERE clause *qualification* determines which rows are updated. The *qualification* consists of one or more expressions that resolve to a single Boolean TRUE or FALSE. The server updates only those rows for which the qualification is TRUE. The expressions in a qualification can include comparison operators, literals, scalar and date functions, column names, arithmetic expressions, and column predicates. Multiple expressions are joined together using logical operators.

## Related statements

[Delete](#), page 87

[Insert](#), page 105

[Register](#), page 108

[Unregister](#), page 153

## Example

This example updates the column called `chef_name`:

```
UPDATE "recipes" SET "chef_name" = 'carol'  
WHERE "chef_name" = 'carole'
```

# Update...Object

**Purpose** Updates an object in the repository.

## Syntax

```
UPDATE [PUBLIC] type_name [(ALL)][correlation_var]OBJECT[S] update_list
[,SETFILE filepath WITH CONTENT_FORMAT=format_name]
{,SETFILE filepath WITH PAGE_NO=page_number}
[IN ASSEMBLY document_id [VERSION version_label]
[NODE component_id] [DESCEND]]
[SEARCH fulltext search condition]
[WHERE qualification]
```

## Syntax description

**Table 2-23. UPDATE...OBJECT syntax description**

Variable	Description
<i>type_name</i>	Identifies the type of the object that you want to update. Valid values are:  dm_assembly and its subtypes dm_user and its subtypes dm_group and its subtypes dm_relation_type dm_sysobject and its subtypes  If you have Sysadmin or Superuser user privileges, you can also update the dmr_content object type.
<i>correlation_var</i>	Defines a qualifier for the type name that is used to clarify attribute references in the statement.

Variable	Description
<i>update_list</i>	<p>Specifies the update operations to perform on the object. Valid formats are:</p> <pre> set attribute_name = value set attribute_name[[index]] = value append [n] attribute_name = value insert attribute_name[[index]] = value remove attribute_name[[index]] truncate attribute_name[[index]] [un]link 'folder path' move [to] 'folder path' </pre> <p>If you specify more than one operation, use commas to separate them.</p>
SETFILE clause WITH CONTENT_FORMAT option	<p>Adds the first content file to the new object. Refer to <a href="#">WITH CONTENT_FORMAT option, page 161</a> for details.</p>
SETFILE clause WITH PAGE_NO option	<p>Adds additional content to the object. Refer to <a href="#">WITH PAGE_NO option, page 162</a> for details.</p>
IN ASSEMBLY clause	<p>Restricts the statement to objects in a particular assembly.</p> <p><i>document_id</i> identifies the document with which the assembly is associated. Use a literal object ID:</p> <pre>ID('object_id')</pre> <p><i>version_label</i> specifies a particular version of the document. Use a symbolic or an implicit version label. If you do not include a version label, the server uses the version identified by the <i>document_id</i> argument.</p> <p><i>component_id</i> specifies a particular component in the assembly. Including the NODE option restricts the statement to the specified component. Use a literal object ID to identify the component:</p> <pre>ID('object_id')</pre> <p>The DESCEND keyword directs the server to update not only all directly contained node components, but also any indirectly contained components that meet the criteria. If you do not include this keyword, the server updates only the directly contained components that meet the criteria in the statement.</p>

Variable	Description
SEARCH clause	Restricts the statement to objects that meet the SEARCH clause <i>fulltext search condition</i> . Refer to <a href="#">The SEARCH clause, page 138</a> for a detailed description of a SEARCH clause.
WHERE clause	Restricts the statement to objects that meet the <i>qualification</i> . The qualification is an expression that evaluates to TRUE or FALSE. It can include comparison operators, literals, attribute names, scalar and date functions, special keywords, predicates, and arithmetic expressions. Multiple expressions can be joined together using logical operators.  Refer to <a href="#">The WHERE clause, page 142</a> for a detailed description of the clause.

## Return value

The UPDATE...OBJECT statement returns a collection whose result object has one attribute, called `objects_updated`, that contains the number of objects updated.

## Permissions

To update an object, you must have Write permission on the object.

To include the SETFILE clause in the statement, you must have Superuser privileges.

## General notes

To update an object, the object cannot belong to a frozen assembly or a frozen or immutable virtual document.

The *type\_name* argument specifies the type of objects to update. The server searches all objects of the specified type and any subtypes for objects that meet any additional criteria you define in the statement.

The keyword PUBLIC restricts the query to those objects with the `r_is_public` attribute set to TRUE. If the query contains a SEARCH clause, the full-text search is restricted to those documents for which ISPUBLIC is TRUE. When the server queries or searches only public objects, it uses only the setting of `r_is_public` for security checks.



The keyword ALL directs the server to consider all versions of each object. If you do not include ALL, the server only considers the version with the symbolic label CURRENT. You must enclose ALL in parentheses.

The IN ASSEMBLY, SEARCH, and WHERE clauses restrict the scope of the statement. The IN ASSEMBLY clause restricts the operation to a particular virtual document assembly or node (component) within the virtual document. The SEARCH clause restricts the operations to indexed objects that meet the fulltext search condition. The WHERE clause restricts the operations to objects that meet the specified qualification. The clauses are applied in the following order:

- SEARCH
- IN ASSEMBLY
- WHERE

If you do not include any of these clauses, the updates are applied to all objects of that type for which you have Write permission.

If any of the objects that are otherwise eligible for updating belong to a frozen assembly or an unchangeable virtual document, then the entire statement is rolled back and no objects are updated.

## The SETFILE clauses

A SETFILE clause adds new content to the end of the sequence of content files in the object or replaces an existing content file. You cannot use SETFILE to insert a content file between two current files. You cannot store the content file you add in a turbo store storage area.

You must have Superuser privileges to include the clause in the statement.

Any object capable of having content may have multiple associated content files. All files must have the same content format and be ordered within the object.

The content format is defined when you add the first content file to the object. All subsequent additions must have the same format. Consequently, specifying the format for content additions after the first file is not necessary. Instead, you must specify the content's position in the ordered list of content files for the object.

To add the first content, use the SETFILE clause with the WITH CONTENT\_FORMAT option.

To add additional content, use the SETFILE clause with the PAGE\_NO option.

You can't include both options in a single SETFILE clause.

### WITH CONTENT\_FORMAT option

Use this SETFILE option to add the first content file to an object. The syntax is:

```
SETFILE 'filepath' WITH CONTENT_FORMAT='format_name'
```

The filepath must identify a location that is visible to Content Server.

The format name is the name found in the name attribute of the format's `dm_format` object.

### WITH PAGE\_NO option

Use this SETFILE option to add additional content to an object or replace existing content. The syntax is:

```
SETFILE 'filepath' WITH PAGE_NO=page_number
```

The filepath must identify a location that is visible to Content Server. The file must have the same file format as the existing content associated with the object.

The page number identifies the file's position of the content file within the ordered contents of the new object. You must add content files in sequence. For example, you cannot add two files and specify their page numbers as 1 and 3, skipping 2. Because the first content file has a page number of 0, page numbers for subsequent additions begin with 1 and increment by 1 with each addition.

To replace a content file, specify the page number of the file you want to replace.

## The update operations

The `update_list` defines the operations to perform. You can:

- Set the value of a single-valued or repeating attribute
- Add a value for a repeating attribute
- Insert a value into the list of values for a repeating attribute
- Remove a value from a repeating attribute
- Truncate a repeating attribute (remove all values)
- Link an object to a folder or cabinet or unlink an object from a folder or cabinet
- Move an object to a new folder or cabinet

Unless you have Superuser user privileges, you can only update read/write attributes. These attributes have names beginning with `a_` or have no prefix at all. If you have Superuser user privileges, you can update read-only attributes. These attributes have names beginning with `r_`.

You can specify more than one update operation in a single statement. When an UPDATE...OBJECT statement includes multiple operations, use commas to separate them. For example, the following statement sets two attributes and inserts a value into a repeating attribute. Notice the commas at the end of each update operation clause:

```
UPDATE "dm_document" OBJECTS
SET "title" = 'A Cake Primer',
SET "subject" = 'cake',
INSERT "authors"[3] = 'georgette'
WHERE "r_object_id" = '090007354140004e'
```

The server processes the update operations in the order listed.

### Setting an attribute value

To set the value of a single-valued or repeating attribute, use the following syntax:

```
set attribute_name[[index]] = value
```

The *attribute\_name* argument is the name of the attribute. If the attribute is a repeating attribute, the *index* identifies the position of the new value in the attribute's list of values. The positions of the values for a repeating attribute are numbered from zero. Enclose the *index* value in square brackets.

The *value* argument is the value to assign to the attribute. The value can be a literal or a subquery. If it is a subquery, it cannot return more than one row. If it returns more than one row, the statement returns an error.

### Appending a value to a repeating attribute

To append a value to a repeating attribute, use the following syntax:

```
append [n]attribute_name = value
```

The *attribute\_name* argument is the name of the attribute to which to append a new value. The *value* argument specifies what value to add to the attribute's ordered list of values. The value is automatically added to the end of the repeating attribute's list of values.

The value can be either a literal value or a subquery. The subquery can return any number of rows. By default, the system appends a maximum of 20 rows. To override the default, use the *[n]* option to define how many rows to append. You can use any integer number (to append that number of rows) or an asterisk (\*) (to append all rows).

### Inserting a value into a repeating attribute

To insert a value into the list of values for a repeating attribute, use the following syntax:

```
insert attribute_name[[index]] = value
```

The *attribute\_name* argument identifies the attribute. The *index* argument defines where to insert the new value. The positions of all values for a repeating attribute are numbered from zero. If you do not include the index, the server automatically inserts the new value in position zero, as the first element in the ordered list. You must enclose the index in square brackets.

The *value* defines the value that you want to insert. The *value* can be either a literal value or a subquery. If it is a subquery, it cannot return more than one row. If it returns multiple rows, the statement returns an error.

When you insert a value, all values that follow the inserted value are renumbered. For instance, when you insert a value at position [5], the value formerly at position [5] is moved up to position [6]. Consequently, if you are inserting more than one value in the attribute (for example, if the statement is in a program loop), be sure to increase the index number each time you insert a value.

## Removing values from repeating attributes

To remove a value from a repeating attribute, use the following syntax:

```
remove attribute_name[[index]]
```

The *attribute\_name* argument identifies the attribute. The *index* argument indicates the position of the value to remove in the attribute's ordered list of values. The positions of all values for a repeating attribute are numbered from zero. If you do not include the index, the server removes the first value (position 0) in the attribute. Enclose the index in square brackets.

When you remove a value, the remaining values are renumbered. For instance, when you remove the value at position [5], the value formerly at position [6] is moved up to position [5]. Consequently, if you are removing more than one value in the attribute (for example, if the statement is in a program loop), be sure to start with the highest index number and decrement the index number each time you delete a value.

For example, if you want to remove the values at positions 6 and 7, remove 7 first and then 6. If you remove 6 first, the value at 7 is moved into position 6 and the value at 8 is moved into position 7. When you remove 7, you are actually removing the value formerly in the position 8.

## Truncating a repeating attribute

To truncate a repeating attribute (remove its values), use the following syntax:

```
truncate attribute_name[[index]]
```

The *attribute\_name* argument identifies the attribute. The *index* argument specifies where in the attribute's list of values to begin the truncation. For example, if you specify `attribute_name[4]`, the server removes all values beginning with `attribute_name[4]`. If you do not specify an index level, the server removes all values.

## Linking or unlinking an object

To link an object to a folder or cabinet, use the following syntax in the `update_list`:

```
link 'folder path'
```

Linking an object adds a new link for the object. Current links are not affected.

To unlink an object from a cabinet or folder use the following syntax:

```
unlink 'folder path'
```

Unlinking removes only the specified link. Other links are not affected.

The *folder path* argument specifies the folder or cabinet to which you are linking the object. A folder path has the following format:

```
cabinet_name{/folder_name}
```

## Moving an object to a new folder or cabinet

To move an object to a new folder or cabinet, use the following syntax:

```
move [to] 'folder path'
```

The *folder path* argument specifies the folder or cabinet to which to link the object. A folder path has the following format:

```
cabinet_name{/folder_name}
```

Moving an object removes all current links and adds a link to the specified cabinet or folder.

## Related statements

[Change...Object, page 61](#)

[Create...Object, page 67](#)

[Delete...Object, page 89](#)

## Examples

The following statement deletes all indexed documents that contain the word *yeast*:

```
UPDATE "dm_document" OBJECTS
SET "keywords" = 'yeasted'
SEARCH DOCUMENT CONTAINS 'yeast'
```

This next statement updates all documents that contain the word *yeast* but not the word *rolls*:

```
UPDATE "dm_document" OBJECTS
SET "keywords" = 'pastries'
SEARCH DOCUMENT CONTAINS 'yeast' AND NOT 'rolls'
```

The following statement updates all cabinets that have either Janine or Jeremy as their owner:

```
UPDATE "dm_cabinet" OBJECTS
SET "is_private" = TRUE
WHERE "owner_name"='janine' OR owner_name='jeremy'
```



## Administration Methods

Administrative methods are methods that perform a variety of administrative and monitoring tasks. (Table 3–1, page 169 lists the methods by task categories.) They are executed in an application by invoking either the DQL EXECUTE statement or the API Apply method. You can also execute them interactively through Documentum Administrator.

This chapter first describes how to invoke the methods. Then it provides an alphabetical reference to the methods. For each method, it provides:

- Invoking syntax
- Arguments
- Return Value
- Permissions
- General Notes
- Related Administration Methods
- Examples

## Invoking administration methods

To execute an administration method manually, use Documentum Administrator.

To execute an administration method in an application, use either the API Apply method or the DQL EXECUTE statement.

You can also use Apply or EXECUTE to invoke an administration method through IAPI or IDQL if for any reason you can't access Documentum Administrator.

## Using Documentum Administrator

All the methods (except `DO_METHOD`, `SYNC_REPLICA_RECORDS`, `WEBCACHE_PUBLISH`, and `ROLES_FOR_USER`) are available through the Methods facility in repository Management. Many of the methods are also available through category-specific pages. For example, you can run `UPDATE_FTINDEX` to update a fulltext index through the Methods facility or through the Fulltext Index management facilities.

**Note:** Because `DO_METHOD` is used to execute user-defined procedures, this method is implemented as part of the Attributes page for user-defined methods.

## Using the EXECUTE statement

The EXECUTE statement is the DQL equivalent of the API Apply method. You can use it to execute any of the administration methods except `PING` or `WEBCACHE_PUBLISH`.

The EXECUTE statement is not case sensitive. You can enter the administration method name and argument names in uppercase, lowercase, or any combination.

You must enclose character string values (including an object ID in a FOR clause) in single quotes when they are included in the EXECUTE statement.

## Using the Apply method

You can use the Apply method to invoke any of the administration methods. The Apply method is case sensitive. The administration method name, argument names, and datatypes must be in uppercase.

## Scope of the administration methods

Administrative methods execute in the context of the current repository scope. You cannot connect to a repository and execute an administration method against a remote repository.



# Administration method operations

Table 3–1, page 169 lists the administration methods by category and describes the operation that you can perform with each method.

**Table 3-1. Administration methods by category**

Category of Operation	Administration Method	Description
Process Management	<a href="#">BATCH_PROMOTE</a> , page 175	Promotes multiple objects to their next lifecycle states.  <b>Note:</b> This method is not available through Documentum Administrator or using DQL EXECUTE.
	<a href="#">CHECK_SECURITY</a> , page 186	Checks a user or group's permissions level for one or more objects.
	<a href="#">GET_INBOX</a> , page 219	Returns items in user's Inbox.
	<a href="#">MARK_AS_ARCHIVED</a> , page 256	Sets the <code>i_is_archived</code> attribute of a <code>dm_audittrail</code> , <code>dm_audittrail_acl</code> , or <code>dm_audittrail_group</code> to T.
	<a href="#">PURGE_AUDIT</a> , page 275	Removes audit trail entries from the repository.
	<a href="#">RECOVER_AUTO_TASKS</a> , page 288	Recovers work items claimed by a workflow agent master session but not yet processed.
	<a href="#">ROLES_FOR_USER</a> , page 302	Returns the roles assigned to a user in a particular client domain.
Execute methods	<a href="#">DO_METHOD</a> , page 197	Executes system-defined procedures such as <code>lpq</code> or <code>who</code> or user-defined procedures.
	<a href="#">HTTP_POST</a> , page 228	Directs the execution of a method to an application server.

Category of Operation	Administration Method	Description
Content storage management	<a href="#">CAN_FETCH</a> , page 177	Determines whether content in a distributed storage area component can be fetched by the server.
	<a href="#">CHECK_RETENTION_EXPIRED</a> , page 182	Finds SysObjects in content-addressed storage that have an expired retention period or no retention period.
	<a href="#">CLEAN_LINKS</a> , page 189	Provides maintenance for linked store storage areas.  On UNIX, this method cleans up unneeded linkrecord objects and directories and links associated with linked storage areas.  On Windows, this method cleans up unneeded linkrecord objects and resets file storage object security.
	<a href="#">DELETE_REPLICA</a> , page 193	Removes a replica from a distributed storage area.
	<a href="#">DESTROY_CONTENT</a> , page 195	Removes a content object and its associated file from the repository. (Do not use this for archiving; use <code>PURGE_CONTENT</code> instead.)
	<a href="#">GET_FILE_URL</a> , page 217	Returns the URL to a content file.
	<a href="#">GET_PATH</a> , page 225	Returns the path to a particular content file in a particular distributed storage area component.
	<a href="#">IMPORT_REPLICA</a> , page 234	Imports an external file as a replica of content already in the repository.
	<a href="#">MIGRATE_CONTENT</a> , page 260	Moves content files from one storage area to another.

Category of Operation	Administration Method	Description
	<a href="#">PURGE_CONTENT</a> , page 283	Deletes a content file from a storage area. (Used as part of the archiving process.)
	<a href="#">PUSH_CONTENT_ATTRS</a> , page 285	Sets the content metadata in a content-addressed storage system for a document stored in that storage system.
	<a href="#">REGISTER_ASSET</a> , page 290	Queues a request for the creation of a thumbnail, proxies, and metadata for a rich media content file. The request is queued to the Media Server.
		This method is only available or useful if you have Documentum Media Transformation Services running.
	<a href="#">REPLICATE</a> , page 295	Copies content in one component of a distributed storage area to another area.
	<a href="#">RESTORE_CONTENT</a> , page 300	Moves a file or files from archived storage to the original storage location.
	<a href="#">SET_CONTENT_ATTRS</a> , page 307	Sets the content_attr_name and content_attr_value attributes in the content object associated with the content file.
	<a href="#">SET_STORAGE_STATE</a> , page 316	Sets the state of a storage area to off-line, on-line, or read-only.
	<a href="#">SYNC_REPLICA_RECORDS</a> , page 321	Synchronizes the replica record objects with the current definition of a distributed storage area.
		<b>Note:</b> This method is not available through Documentum Administrator.

Category of Operation	Administration Method	Description
	<a href="#">TRANSCODE_CONTENT</a> , page 325	Queues a request for a content transformation to the Media Server.  This method is only available or useful if you have Documentum Media Transformation Services running.
Database Methods	<a href="#">DB_STATS</a> , page 191	Provides database operation statistics for a session.
	<a href="#">DROP_INDEX</a> , page 206	Drops an index.
	<a href="#">EXEC_SQL</a> , page 211	Executes SQL statements.
	<a href="#">FINISH_INDEX_MOVES</a> , page 215	Completes an interrupted move operation for an object type index.
	<a href="#">MAKE_INDEX</a> , page 253	Creates an object type index.
	<a href="#">MOVE_INDEX</a> , page 271	Moves an object type index from one tablespace or segment to another.  <b>Note:</b> This is not supported on DB2.
	<a href="#">REORGANIZE_TABLE</a> , page 292	Reorganizes a database table for query performance.
	<a href="#">UPDATE_STATISTICS</a> , page 329	Updates the statistics in a database table.
Full-Text Methods	<a href="#">ESTIMATE_SEARCH</a> , page 208	Returns the number of results matching a particular SEARCH condition.
	<a href="#">MARK_FOR_RETRY</a> , page 258	Finds all content objects that have a particular negative update_count and marks them as awaiting indexing.
	<a href="#">MODIFY_TRACE</a> , page 269	Sets the tracing level for full-text indexing operations.

Category of Operation	Administration Method	Description
Session Management	<a href="#">CHECK_CACHE_CONFIG</a> , page 179	Requests a consistency check on a particular cache config object.
	<a href="#">GET_LAST_SQL</a> , page 223	Returns the last SQL statement issued.
	<a href="#">GET_SESSION_DD_LOCALE</a> , page 227	Returns the locale in use for the current session.
	<a href="#">LIST_AUTH_PLUGINS</a> , page 238	Lists the authentication plugins loaded by Content Server.
	<a href="#">LIST_RESOURCES</a> , page 240	Provides information about the server operating system environment.
	<a href="#">LIST_SESSIONS</a> , page 244	Provides information about current, active sessions.
	<a href="#">LIST_TARGETS</a> , page 248	Lists the connection brokers defined as targets for the server.  The information is returned in a collection with one result object whose attributes list the connection brokers defined as targets for the server.
	<a href="#">LOG_ON</a> , page 251 and <a href="#">LOG_OFF</a> , page 250	Turn server logging of information about RPC calls on or off.
	<a href="#">PING</a> , page 273	Determine if a client has an active server connection.
	<a href="#">SET_APIDEADLOCK</a> , page 304	Sets a deadlock trigger on a particular API method or operation.
<a href="#">SET_OPTIONS</a> , page 312	Turn various tracing options on or off.	

Category of Operation	Administration Method	Description
	<a href="#">SHOW_SESSIONS</a> , <a href="#">page 319</a>	Provides information about current, active sessions and a user-specified number of timed-out sessions.
Web Publishing Management	<a href="#">WEBCACHE_PUBLISH</a> , <a href="#">page 332</a>	Invokes the dm_webcache_publish method to publish documents to a Web site.

# BATCH\_PROMOTE

**Purpose** Promotes multiple objects to their next state.

## Syntax

```
dmAPIGet ("apply, session, NULL, BATCH_PROMOTE,
ARGUMENTS, S, 'object_ID{, object_ID}' ")
```

This method cannot be executed using the DQL EXECUTE statement.

## Arguments

Table 3-2. BATCH\_PROMOTE arguments

Argument	Datatype	Value	Description
ARGUMENTS	S	<i>comma-separated list of object IDs</i>	Identifies the objects to promote. Use the objects' object IDs. You must enclose the list of object IDs in single quotes.

## Return value

BATCH\_PROMOTE launches a method, dm\_bp\_batch. If the method executes successfully, BATCH\_PROMOTE returns a collection with one result object. If the method is not successfully launched, BATCH\_PROMOTE returns nothing. If BATCH\_PROMOTE returns nothing, use GetMessage to retrieve the error message.

If BATCH\_PROMOTE returns a collection, check the value of the method\_return\_val attribute of the result object. A value of 36 indicates that all objects were successfully promoted. Any value other than 36 means that an error occurred and processing stopped at some point. Use GetMessage to retrieve the error mesesage.

## Permissions

The user issuing the method must have the necessary permissions to promote the objects listed in the argument.

## General notes

BATCH\_PROMOTE allows you to promote multiple objects with one command. The objects can be attached to different lifecycles and can be in different states.

The method checks the permissions on each object specified to ensure that the user issuing the method has the correct permissions to promote the object. Then the method checks that associated lifecycle or lifecycles are valid. Finally, the method executes any entry criteria specified for the objects' next states. If BATCH\_PROMOTE encounters an error at this stage, the method exits and returns nothing.

If the permissions are correct, the lifecycles are valid, and any entry criteria are fulfilled, BATCH\_PROMOTE launches the dm\_bp\_batch method. The method performs any action procedures needed, as a single transaction. If an error occurs, the method exits and reports an error. Any objects promoted before the action procedure error remain promoted; the object whose procedure caused the error and any objects in the list after it are not promoted. For example, if the argument list includes 6 objects and an error occurs on an action procedure for object 4 in the list, objects 1, 2, and 3 are promoted. Objects 4, 5, and 6 are not.

There are two limitations on the use of BATCH\_PROMOTE:

- BATCH\_PROMOTE cannot run actions on behalf of the lifecycle\_owner. If the value in the a\_bpaction\_run\_as attribute in the docbase config object is set to lifecycle\_owner, BATCH\_PROMOTE exits with an error.
- You can specify a maximum of 200 objects in each execution of BATCH\_PROMOTE.

## Related administration methods

None



---

# CAN\_FETCH

**Purpose** Determines if the server can fetch a specified content file.

## Syntax

```
EXECUTE can_fetch FOR 'content_object_id'  
dmAPIGet("apply,session,content_obj_id,CAN_FETCH")
```

## Arguments

CAN\_FETCH has no arguments.

## Return value

CAN\_FETCH returns a collection with one query result object. The object has one Boolean attribute that is set to TRUE if the fetch is possible or FALSE if it is not.

## Permissions

Anyone can use this method.

## General notes

If a repository has a distributed storage area, it is possible to configure the servers so that they can fetch from all distributed storage area components, from a subset of the components, or only from the local component. (For information about configuring this capability, refer to the *Distributed Configuration Guide*.) In such repositories, you can use the CAN\_FETCH method to determine whether a server can fetch from a particular distributed storage area component.

In a content server configuration, the CAN\_FETCH method is executed by the content server.

## Related administration methods

[GET\\_PATH](#), page 225

## Examples

The following examples determine whether Content Server can fetch the content file associated with the content object 06000002472185e1:

```
EXECUTE can_fetch FOR '06000002472185e1'  
dmAPIGet ("apply, s0, 06000002472185e1, CAN_FETCH")
```

# CHECK\_CACHE\_CONFIG

**Purpose** Requests a consistency check on a particular cache config object.

## Syntax

```
EXECUTE check_cache_config [FOR 'cache_config_id']
[WITH argument = value ][,argument = value]
```

With EXECUTE, do not include the FOR clause if you include the CONFIG\_NAME argument.

```
dmAPIGet ("apply, session, cache_config_id, CHECK_CACHE_CONFIG
[, argument, datatype, value ][, argument, datatype, value] ")
```

With Apply, specify the cache config ID as NULL if you include the CONFIG\_NAME argument.

## Arguments

**Table 3-3. CHECK\_CACHE\_CONFIG arguments**

Argument	Datatype	Value	Description
CONFIG_NAME	S	<i>name of cache config object</i>	Identifies the cache config object on which to perform the consistency check. Use the cache config's object name.  This is an optional argument. If you do not include it, you must include the object ID of the cache config object.
FORCE_CHECK	B	T (TRUE) or F (FALSE)	TRUE directs Content Server to re-execute the queries regardless of the server_check_interval value. (Refer to the General Notes for a description of the use of the server_check_interval in the context of CHECK_CACHE_CONFIG.)

This is an optional argument.  
The default is F (FALSE).

## Return value

The method returns a collection with one query result object. The object has the attributes listed in [Table 3-4, page 180](#).

**Table 3-4. Attributes in the CHECK\_CACHE\_CONFIG result object**

Attribute	Datatype	Description
r_object_id	ID	Object ID of the cache config object
r_last_changed_date	Date/time	The date and time at which Content Server last validated the query results and found that something had changed.
r_last_checked_date	Date/time	The date and time at which Content Server last ran the queries to determine whether the results had changed.
server_check_interval	integer	How long the server waits between validations of the query results. The time is expressed in seconds.
client_check_interval	integer	The consistency check interval used by clients when issuing Fetch or Query_cmd methods that include this cache confi object as an argument.
server_time	Date/time	Current server time
server_time_secs	integer	Current server time in seconds
cache_config_exists	Boolean	F (FALSE) if the specified cache config object is not found. T (TRUE) otherwise.

If an error occurs, the method returns an error rather than the collection.

## Permissions

The cache config object must be owned by a superuser.

You must have at least Browse permission on the cache config object to issue this method. To issue the method with FORCE\_CHECK set to TRUE, you must be a Superuser or have Execute Procedure permission on the cache config object.

## General notes

CHECK\_CACHE\_CONFIG directs Content Server to check whether the data defined by a particular cache config object is current. (Cache config objects define queries whose results are cached persistently on the client.) To determine if the data is current, the server compares the current time to the date and time in the object's `r_last_checked_date` attribute. If the difference is less than the interval defined in `server_check_interval`, indicating the interval has not expired, the data is considered current and the server does not recompute the data. If the interval has expired, the data is considered out of date. The server executes the `dm_CheckCacheConfig` method to recompute the data. The method executes the queries defined in `cache_element_queries`, computes a hash of the results, and stores the hash in the `i_query_result_hash` attribute. The method returns the new computation date and time in the query result object's `r_last_checked_date` attribute.

You can use the `FORCE_CHECK` argument to override the server check interval and force Content Server to execute the queries.

For information about persistent client caching and cache config objects and their use, refer to [Persistent client caches, page 45](#) in *Content Server Fundamentals*. You can execute the `dm_CheckCacheConfig` method manually, using a `DO_METHOD` administration method. For an example of using `DO_METHOD` to call `dm_CheckCacheConfig`, refer to [Automating cache config data validation, page 197](#), in the *Content Server Administrator's Guide*. However, explicit calls to the method are not normally necessary. The calls occur automatically, as a side effect of referencing cache config objects in `Fetch` and `Query_cmd` methods.

## Related administration methods

None

## Examples

The following example identifies the cache config object by its object ID and forces the recomputation of the data:

```
EXECUTE check_cacke_config FOR '08000002723ae36b'
WITH force_check = true
```

This next example identifies the cache config object by its owner and name.

```
dmAPIGet ("apply,s0,NULL,CHECK_CACHE_CONFIG
,CONFIG_NAME,S,johndoe.report_app_config")
```

# CHECK\_RETENTION\_EXPIRED

**Purpose** Generates a list of objects whose content, stored in content-addressed storage, has an expired retention period or a zero retention period.

## Syntax

```
EXECUTE CHECK_RETENTION_EXPIRED
WITH QUERY='where_clause'
[,SELECT_LIST='attribute_list']
[,INCLUDE_ZERO_RETENTION_OBJECTS=T|F]

dmAPIGet ("apply,session,NULL,CHECK_RETENTION_EXPIRED,
QUERY,S,'where_clause'[,SELECT,S,'attribute_list']
[,INCLUDE_ZERO_RETENTION_OBJECTS,B,T|F]")
```

## Arguments

Table 3-5. CHECK\_RETENTION\_EXPIRED arguments

Argument	Datatype	Value	Description
QUERY	S	<i>where_clause</i>	Defines which SysObjects are to be checked for an expired or no retention period. Consists of a DQL SELECT where clause. Only literal values and attributes defined for the dm_sysobject type may be referenced. The string must be enclosed in single quotes.
SELECT_LIST	S	<i>attribute_list</i>	Identifies additional attributes whose values are to be returned in the result object.  Only attributes defined for the dm_sysobject object type may be included.  Separate multiple attribute names with commas and

Argument	Datatype	Value	Description
			enclose the entire list in single quotes.
INCLUDE_ZERO_RETENTION_OBJECTS	B	T (TRUE) or F (FALSE)	<p>If set to T, the method checks objects stored in content-addressed storage areas that allow but do not require a retention period.</p> <p>The default is F, meaning that only objects stored in content-addressed storage areas that require a retention period are checked.</p>

## Return value

CHECK\_RETENTION\_EXPIRED returns a collection of result objects, each representing one SysObject whose retention has expired or that had no retention period. The result objects have five default attributes plus any specified in the SELECT\_LIST argument. The default attributes of the result object are:

- r\_object\_id
- object\_name
- a\_storage\_type
- r\_creation\_date
- retention\_date

The retention\_date attribute is a computed attribute. The date value is the GMT equivalent of the retention period as it is defined in the time zone of the Centera storage system. For example, suppose the Centera system is in the Eastern time zone (EST) and the client on which user is working is in the Pacific time zone (PST). If the user sets the retention value to March 15, 2004 11 a.m. PST, the retention value is stored as March 15, 2004 2 p.m.—the Eastern time zone equivalent of March 15, 2004 11 a.m. The value returned in the computed attribute is the GMT equivalent of March 15, 2004 2 p.m.

## Permissions

You must have Sysadmin or Superuser privileges to execute this method.

## General notes

The CHECK\_RETENTION\_EXPIRED method is used by the RemoveExpiredRetnObjects administration job to find the content in content-addressed storage that has an expired retention period. The method returns a list of the objects. It does not remove the content from the storage area, nor does it remove from the repository any of the objects that contain the content.

A content-addressed storage area can have three possible retention period configurations:

- The storage area may require a retention period.  
In this case, the a\_retention\_attr name attribute is set and the a\_retention\_attr\_req is set to T.
- The storage area may not allow a retention period.  
In this case, the a\_retention\_attr name attribute is not set and the a\_retention\_attr\_req is set to F.
- The storage area may allow but not require a retention period.  
In this case, the a\_retention\_attr name attribute is set , but the a\_retention\_attr\_req is set to F.

By default, the method operates only on objects that are stored in content-addressed storage areas that require a retention period. The method never includes objects stored in content-addressed storage areas that do not allow retention periods. If you want it to examine objects stored in content-addressed storage areas that allow but do not require a retention period, you must set the INCLUDE\_ZERO\_RETENTION\_OBJECTS argument to true (refer to [INCLUDE\\_ZERO\\_RETENTION\\_OBJECTS argument](#), page 184, for more information).

### QUERY argument

The QUERY argument's where clause is a DQL where clause qualification. It is used to select objects for possible inclusion in the results returned by the method. The objects that fulfill the QUERY argument and are stored in an appropriate content-addressed storage area are then examined to determine whether the retention period has expired. If so, the object is included in the results.

The QUERY argument can reference only literal values or attributes defined for the dm\_sysobject object type. If the where clause includes single quotes, you must escape them with single quotes. For example:

```
...QUERY,S,'a_storage_type='castore_1'  
and r_creation_date > DATE(01/01/2003)'
```

### INCLUDE\_ZERO\_RETENTION\_OBJECTS argument

By default, the method does not include objects whose content has a 0 retention period because the assumption is that such content is meant to be kept forever. However, in a storage area that allows but does not require a retention period, a 0 retention period can be result from two possible causes:



- The user deliberately set no retention period, and consequently, the server set the retention period to 0
- The user specified a retention date that had already elapsed. When this occurs, the server sets the retention period to 0.

Because the meaning of 0 is ambiguous in such storage areas, the method supports the `INCLUDE_ZERO_RETENTION_OBJECTS` argument to allow you to include content with a zero retention in storage areas that allow but do not require a retention period.

If you set `INCLUDE_ZERO_RETENTION_OBJECTS` to T, when the method examines objects in storage areas that allow but do not require a retention period and it will include in the results any object with an expired or zero retention period.

### SELECT\_LIST argument

The `SELECT_LIST` argument allows you to include additional SysObject attributes in the result objects. You must enclose the argument's value in single quotes.

## Related administration methods

None

## Examples

The following example checks SysObjects stored in the content-addressed storage area named "castore\_2" owned by the user named "John Arthur". It also adds the title and subject attribute values to the result objects.

```
EXECUTE check_retention_expired
WITH QUERY='a_storage_type='castore_2' and
owner_name='John Author'',
SELECT_LIST='title,subject'

dmAPIGet ("apply,S0,NULL,CHECK_RETENTION_EXPIRED,
QUERY,S,'a_storage_type='castore_2' and
owner_name='John Author'',SELECT_LIST,S,'title,subject'")
```

This example directs the method to include objects in a storage area that allows but doesn't require a retention period:

```
EXECUTE check_retention_expired
WITH QUERY='a_storage_type='castore_3' and
owner_name='Mary Writer'',
SELECT_LIST='title,subject',
INCLUDE_ZERO_RETENTION_OBJECTS=true

dmAPIGet ("apply,S0,NULL,CHECK_RETENTION_EXPIRED,
QUERY,S,'a_storage_type='castore_3' and
owner_name='Mary Writer'',SELECT_LIST,S,'title,subject',
INCLUDE_ZERO_RETENTION_OBJECTS,B,T")
```

# CHECK\_SECURITY

**Purpose** Checks a user's or group's access permissions on one or more objects or checks a user's or group's permission level in one or more ACLs.

## Syntax

```
EXECUTE check_security WITH user_name='name'|group_name='name',
level=security_level,object_list='list_of_objectids'

dmAPIGet ("apply, session, NULL, CHECK_SECURITY,
USER_NAME|GROUP_NAME, S, name, LEVEL, I, security_level,
OBJECT_LIST, S, list_of_objectids
```

## Arguments

**Table 3-6. CHECK\_SECURITY arguments**

Argument	Datatype	Value	Description
USER_NAME	S	<i>name</i>	Name of the user for whom you are checking permissions. If you include this argument, do not include GROUP_NAME.
GROUP_NAME	S	<i>name</i>	Name of a group for which you are checking permissions. If you include this argument, do not include USER_NAME.
LEVEL	I	<i>security_level</i>	Minimum access permission level for which you are checking. Valid values are:  1, for None 2, for Browse 3, for Read 4, for Relate 5, for Version 6, for Write 7, for Delete
OBJECT_LIST	S	<i>list of object IDs</i>	The objects being checked. This can be a list of SysObject

Argument	Datatype	Value	Description
			object IDs, a list of ACL object IDs, or both. Use a space to delimit the individual items.

## Return value

CHECK\_SECURITY returns a collection of query result objects. The objects have one attribute, `r_object_id`. The attribute contains the object IDs of those objects submitted in the `OBJECT_LIST` argument for which the user or group has at least the permission level identified in the `LEVEL` argument.

## Permissions

You must have Superuser privileges to use this method.

## General notes

There are two uses for CHECK\_SECURITY:

- You can use it to determine whether a user or group has a particular access permission or better for a group of SysObjects
- You can use to determine whether a user or group has entries in one or more ACLs that give the user or group a particular access permission (or better).

To determine whether a user or group has at least the permission identified in the `LEVEL` argument for a particular document, include the document's object ID in the `OBJECT_LIST` argument.

To determine whether a user or group has entries in an ACL that give at least the permission level identified in the `LEVEL` argument or better to the user or group, include the object ID of the ACL in the `OBJECT_LIST` argument.

The method ignores all object IDs that do not represent SysObjects or ACLs or object IDs that are incorrectly identified (too few characters, too many characters, and so forth).

## Related administration methods

None

## Examples

This example checks to determine whether the user LibbyLoo has at least Write permission on three documents:

```
EXECUTE check_security WITH user_name='LibbyLoo',  
level=5,object_list='09000001734912ac  
0900000153813f2b 0900000116572af3'
```

This example determines whether the group Engineering has accessor entries that give the group at least Read permission in the ACLs included in the object list:

```
dmAPIGet ("apply,s0,NULL,CHECK_SECURITY,  
GROUP_NAME,S,Engineering,LEVEL,I,3,  
OBJECT_LIST,S,4500000112562e4c 450000017351213c
```

# CLEAN\_LINKS

**Purpose** On Windows, this method removes unneeded dmi\_linkrecord objects and resets security on file store objects to the original state. On UNIX, this method removes unneeded dmi\_linkrecord objects and the auxiliary directories and symbolic links.

## Syntax

```
EXECUTE clean_links [WITH force_active=true|false]
dmAPIGet ("apply, session, NULL, CLEAN_LINKS
[, FORCE_ACTIVE, B, T|F] ")
```

## Arguments

Argument Name	Datatype	Value	Description
FORCE_ACTIVE	B	T (TRUE) or F (FALSE)	TRUE directs the server to clean the links in all sessions. FALSE directs the server to clean links only in inactive sessions. The default is FALSE.

## Return value

CLEAN\_LINKS returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Superuser privileges to use this method.

## General notes

Linked storage areas are handled differently on UNIX and Windows platforms. Consequently, the behavior of CLEAN\_LINKS is slightly different on each platform.

On UNIX, when a linked storage area is referenced, the server creates a `dmi_linkedrecord` object and auxiliary directories and symbolic links. On Windows, the server creates a `dmi_linkrecord` object and resets the security of the linked storage object. Generally, the server removes unneeded linkrecord objects and, on UNIX, any unneeded auxiliary directories and symbolic links. On Windows, the server typically resets the linked storage object's security as needed. However, there are conditions that do not allow the server to do this work.

You can use `CLEAN_LINKS` to perform this work manually. We recommend running `CLEAN_LINKS` regularly. (Note that `CLEAN_LINKS` is run automatically whenever the server is restarted.)

To determine if you need to run `CLEAN_LINKS`, run `LIST_SESSIONS` and compare the reported session IDs (in `session[x]`) to the values in the `session_id` attributes of the `dmi_linkrecord` objects. If the `session_id` attribute in a link record object references an inactive session (a session not reported in `LIST_SESSIONS`), that link record object is not needed. Depending on how many unneeded link record objects you find, you may want to run `CLEAN_LINKS` to remove them and perform the other, associated clean up work.

## Related administration methods

[LIST\\_SESSIONS, page 244](#)

## Examples

This example runs `CLEAN_LINKS` against only inactive sessions because the `FORCE_ACTIVE` argument is defaulted to `FALSE`:

```
EXECUTE clean_links
```

The following example removes the unneeded link record objects for all repository sessions, active and inactive:

```
dmAPIGet ("apply, c, NULL, CLEAN_LINKS, FORCE_ACTIVE, B, T")
```

# DB\_STATS

**Purpose** Returns a set of statistics about database operations for the current session.

## Syntax

```
EXECUTE db_stats [WITH clear = true|false]
dmAPIGet ("apply, session, NULL, DB_STATS
[, CLEAR, B, T|F] ")
```

## Arguments

**Table 3-7. DB\_STATS arguments**

Argument	Datatype	Value	Description
CLEAR	B	T (TRUE) or F (FALSE)	Indicates whether you want to clear the counters. The default is FALSE.

## Return value

DB\_STATS returns a collection with one result object, described in [Table 3-8, page 191](#).

**Table 3-8. Query result object attributes for DB\_STATS administration method**

Attribute	Datatype	Description
updates	integer	Number of SQL update operations performed
inserts	integer	Number of SQL insert operations performed
deletes	integer	Number of SQL delete operations performed
selects	integer	Number of SQL select operations performed

Attribute	Datatype	Description
rpc_ops	integer	Number of RPC calls between Content Server and the RDBMS  <b>Note:</b> This may not be implemented for all databases.
ddl_ops	integer	Number of data definition operations performed  Data definition operations are operations such as create table or drop table.
max_cursors	integer	Maximum number of concurrently open cursors

## Permissions

Anyone can use this method.

## General notes

DB\_STATS returns a set of statistics about database operations for the current session. The statistics are counts of the numbers of:

- Inserts, updates, deletes, and selects executed
- Data definition statements executed
- RPC calls to the database
- Maximum number of cursors opened concurrently during the session

## Related administration methods

None

## Examples

The following example uses EXECUTE to invoke DB\_STATS. It returns the statistics for the current database session and resets the counters to zero:

```
EXECUTE db_stats WITH clear=true
```

This example invokes DB\_STATS using the Apply method. It returns the statistics for the current database session and resets the counters to zero:

```
dmAPIGet ("apply, c, NULL, DB_STATS, CLEAR, B, T")
```



# DELETE\_REPLICA

**Purpose** Removes a content file from a component of a distributed storage area.

## Syntax

```
EXECUTE delete_replica FOR 'content_object_id'
WITH STORE='storage_name'

dmAPIGet ("apply, session, content_obj_id, DELETE_REPLICA,
STORE, S, storage_name")
```

## Arguments

**Table 3-9. DELETE\_REPLICA arguments**

Argument	Datatype	Value	Description
STORE	S	<i>storage_name</i>	Specifies the component storage area that contains the file to remove. This is a required argument. Use the storage area's name as defined in its storage object.

## Return value

DELETE\_REPLICA returns a collection with one query result object. The object has one Boolean attribute that indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Superuser privileges to use this method.

## General notes

DELETE\_REPLICA removes a content file from a storage area that is a component of a distributed storage area. Using DELETE\_REPLICA also modifies the replica record object associated with the content. The replica record maintains the information that allows the

server to fetch the content from any of the component storage areas. When you remove a file from a component storage area using DELETE\_REPLICA, this record is updated also.

To use DELETE\_REPLICA, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use DELETE\_REPLICA in the session.

## Related administration methods

[IMPORT\\_REPLICA](#), page 234

## Examples

This example removes the content file associated with the content object identified by 06000001684537b1 from the distcomp\_2 storage area:

```
EXECUTE delete_replica FOR '06000001684537b1'  
WITH STORE='distcomp_2'
```

The next example uses the Apply method to perform the same operation:

```
dmAPIGet ("apply,c,06000001684537b1,DELETE_REPLICA,  
STORE,S,distcomp_2")
```

# DESTROY\_CONTENT

**Purpose** Removes content objects from the repository and their associated content files from storage areas.

## Syntax

```
EXECUTE destroy_content FOR 'content_obj_id'  
dmAPIGet ("apply, session, content_obj_id, DESTROY_CONTENT")
```

## Arguments

DESTROY\_CONTENT has no arguments.

## Return value

DESTROY\_CONTENT returns a collection with one query result object. The object has one Boolean attribute that indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

DESTROY\_CONTENT removes orphaned content objects. An orphaned content object is a content object that has no values in its parent\_id attribute. The method also removes the content file referenced by the orphaned content object if there are no other content objects that reference that file.

The DESTROY\_CONTENT method is the method invoked by dmclean.

To run DESTROY\_CONTENT, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use DESTROY\_CONTENT in the session.

**Note:** Do not use this method to archive content. It removes the file's content object, rather than marking it off-line.

## Destroying content in content-addressed storage

If the storage area defined in the content object is a content-addressed storage system (a ca store storage area), Content Server first determines whether the storage area allows content to be deleted. If not, DESTROY\_CONTENT fails with an error. If the storage system allows deletions, DESTROY\_CONTENT checks the retention period specified for the content represented by the content object. If there are multiple addresses recorded for the content, Content Server checks the retention period stored with each address. The retention period for all addresses must be expired before Content Server can destroy the content.

**Note:** Some operations, such as those that modify the metadata or object replication, result in the generation of a new content address for a content file. These additional addresses are stored in the `i_contents` attribute of a subcontent object.

If that period has not expired, DESTROY\_CONTENT cannot remove the file from the storage area. The method fails with an error. If the retention period has expired or there is none set, the method removes the content.

## Related administration methods

[PURGE\\_CONTENT](#), page 283

## Examples

This example uses the EXECUTE statement to invoke DESTROY\_CONTENT:

```
EXECUTE destroy_content FOR '06000001684537b1'
```

This example uses the Apply method to invoke DESTROY\_CONTENT:

```
dmAPIGet ("apply,s0,06000001684537b1,DESTROY_CONTENT")
```

# DO\_METHOD

**Purpose** Executes an external program, a Docbasic script, or a Java method.

## Syntax

```
EXECUTE do_method WITH METHOD='method_name'
{, arguments=value}

dmAPIGet ("apply, session, NULL, DO_METHOD, METHOD, S, method_name,
{, arguments, datatype, value}")
```

## Arguments

**Table 3-10. DO\_METHOD arguments**

Argument	Datatype	Value	Description
SAVE_RESULTS	B	T (TRUE) or F (FALSE)	Indicates if you wish to save the results into a document.
ARGUMENTS	S	<i>command-line arguments</i>	<p>Specifies the command-line arguments for the program.</p> <p>If the method is to be executed on the application server, UTF-8 characters are accepted for the argument strings.</p> <p>If the method is to be executed on the method server or Content Server, only characters from the server os code page are accepted for the argument strings.</p> <p>Refer to <a href="#">Specifying the arguments, page 201</a>, for more information.</p>

Argument	Datatype	Value	Description
TIME_OUT	I	<i>value</i>	Specifies the length in seconds of the default time-out period for the program that you are executing. Refer to <a href="#">Defining a time out period , page 202</a> for more information.
METHOD	S	<i>method_name</i>	Name of the method object representing the script, program, or Java method to execute.  This argument is required.
LAUNCH_DIRECT	B	TRUE or FALSE	Indicates whether to execute the program using the Windows (UNIX) system API or the exec API. Set this to TRUE to use the system API. By default, it is FALSE, which uses the system call. Refer to <a href="#">Launching directly, page 202</a> for more information.  This argument is ignored if the method's use_method_server attribute is TRUE.
LAUNCH_ASYNC	B	TRUE or FALSE	Indicates whether to execute the program asynchronously. TRUE executes the method asynchronously. The default is FALSE.  Setting this argument to TRUE is ignored if the SAVE_RESULTS argument is also TRUE. Refer to <a href="#">Launching asynchronously, page 203</a> for more information.

Argument	Datatype	Value	Description
RUN_AS_SERVER	B	TRUE or FALSE	<p>Indicates whether to run the method under the server's account. If the argument is not set, the server uses the value of the method's run_as_server attribute.</p> <p>You must have Sysadmin or Superuser privileges to set this to TRUE on the command line if the method identified in the METHOD argument has its run_as_server attribute set to FALSE.</p> <p>This argument must be TRUE for methods that have their use_method_server attribute set to TRUE.</p>
TRACE_LAUNCH	B	TRUE or FALSE	<p>Indicates whether to generate tracing information for the method.</p> <p><a href="#">Recording tracing information and program output, page 204</a> describes where the information is stored.</p>

## Return value

A DO\_METHOD function returns a collection that contains one query result object. [Table 3–11, page 200](#) lists the attributes of this object.

**Table 3-11. Attributes of the query result object returned by DO\_METHOD**

Attribute	Datatype	Description
result	Integer or ID	<p>If the program you launched was dmfilesan or dmclean, this attribute holds the object ID of the script run by the utility.</p> <p>Otherwise, the attribute contains an integer value indicating the success or failure of the program.</p> <p>For methods executed on an application server, the values are:</p> <p>0, meaning a status of HTTP/1.1 2xx  1, meaning a status of HTTP/1.1 5xx  -1, for any other status</p>
result_doc_id	ID	<p>Contains the object ID of the document that contains the output of the program. (Note that the output is captured even if the program times out.)</p> <p>This attribute is present only when the SAVE_RESULTS argument is set to TRUE.</p>
launch_failed	Boolean	<p>Indicates whether the program was successfully executed. T means that the program was not launched and the method_return_val attribute is meaningless. F means that the program was launched and the return value is the exit code of the process when the method terminated.</p>
method_return_val	Integer	<p>Contains the return value of the executing program.</p> <p>For methods executed on the application server, the values are:</p> <p>0, for a status of HTTP/1.1 2xx  1, for any other status</p> <p>For all other methods, if the program times out, this value is 0.</p> <p>If launch_async is T, then method_return_val is always 0.</p>
os_system_error	String	<p>Contains an operating system error if one occurs. This attribute can be empty.</p>
timed_out	Boolean	<p>Indicates whether the DO_METHOD was terminated due to a timeout. T means the method was terminated while waiting for the</p>



Attribute	Datatype	Description
		program to complete. F means the method received a response.
time_out_length	Integer	The length of the time-out period.

## Permissions

Anyone can use this method.

If the method you are executing has the `run_as_server` attribute set to `FALSE` in its method object, you must have at least Sysadmin privileges to override that setting in the `DO_METHOD` command line.

## General notes

Use `DO_METHOD` to execute a Docbasic script, a Java method, or other executable program. (Docbasic is Documentum's interpretive language.) You can direct the execution of the method to the method server, the application server, or Content Server. Which you choose depends on the language in which the program is written and how your site is configured. For details about each of the execution agents and how to direct a method to the agents, refer to [Chapter 4, Methods and Jobs](#), in *Documentum Content Server Administrator's Guide*.

To use `DO_METHOD`, the invoked script or program must be defined in the repository by a method object. A method object contains attributes that tell the server the program's command line and arguments and provide parameters for executing the program. If the program is a Docbasic script, the script is stored as the content of the method. (For information about creating method objects, refer to [Creating a method object, page 141](#) of the *Content Server Administrator's Guide*. Information about Docbasic scripts can be found in Documentum's *Docbasic User Guide and Reference Manual*.)

## Specifying the arguments

There are no restrictions on the format of the argument list for methods written in `dmbasic` or methods to be executed through the Content Server or the method server.

If the method is to be executed using the application server (`use_method_server` is `T` and `method_type` is `java`), you must pass both argument names and values in the `ARGUMENTS value`. The format for `value` is:

```
-argument_name argument_value
```

Separate multiple arguments with a single space. For example:

```
dmAPIGet ("apply,c,NULL,DO_METHOD,METHOD,S,payroll_report,
SAVE_RESULTS,B,T,ARGUMENTS,S,-docbase accounting
-user auditor
```

```
-ticket DM_TICKET=0000000222a02024.accounting@host01")
```

or

```
EXECUTE do_method WITH METHOD='payroll_report',  
ARGUMENTS='-docbase accounting  
-user auditor  
-ticket DM_TICKET=0000000222a02024.accounting@host01'
```

If you are directing the DO\_METHOD to the application server, Content Server encodes the arguments using application/x-www-form-urlencoded format. For example, suppose you issued the following DO\_METHOD:

```
dmAPIGet ("apply,c,NULL,DO_METHOD,METHOD,S,paymethod,  
SAVE_RESULTS,B,T,TRACE_LAUNCH,B,T,  
ARGUMENTS,S,-docbase accounting -user paymaster  
-ticket DM_TICKET=0000000222a02054.accounting@host01  
-document "weekly record"
```

Content Server sends the arguments in the following format:

```
docbase=accounting&user=paymaster&ticket=DM_TICKET%  
3D0000000222a02054.accounting%4Dhost01&document=weekly+record
```

## Defining a time out period

The TIME\_OUT argument defines a time out period for the program or script you are executing. Assigning a value to this argument overrides any value assigned to the timeout\_default attribute in the program's method object.

The value that you specify cannot be greater than the value assigned to the method object's timeout\_max attribute or less than the value assigned to the object's timeout\_min attribute. If the maximum or minimum time that you specify on the DO\_METHOD command line violates this rule, the server ignores your specification and uses the timeout\_max or timeout\_min value specified in the method object.

## Launching directly

When you execute DO\_METHOD using Content Server as the execution agent, the server calls the Windows or Unix (depending on the platform on which the method is running) system API to execute it by default. If you set LAUNCH\_DIRECT to TRUE, the server calls the exec API instead. This API executes the program or script directly, instead of calling a shell script, which provides the advantage of better error reporting.

To execute with LAUNCH\_DIRECT set to TRUE, the method's method\_verb attribute must contain a fully qualified pathname.

## Launching asynchronously

There are two ways to launch a program or script asynchronously:

- Use the ARGUMENTS argument to DO\_METHOD to append an ampersand (&) to the end of the command line arguments. If there are multiple command-line arguments, the ampersand must be the last argument in the list.

For example,

```
EXECUTE do_method
WITH method = 'validate', arguments = '&'
```

- Set the DO\_METHOD's LAUNCH\_ASYNC argument to TRUE.

For example,

```
EXECUTE do_method
WITH method = 'validate', launch_async = TRUE
```

If you launch asynchronously and the method is executing on the application server, it is not possible to capture the method's output.

## Saving results

For DO\_METHOD methods that are executing on the application server, Documentum provides a simple, example interface that captures the program output so the output can be saved into the repository if SAVE\_RESULTS is TRUE. This interface is described fully in [Chapter 4, Methods and Jobs](#), in the *Content Server Administrator's Guide*.

## Running as the server account

By default, the program invoked by the DO\_METHOD runs as the logged-in user. If you want the program to execute under the Content Server's account, you can:

- Set the run\_as\_server attribute in the associated method object to TRUE and set the RUN\_AS\_SERVER argument in the command line to TRUE.
- Set only the RUN\_AS\_SERVER argument in the command line to TRUE.

By default, both the run\_as\_server attribute and the RUN\_AS\_SERVER argument are FALSE. To run as the server account, either both must TRUE or you must override the attribute setting by setting the RUN\_AS\_SERVER argument to TRUE. Overriding the attribute in the DO\_METHOD command line by setting the RUN\_AS\_SERVER argument to TRUE requires Sysadmin or Superuser privileges. (Overriding the attribute if it is set to TRUE by setting the RUN\_AS\_SERVER argument to FALSE requires no special privileges.)

If you execute DO\_METHOD using either the method server or the application server as the execution agent, you must set both the method's run\_as\_server attribute to TRUE and the RUN\_AS\_SERVER argument to TRUE.

**Notes:**

- If LAUNCH\_DIRECT is set to TRUE, either on the command line or in the method's attribute, RUN\_AS\_SERVER must also be set to TRUE. If it is not, the method does not execute correctly.
- Content Server uses the assume user program to run procedures and print jobs as the logged-in user. If you disable the assume user program (by setting the assume\_user\_location attribute in the server config object to a blank), Content Server runs all procedures and all print jobs under its account.

**Ensuring security on the application server**

There are two security issues to consider when using an application server to execute a DO\_METHOD that invokes a Java servlet or method:

- Determining the origin of the HTTP\_POST request
- Login without passwords (this is only possible on Windows platforms)

Issuing a DO\_METHOD to execute on the application server sends an internal HTTP\_POST request to the application server. The invoked servlet ensures that the generated request comes from a machine that hosts a Content Server by checking the IP address of the sender against a list of repositories. This list is set up when the application server is installed and configured. If the sender's IP address doesn't match the IP address of a repository host, the request fails.

The application server runs as the Content Server installation owner. Consequently, the servlet it invokes to execute DO\_METHOD calls also runs as the installation owner. On Windows platforms, the current operating system user is allowed to log in to the repository without providing a password. Consequently, a servlet or an invoked Java method can log into the repository with superuser privileges without providing a password.

If you write a method that logs in in that manner, you may want to ensure that the actual user who issues the DO\_METHOD to invoke the method has the appropriate privileges to execute the program as a superuser. To do this, send the current user's name and a login ticket to the method in the arguments. The method can use these to log in and check the privileges of the user before connecting as the current operating system user.

**Recording tracing information and program output**

Trace information and program output are two different sets of information. Trace information is information about the DO\_METHOD invocation and its success or failure. The program output is the results returned by the program called by the DO\_METHOD.

Setting TRACE\_LAUNCH to TRUE logs tracing information, of up to 2047 characters, to the repository log. Setting SAVE\_RESULTS to TRUE saves the execution output of the invoked program.

## Related administration methods

[EXEC\\_SQL](#), page 211

[HTTP\\_POST](#), page 228

## Examples

This example executes the `update_accounts` procedure, specifying a timeout period of 120 seconds (2 minutes):

```
EXECUTE do_method
WITH method_name='update_accounts',
time_out=120
```

The following example executes the user-defined procedure `run_rpt_newaccts` and saves the results to a document:

```
dmAPIGet ("apply, s0, NULL, DO_METHOD,
METHOD, S, run_rpt_newaccts, SAVE_RESULTS, B, T")
```

This example executes a method on the application server:

```
dmAPIGet ("apply, s0, NULL, DO_METHOD,
METHOD, S, check_vacation, SAVE_RESULTS, B, T,
ARGUMENTS, S, -doibase HumanRSRC -user adminasst
-ticket DM_TICKET=0000000221c02052.HumanRSRC@hr025
-document "monthly payroll")
```

# DROP\_INDEX

**Purpose** Destroys a user-defined object type index.

## Syntax

```
EXECUTE drop_index [[FOR] 'dmi_index_id']  
[WITH name = 'index_name']
```

With EXECUTE, do not include the FOR clause if you include the NAME argument.

```
dmAPIGet ("apply, session, dmi_index_id, DROP_INDEX  
[, NAME, S, index_name] ")
```

With Apply, specify the index object ID as NULL if you include the NAME argument.

## Arguments

**Table 3-12. DROP\_INDEX arguments**

Argument	Datatype	Value	Description
NAME	S	<i>index_name</i>	Identifies the index by the name of its index (dmi_index) object.

## Return value

DROP\_INDEX returns a collection that contains one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Superuser privileges to use this method.

## General notes

You can obtain an object type index's name or object ID from the dmi\_index type table. Each row in this table represents one index in the repository.

## Related administration methods

[FINISH\\_INDEX\\_MOVES](#), page 215

[MAKE\\_INDEX](#), page 253

[MOVE\\_INDEX](#), page 271

## Examples

These examples illustrate using EXECUTE to drop a user-defined index on the dm\_user object type. The first example identifies the index by its name, user\_index, and the second example identifies the index by its object ID.

```
EXECUTE drop_index WITH name='user_index'  
EXECUTE drop_index FOR '1f00000011231563a'
```

These examples illustrate using the Apply method to drop a user-defined index on the dm\_user object type. The first example identifies the index by its name, user\_index, and the second example identifies the index by its object ID.

```
dmAPIGet ("apply, s0, NULL, DROP_INDEX, NAME, S, user_index")  
dmAPIGet ("apply, s0, 1f00000011231563a, DROP_INDEX")
```

# ESTIMATE\_SEARCH

**Purpose** Returns the number of results matching a particular SEARCH condition.

## Syntax

```
EXECUTE estimate_search [[FOR] 'fulltext_index_obj_id']
WITH [name = 'index_name'] [,type = 'object_type']
[,query = 'value']
```

With EXECUTE, do not include the FOR clause if you include the NAME argument.

```
dmAPIGet("apply,session,fulltext_index_obj_id,ESTIMATE_SEARCH
[,NAME,S,index_name] [,TYPE,S,object_type]
[,QUERY,S,value] ")
```

With Apply, specify the fulltext index object ID as NULL if you include the NAME argument.

## Arguments

**Table 3-13. ESTIMATE\_SEARCH arguments**

Argument	Datatype	Value	Description
NAME	S	<i>index_name</i>	Identifies the index to be searched. This is optional. You can specify the object ID of the fulltext index object instead.
TYPE	S	<i>object_type</i>	Identifies the type of objects to be searched. All subtypes of the specified type are included in the search also.  If not included, the method searches all object types in the index.
QUERY	S	<i>value</i>	Defines the word or phrase for which you are searching. You can use a Boolean Plus expression for the value or a particular word or phrase.



Argument	Datatype	Value	Description
			If not included, the method's return value represents all objects in the index.

## Return value

ESTIMATE\_SEARCH returns one of the following:

- The exact number of matches that satisfy the SEARCH condition if the user running the method is a superuser or there are more than 25 matches.
- The number 25 if there are 0-25 matches and the user running the method is not a superuser.
- The number -1 if an error occurs during execution.

Errors are logged in the session log file.

## Permissions

Any user can execute this method. However, the return value is affected by the user's privilege level. Refer to the General Notes for an explanation.

To execute this method, the server.ini key, use\_estimate\_search, must be set to TRUE. The key defaults to FALSE.

## General notes

ESTIMATE\_SEARCH is a useful tool for fine-tuning a SEARCH condition in a SELECT statement. ESTIMATE\_SEARCH provides an estimate of the number of results a query will return. Use it to determine how selective or unselective a particular query is. Do not use it to determine the exact number of results a particular query will return. It is intended only as a way to tune queries. Factors such as security affect the actual number of results returned when the actual query is run.

If the user executing the method is a superuser, the method returns the exact number of matches regardless of how few or how many matches are returned.

If the user is not a superuser, the method returns the exact number of matches if the number is greater than 25. If the number of matches is 0-25, the method always returns the number 25.

## Related administration methods

None

## Example

```
EXECUTE estimate_search WITH name='filestore2_indx',  
type='dm_document',query='Competitor Evaluation'  
  
dmAPIGet("apply,c,NULL,ESTIMATE_SEARCH  
,NAME,S,filestore2_indx,TYPE,S,dm_document  
,QUERY,S,Competitor Evaluation")
```

# EXEC\_SQL

**Purpose** Executes SQL statements.

## Syntax

```
EXECUTE exec_sql WITH query='sql_query'
dmAPIGet ("apply, session, NULL, EXEC_SQL,
QUERY, S, sql_query]")
```

## Arguments

Table 3-14. EXEC\_SQL arguments

Argument	Datatype	Value	Description
QUERY	S	<i>sql_query</i>	Defines the query that you want to execute.

## Return value

EXEC\_SQL returns a collection that contains one query result object. The object has one Boolean attribute whose value is TRUE if the query succeeded and FALSE if it was unsuccessful.

## Permissions

You must have superuser privileges to use this method.

## General notes

EXEC\_SQL executes any SQL statement with the exception of SQL SELECT statements.

If you use the Apply method to execute the method and the query contains commas, you must enclose the entire query in single quotes. For example:

```
dmAPIGet ("apply, s0, NULL, EXEC_SQL, QUERY, S, 'create table
mytable (name char(32), address char(64))'")
```

In the EXECUTE statement, character string literals must always be single-quoted:

```
EXECUTE exec_sql  
with query='create table mytable (name char(32), address char(64))'
```

## Related administration methods

[DO\\_METHOD](#), page 197

## Examples

Refer to the General Notes.

# EXPORT\_TICKET\_KEY

**Purpose** Returns a login ticket key.

## Syntax

```
EXECUTE export_ticket_key WITH PASSWORD='password'
dmAPIGet ("apply, session, NULL, EXPORT_TICKET_KEY, PASSWORD, S, password")
```

## Arguments

**Table 3-15. EXPORT\_TICKET\_KEY arguments**

Argument	Datatype	Value	Description
PASSWORD	S	<i>password</i>	User-defined password used to encrypt the returned login ticket key.

## Return value

If successful, the method returns the login ticket key used by the repository as an encrypted and ASCII-encoded string. The returned key is encrypted using the password provided as an argument and then encoded as an ASCII string.

If the method fails, the method returns NULL.

## Permissions

You must have Superuser privileges to execute this method.

## General notes

Use EXPORT\_TICKET\_KEY when you want to copy a login ticket key from one repository to another, to configure a trust relationship between the two repositories. (For a description of trusted repositories, refer to [Trusted repositories, page 37](#), in *Content Server Fundamentals*.)

## Related methods

[IMPORT\\_TICKET\\_KEY](#), page 236

[RESET\\_TICKET\\_KEY](#), page 298

`IDfSession.exportTicketKey()`

## Example

```
EXECUTE export_ticket_key WITH PASSWORD='myword'  
ticket=dmAPIGet("apply,s0,EXPORT_TICKET_KEY,  
PASSWORD,S,myword")
```

---

# FINISH\_INDEX\_MOVES

**Purpose**            Completes all unfinished object type index moves.

## Syntax

```
EXECUTE finish_index_moves  
dmAPIGet ("apply, session, NULL, FINISH_INDEX_MOVES")
```

## Arguments

FINISH\_INDEX\_MOVES has no arguments.

## Return value

FINISH\_INDEX\_MOVES returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Superuser privileges to use this method.

## General notes

FINISH\_INDEX\_MOVES completes an interrupted move operation for an object type index.

Moving an object type index is not an atomic operation. Consequently, if a move operation is interrupted, the repository may be left with a `dmi_index` object that has no associated index. To resolve this situation, use FINISH\_INDEX\_MOVES. This method scans the `dmi_index` table and completes all unfinished index moves.

## Related administration methods

[DROP\\_INDEX](#), page 206

[MAKE\\_INDEX](#), page 253

[MOVE\\_INDEX](#), page 271

## Example

Refer to the syntax description.



# GET\_FILE\_URL

**Purpose** Returns the URL to a particular content file.

## Syntax

```
EXECUTE get_file_url FOR object_id
WITH format='format_name' [,page=page_number,]
[page_modifier='value']

dmAPIGet ("apply,session,object_id,GET_FILE_URL,
FORMAT,S,format_name[,PAGE,I,page_number]
[,PAGE_MODIFIER,S,value]")
```

*object\_id* is the object ID of the document that contains the content file.

## Arguments

**Table 3-16. GET\_FILE\_URL arguments**

Argument	Datatype	Value	Description
FORMAT	S	<i>format_name</i>	Name of the content format, as specified in the format object.
PAGE	I	<i>page_number</i>	Page number of the content file.  Setting this to -1 directs the server to return the URLs for all primary content pages and renditions that have the specified format.  The default is 0.

## Return value

GET\_FILE\_URL returns a collection consisting of one object with five attributes. One attribute indicates the success or failure of the method call. If the call is successful, the other attribute values can be concatenated to compose the URL.

The attributes are:

result	Boolean attribute whose value indicates whether the method was successful. T (TRUE) indicates successful completion and F (FALSE) indicates failure.
base_url	The value in the base_url attribute is the base URL used by the Web server to retrieve the content. This value is retrieved from the base_url attribute of the storage area that contains the file.
store	The value in the store attribute is the name of the storage area that contains the file.
path	The value in the path attribute is a path to the file. The path is relative to the storage area identified in the store attribute.
ticket	Contains an encryption of the path value plus a time stamp.

## Permissions

You must have at least Read permission on the object or Sysadmin privileges to use this method.

## General notes

Use GET\_FILE\_URL in an application when you want to access content using a Web server or a streaming server.

## Related administration methods

None

## Example

```
EXECUTE get_file_url FOR 090000215400ac12
WITH format='jpeg_th',page=0
dmAPIGet ("apply,s0,090000215400ac12,GET_FILE_URL,
FORMAT,S,jpeg_th,PAGE,I,0")
```

# GET\_INBOX

**Purpose** Returns items from an Inbox.

## Syntax

```
EXECUTE get_inbox [WITH name='user_name'][,category=value]
[,batch=value]{,order_by='attr_name [asc|desc]'}]
dmAPIGet ("apply,session,NULL,GET_INBOX
[,NAME,S,user_name] [,CATEGORY,I,value] [,BATCH,I,value]
{,ORDER_BY,S,attr_name [asc|desc]}")
```

## Arguments

**Table 3-17. GET\_INBOX arguments**

Argument	Datatype	Value	Description
NAME	S	<i>user_name</i>	User whose Inbox items you want to retrieve. The default value is the session user.
CATEGORY	I	<i>value</i>	<p>The kind of items to retrieve. Valid values are:</p> <ul style="list-style-type: none"> <li>1, for workflow tasks</li> <li>2, for router tasks</li> <li>4, for notifications</li> <li>8, for completed router tasks</li> </ul> <p>To retrieve multiple kinds of items, use the sum of the integers representing the items. For example, to retrieve workflow tasks and notifications, specify the value as 5.</p> <p>The default is 3, workflow tasks and router tasks.</p>

Argument	Datatype	Value	Description
BATCH	I	<i>value</i>	<p>Number of items returned by each Next method issued against the collection returned by GET_INBOX.</p> <p>The default value is 0, meaning return all rows.</p>
ORDER_BY	S	<i>attr_name</i>	<p>Attribute by which to order the returned items.</p> <p>The attribute must be an attribute of the dmi_queue_item object type.</p> <p>Including asc sorts the returned items in ascending order. desc sorts the items in descending order. The default is ascending order.</p> <p>The default ordering is by r_object_id</p>

## Return value

GET\_INBOX returns a collection containing the Inbox items in query result objects. The query result objects have 49 attributes. The first 41 are the attributes of the dmi\_queue\_item object representing the Inbox item. The attribute names in the query result object match the names of the attributes of the queue items. The remaining 8 attributes identify the SysObject associated with the Inbox item, if any. Generally, these attributes have values if the Inbox item is a workflow or router task. Notification items have no values in these 8 attributes. (Refer to the General Notes for a more detailed explanation.)

[Table 3–18, page 221](#), lists the SysObject-related attributes and their corresponding SysObject attribute.

**Table 3-18. SysObject-related attribute names for GET\_INBOX query result objects**

Query Result Attribute Name	Associated SysObject Attribute
pkg_object_id	r_object_id
pkg_object_name	object_name
pkg_object_type	r_object_type
pkg_content_type	a_content_type
pkg_application_type	a_application_type
pkg_link_cnt	r_link_cnt
pkg_lock_owner	r_lock_owner
pkg_is_virtual_doc	r_is_virtual_doc

## Permissions

Any user can issue this method to return either his or her own Inbox items or the Inbox items of another user.

## General notes

With one exception, you can specify the arguments in any order. The exception is `ORDER_BY`. This argument must appear last.

Generally, `GET_INBOX` returns one query result object for each item in user's Inbox. However, if a particular task has multiple objects associated with it, the method returns one query result object for each object associated with the task. The queue item attribute values for the multiple objects will be the same. Only the values of the eight SysObject-related attributes will be different.

The eight SysObject-related attributes contain null values in the following cases:

- The Inbox item is an event notification and therefore has no associated object.
- The Inbox item is a task, but its associated object has been deleted from the repository.
- The Inbox item is a task but the user who issues the method doesn't have at least Browse permission on the associated object.
- The Inbox item is a workflow task, such as a Beginning task, that has no package attached to it.
- The Inbox item represents an object in a remote repository.

## Related administration methods

None

## Examples

This example returns all tasks and notifications for the user issuing the method:

```
EXECUTE get_inbox WITH category=7
```

This example returns just notifications for virtual user named " my\_app" in batches of 20:

```
dmAPIGet ("apply, c, NULL, GET_INBOX, NAME, S, my_app,  
CATEGORY, I, 4, BATCH, I, 20")
```

# GET\_LAST\_SQL

**Purpose**           Retrieves the SQL translation of the last DQL statement issued.

## Syntax

```
EXECUTE get_last_sql  
dmAPIGet ("apply, session, NULL, GET_LAST_SQL ")
```

## Arguments

None

## Return value

GET\_LAST\_SQL returns a collection with one query result object. The result object has one attribute whose value is the last SQL statement. To see the statement, issue a Next on the collection and then dump the collection.

If the last SQL tracing option is turned off, this method returns the following error message:

No SQL Statement is available because Last SQL Trace is disabled.

## Permissions

Any one can issue this method.

## General notes

The last SQL tracing feature is turned on by default when a server starts. If the feature is turned off, you can turn it on using the `last_sql_trace` option of the `SET_OPTIONS` method. (Refer to [SET\\_OPTIONS](#), page 312, for instructions.)

## Related administration methods

None

## Examples

```
EXECUTE get_last_sql  
dmAPIGet ("apply,S0,NULL,GET_LAST_SQL")
```



# GET\_PATH

**Purpose** Returns the directory location of a content file stored in a distributed storage area.

## Syntax

```
EXECUTE get_path [FOR] 'content_obj_id'
[WITH store = 'value']

dmAPIGet ("apply, session, content_obj_id, GET_PATH
[, STORE, S, value] ")
```

## Arguments

**Table 3-19. GET\_PATH arguments**

Argument	Datatype	Value	Description
STORE	S	<i>storage_</i> <i>component</i> <i>_name</i>	Specifies a storage area that contains the file whose full path you want to determine. This is an optional argument. Use the storage area's name as defined in its storage object.

## Return value

Returns the directory path of the specified content file.

## Permissions

Anyone can use this method.

## General notes

In a content server configuration, the GET\_PATH function is executed by the content server.

If you do not include the STORE argument, the method looks in the local component of the distributed storage area. If the file isn't found in the local area, the method attempts to create a replica of the file in the local area and returns the path of the local replica.

## Related administration methods

None

## Examples

The following examples return the file path for the content file represented by the content object whose object ID is 060000027435127c in the storage1 storage area:

```
EXECUTE get_path FOR '060000027435127c'  
WITH store='storage1'  
  
dmAPIGet ("apply,c,060000027435127c,GET_PATH,  
STORE,S,storage1")
```

# GET\_SESSION\_DD\_LOCALE

**Purpose** Returns the locale in use for the current session.

## Syntax

```
EXECUTE get_session_dd_locale  
dmAPIGet ("apply, session, NULL, GET_SESSION_DD_LOCALE")
```

## Arguments

None

## Return value

The method returns a collection with one result object. The result object has one attribute, named `dd_locale`. The value of this attribute is the locale for the session.

## Permissions

Anyone can use this method.

## General notes

None

## Related administration methods

None

## Example

```
dmAPIGet ("apply, S0, NULL, GET_SESSION_DD_LOCALE")
```

# HTTP\_POST

**Purpose** Sends an HTTP\_POST request invoking a Java servlet.

## Syntax

```
EXECUTE http_post WITH app_server_name='name'
[,arguments=argument_list][,save_response=value]
[,time_out=value][,launch_asynch=value][,trace_launch=value]

dmAPIGet ("apply, session, NULL, HTTP_POST,
APP_SERVER_NAME, S, name [, ARGUMENTS, S, argument_list]
[, SAVE_RESPONSE, I, value] [, TIME_OUT, I, value]
[, LAUNCH_ASYNCH, B, value] [, TRACE_LAUNCH, B, value] ")
```

## Arguments

**Table 3-20. HTTP\_POST arguments**

Argument	Datatype	Value	Description
APP_SERVER_NAME	S	<i>name</i>	Name of the Java servlet.  This must match a servlet identified in the <code>app_server_name</code> attribute of the server config object.
ARGUMENTS	S	<i>argument list</i>	Defines the command line arguments passed to the servlet or Java method.  UTF-8 characters are acceptable for the argument strings.

Argument	Datatype	Value	Description
SAVE_RESPONSE	I	<i>integer</i>	Indicates whether to save the response in a document in the repository. Value values are:  0, do not save 1, save the results -1, save the results if there is an error  The default is 0, do not save.
TIME_OUT	I	<i>integer</i>	The length of time, in seconds, to wait for a response to the HTTP_POST request.  The default is 60 seconds.
LAUNCH ASYNCH	B	T (TRUE) or F (FALSE)	Indicates whether to execute the HTTP_POST request asynchronously. TRUE executes the HTTP_POST asynchronously. The default is FALSE.  This argument is ignored if SAVE_RESPONSE is TRUE.
TRACE_LAUNCH	B	T (TRUE) or F (FALSE)	Indicates whether to generate tracing information for the HTTP_POST request. If TRACE_LAUNCH is TRUE, the information is stored in the server log. The default is FALSE.

## Return value

HTTP\_POST returns a collection with one query result object that has seven attributes. [Table 3–21, page 230](#), lists the attributes:

**Table 3-21. Query result object attributes for HTTP\_POST**

Attribute Name	Description
result	Indicates the status of the HTTP_POST. Possible values are:  0, indicating the status HTTP/1.1 2xx 1, indicating the status HTTP/1.1 5xx -1, indicating any other status
http_response_status	The response returned for the HTTP_POST.  Some example values are:  HTTP/1.1 2xx OK HTTP/1.1 5xx Internal Server Error  For a complete list of responses, refer to the HTTP protocol specification.
request_failed	Indicates whether the method sent an HTTP request to the application server. T (TRUE) means the method failed to send a request. F (FALSE) means the request was successfully sent to the application server.
response_doc_id	Object ID of the document in which the response is stored. This only has a value if SAVE_RESPONSE was set to TRUE.
time_taken	Length of time, in seconds, between when the request was sent and when a response was received. This represents the execution time of the Java method plus the time used for communication between Content Server and the application server.
timed_out	Indicates whether the method timed out. T (TRUE) means that the HTTP_POST method timed out while waiting for a response. F (FALSE) means that a response was received.
time_out_length	Time, in seconds, of the time out period.

## Permissions

You must have Superuser privileges to execute this method. (Refer to [Preserving security, page 231](#), for more information about security when using this method.)

## General notes

Use HTTP\_POST to invoke a servlet or Java method installed in an application server.

If you set LAUNCH\_ASYNC to TRUE, Content Server closes the connection to the application server immediately after sending the request. If LAUNCH\_ASYNC is FALSE, the server waits for a response until a response is received or the time out period is reached.

Setting TRACE\_LAUNCH to TRUE logs tracing information about the invocation of the HTTP\_POST method and its success or failure.

The ARGUMENTS argument can contain only UTF-8 characters. Content Server encodes the arguments using application/x-www-form-urlencoded format. For example, suppose you issued the following HTTP\_POST method:

```
dmAPIGet ("apply,c,NULL,HTTP_POST,APP_SERVER_NAME,S,payroll,
SAVE_RESPONSE,B,T,TRACE_LAUNCH,B,T,
ARGUMENTS,S,-docbase accounting -user paymaster
-ticket DM_TICKET=0000000222a02054.accounting@host01
-document "weekly record"
```

Content Server sends the arguments in the following format:

```
docbase=accounting&user=paymaster&ticket=DM_TICKET%
3D0000000222a02054.accounting%4Dhost01&document=weekly+record
```

## Preserving security

The Tomcat application server, all servlets that it invokes, and the methods invoked by the servlets run as the Content Server installation owner, which is an account with Superuser privileges. Consequently, it is important that they are written and execute using adequate security precautions.

There are two primary security issues:

- Determining origin of HTTP\_POST requests
- Login without passwords (this is only possible on Windows platforms)

The application server or the servlet has no inherent way to know whether the HTTP\_POST request was sent from a Documentum Server. When you write a servlet, you must include security checking to ensure that the request comes from a Documentum Server. One recommended way is have the servlet ensure that the generated request comes from a machine that hosts a Content Server by checking the IP address of the

sender against a list of valid IP addresses. (This is how the `do_method` servlet checks the origin. Refer to *Installing Content Server* for information about how that is set up.)

On Windows platforms, the current operating system user is allowed to log in to the repository without providing a password. Consequently, a servlet or an invoked Java method can log into the repository with Superuser privileges without providing a password.

If you write a servlet or method that logs in in that manner, you may want to ensure that the actual user who issues the HTTP\_POST to invoke the program has the appropriate privileges to execute the program as a superuser. To do this, send the current user's name and a login ticket to the method in the arguments. The method can use these to log in and check the privileges of the user before connecting as the current operating system user.

## Sample servlet and method

Documentum provides a sample servlet for handling HTTP\_POST method calls and a sample method. The servlet is named `DmSampleServlet.java` and the method is named `DmSampleMethod.java`. The sample servlet and method are located in the `classes` folder under the `WEB-INF` folder (`%DM_HOME%\tomcat\webapps\DmMethods\WEB-INF\classes` or `$DM_HOME/tomcat/webapps/DmMethods/WEB-INF/classes`) The `WEB-INF` folder was set up when you installed the Apache Tomcat application server.

## Recording output

If you set `SAVE_RESPONSE` to `TRUE`, then anything that the invoked servlet writes to `HttpServletResponse.OutputStream` is saved to a document in the repository.

## Related administration methods

[DO\\_METHOD](#), page 197

## Examples

The following examples send an HTTP\_POST request to the Java servlet called `DevAppSrv`. The content of the request passes a repository name, user name, a login ticket, and a document name to the Java servlet.

```
EXECUTE http_post WITH app_server_name='DevAppSrv',
save_response=1,trace_launch=T,time_out=60,
arguments='-docbase DevTest -user test_user
-ticket DM_TICKET=0000000221c02052.DevTest@dev012
-document "A Test"'
```



```
dmAPIGet ("apply, c, NULL, APP_SERVER_NAME, S, DevAppSrv,  
SAVE_RESPONSE, I, 1, TRACE_LAUNCH, B, T, TIME_OUT, I, 60,  
ARGUMENTS, S, -doibase DevTest -user test_user  
-ticket DM_TICKET=0000000221c02052.DevTest@dev012  
-document 'A Test'")
```

# IMPORT\_REPLICA

**Purpose** Imports files from one distributed storage area into another distributed storage area.

## Syntax

```
EXECUTE import_replica FOR 'content_object_id'
WITH store='storage_name',file='file_name'
[,other_file='other_file_name']

dmAPIGet("apply,session,content_object_id,IMPORT_REPLICA,
STORE,S,storage_name,FILE,S,file_name
[,OTHER_FILE,S,other_file_name]")
```

## Arguments

**Table 3-22. IMPORT\_REPLICA arguments**

Argument	Datatype	Value	Description
STORE	S	<i>storage_name</i>	Identifies the storage area in which to place the imported content file. This must be a component of a distributed storage area. Use the storage area's name.
FILE	S	<i>file_name</i>	Identifies the file to replicate.
OTHER_FILE	S	<i>other_file_name</i>	The OTHER_FILE argument is optional. It directs the server to copy the specified Macintosh resource fork file in addition to the data file. Specify the name of the resource fork file.  Use this only when the file you are importing was created on a Macintosh.

## Return value

IMPORT\_REPLICA returns a collection with one query result object. The object has one attribute whose value indicates success (TRUE) or failure (FALSE).

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

IMPORT\_REPLICA imports files from one distributed storage area into another distributed storage area. The files are considered replicas in the target storage area.

To use IMPORT\_REPLICA, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use IMPORT\_REPLICA in the session.

## Related administration methods

[CAN\\_FETCH](#), page 177  
[DELETE\\_REPLICA](#) , page 193  
[GET\\_PATH](#), page 225  
[REPLICATE](#), page 295

## Examples

The following examples import the file mydoc into the storage area named distcomp\_2:

```
EXECUTE import_replica FOR '06000001402371e1'  
WITH store='distcomp_2',FILE='mydoc'  
  
dmAPIGet ("apply,c,06000001402371e1,  
IMPORT_REPLICA,STORE,S,distcomp_2,FILE,S,mydoc")
```

# IMPORT\_TICKET\_KEY

**Purpose** Imports a login ticket key into a repository.

## Syntax

```
EXECUTE import_ticket_key
WITH KEY_STRING='string',
PASSWORD='password'

dmAPIGet ("apply, session, NULL, IMPORT_TICKET_KEY,
KEY_STRING, S, string, PASSWORD, S, password")
```

## Arguments

**Table 3-23. IMPORT\_TICKET\_KEY arguments**

Argument	Datatype	Value	Description
KEY_STRING	S	<i>string</i>	The ASCII-encoded string returned by an EXPORT_TICKET_KEY method.
PASSWORD	S	<i>password</i>	Password used when the key was exported.

## Return value

This method returns a collection with one query result object. The object has one attribute, named *result*, that contains T (TRUE) if the method was successful or F (FALSE) if unsuccessful.

## Permissions

You must have Superuser privileges to execute this method.

## General notes

Use IMPORT\_TICKET\_KEY to import a login ticket key into a repository. The key must have been exported from a repository using the EXPORT\_TICKET\_KEY method. When you import the key, the password argument in IMPORT\_TICKET\_KEY must be the same password that you used in the EXPORT\_TICKET\_KEY method.

Keys are typically exported from one repository and imported into another as part of the configuration process when setting up trusted repositories, to allow use of global login tickets or tokens. (For a description of trusted repositories, global login tickets, and global tokens, refer to [Chapter 2, Session and Transaction Management](#), in *Content Server Fundamentals*.)

## Related administration methods

[EXPORT\\_TICKET\\_KEY](#), page 213

[RESET\\_TICKET\\_KEY](#), page 298

IDfSession.importTicketKey()

## Example

```
EXECUTE import_ticket_key
WITH KEY_STRING='ticket_string',
PASSWORD='myword'

status=dmAPIGet("apply,s0,IMPORT_TICKET_KEY,
KEY_STRING,S,ticket_string,PASSWORD,S,myword")
```

# LIST\_AUTH\_PLUGINS

**Purpose** Lists the authentication plugins that the server has loaded.

## Syntax

```
EXECUTE list_auth_plugins  
dmAPIGet ("apply, session, NULL, LIST_AUTH_PLUGINS")
```

## Arguments

None

## Return value

The method returns a collection with one query result object. The object has two repeating string attributes, `plugin_id` and `plugin_filepath`. The values in `plugin_id` are the unique plugin identifiers and the values in `plugin_filepath` are the full paths of the plugins. The values at corresponding index positions represent one plugin.

## Permissions

You must have Sysadmin or Superuser privileges to execute this method.

## General notes

This method is useful only at sites that have a Trusted Content Services license because the ability to use authentication plugins is a feature of Trusted Content Services.

## Related administration methods

None

## Example

The following IAPI excerpt shows how this method is executed and the results obtained.

```
API>apply, c, NULL, LIST_AUTH_PLUGINS
...
q0
API>next, c, Q0
...
OK
API>dump, c, q0
...
USER_ATTRIBUTES
plugin_id[0]: customauth
[1]: custauth2
[2]: latin1auth
plugin_filepath[0]: C:\Documentum\product\5.2\bin\
    auth\customauth.dll
[1]: C:\Documentum\product\5.2\bin\auth\custauth2.dll
[2]: C:\Documentum\product\5.2\bin\auth\latin1auth2.dll

SYSTEM_ATTRIBUTES
APPLICATION_ATTRIBUTES

INTERNAL_ATTRIBUTES
API>close, c, q0
```

# LIST\_RESOURCES

**Purpose** Lists a variety of information about the server's operating system environment and the server.

## Syntax

```
EXECUTE list_resources [WITH reset=true|false]
dmAPIGet ("apply, session, NULL, LIST_RESOURCES
[RESET, B, T|F]")
```

## Arguments

**Table 3-24. LIST\_RESOURCES arguments**

Argument	Datatype	Value	Description
RESET	B	T (TRUE) or F (FALSE)	TRUE reinitializes the file handle and heap size counters. FALSE keeps the present values of the counters.

## Return value

LIST\_RESOURCES returns a collection with one query result object. On Windows platforms, the query result object has twelve attributes, described in [Table 3-25, page 240](#).

**Table 3-25. Collection attributes for LIST\_RESOURCES**

Attribute	Datatype	Description
file_handles_in_use	integer	The number of file handles in use by the main thread and all session threads (Windows) or the current child process (UNIX).
file_handles_max	integer	The configured limit at the operating-system level on the number of file handles the process can open.



Attribute	Datatype	Description
file_handles_new	integer	A counter that indicates how many file handles have been created or destroyed since the last LIST_RESOURCES with RESET = T. If the number is negative, it means that there are fewer handles open than there were at the last LIST_RESOURCES call. (Issuing LIST_RESOURCES with RESET=T reinitializes file_handles_new to zero.)
session_heap_size_max	integer	Maximum size, in bytes, of a thread's session heap. When a session is started, its maximum heap size corresponds to this value. A value of -1 indicates that the size of the session heap will be unconstrained (the heap will grow to whatever size the server machine resources will allow).
current_heap_size_max	integer	Maximum size of the thread's session heap. This reflects the value that was in session_heap_size_max when the session was started, and is the size of the heap available to the session.
session_heap_size_in_use	integer	How much, in bytes, of the currently allocated heap (virtual memory) is in use by the session.
session_heap_size_new	integer	A count of the bytes that the heap has grown or shrunk since the last LIST_RESOURCES call. (Issuing LIST_RESOURCES with RESET=T reinitializes session_heap_size_new to zero.)
root_heap_size_in_use	integer	How much, in bytes, of the main server thread's heap is in use.

Attribute	Datatype	Description
root_heap_size_new	integer	A count of the bytes that the heap has grown or shrunk since the last LIST_RESOURCES call.
		Issuing LIST_RESOURCES with RESET=T reinitializes session_heap_size_new to zero.
max_processes	integer	The maximum number of processes that can be created by the account under which the server is running.
server_init_file	string(255)	The full path to the server's server.ini file.
initial_working_directory	string(255)	The full path to the directory containing the server executable.

On UNIX, the result object has eight attributes, described in [Table 3–26, page 242](#).

**Table 3-26. Collection attributes for LIST\_RESOURCES**

Attribute	Datatype	Description
file_handles_in_use	integer	The number of file handles in use by the main thread and all session threads (Windows) or the current child process (UNIX).
file_handles_max	integer	The configured limit at the operating-system level on the number of file handles the process can open.
file_handles_new	integer	A counter that indicates how many file handles have been created or destroyed since the last LIST_RESOURCES with RESET = T. If the number is negative, it means that there are fewer handles open than there were at the last LIST_RESOURCES call. (Issuing LIST_RESOURCES with RESET=T reinitializes file_handles_new to zero.)
heap_size_in_use	integer	The size, in bytes, of the session heap.

Attribute	Datatype	Description
heap_size_new	integer	A count of the bytes that the heap has grown or shrunk since the last LIST_RESOURCES call.  Issuing LIST_RESOURCES with RESET=T reinitializes heap_size_new to zero.
max_processes	integer	The maximum number of processes that can be created by the account under which the server is running.
server_init_file	string(255)	The full path to the server's server.ini file.
initial_working_directory	string(255)	The full path to the directory containing the server executable.

## Permissions

Anyone can use this method.

## General notes

LIST\_RESOURCES lists a variety of information about the server operating system environment and the server (the information is described in [Return value, page 240](#)).

## Related administration methods

[DB\\_STATS, page 191](#)  
[LIST\\_SESSIONS, page 244](#)  
[LIST\\_TARGETS, page 248](#)  
[SHOW\\_SESSIONS, page 319](#)

## Examples

The following examples return the resources information and reset the heap size counters:

```
EXECUTE list_resources WITH reset=true
dmAPIGet ("apply,c,NULL,LIST_RESOURCES,RESET,B,T")
```

# LIST\_SESSIONS

**Purpose** Lists information about all the currently active sessions and a user-specified number of historical sessions.

## Syntax

```
EXECUTE list_sessions[WITH brief_info=true|false]
dmAPIGet ("apply, session, NULL, LIST_SESSIONS
[, BRIEF_INFO, B, T|F] ")
```

## Arguments

**Table 3-27. LIST\_SESSIONS arguments**

Argument	Datatype	Value	Description
BRIEF_INFO	B	T (TRUE) or F (FALSE)	Indicates whether you want a subset of the attributes in the result object. The default is FALSE.

## Return value

LIST\_SESSIONS returns a collection. Each result object in the collection describes one currently active session or one historical session.

### If BRIEF\_INFO is FALSE

[Table 3-28, page 244](#) lists the information returned if BRIEF\_INFO is FALSE.

**Table 3-28. Complete information returned by LIST\_SESSIONS (BRIEF\_INFO=F)**

Attribute	Description
root_pid	On Windows, the process ID of the Content Server process.  On UNIX, the process ID of the root Content Server process.

Attribute	Description
shared_mem_id	The ID of the shared memory segment used by the servers.
semaphore_id	This attribute is returned only on UNIX platforms. The ID of the semaphore used by the servers.
root_start	This attribute is returned only on UNIX platforms. The time when the main server thread (root server process) was started.
session	The object ID of the session begun by user_name.
db_session_id	The ID of the database session.
pid	The ID of the session thread (process).
user_name	The user name of the user who started the session.
user_authentication	The user authentication state for the session. Possible values are password, ticket, trusted client, change password, or in progress.  If the value is change password, the user logged in for that session can only perform the change password operation. No other operations are allowed.
client_host	The host name of the machine on which the session was started.
client_lib_ver	The version number of the dmcl in use for the session.
client_locale	A verbose description of the client's environment, including the operating system version, the character set in use, the language in use, and the date format.
start	The starting time of the session.
last_used	The last time the client session contacted the server.

Attribute	Description
session_status	The session status. Active means that the session is connected and has not timed out. Inactive means that the session has timed out.
shutdown_flag	Records the setting of the immediacy_level argument in the Kill method.

### If BRIEF\_INFO is TRUE

Table 3–29, page 246, lists the information returned if BRIEF\_INFO is TRUE.

**Table 3-29. Brief information returned by LIST\_SESSIONS (BRIEF\_INFO=T)**

Attribute	Description
session	The object ID of the session begun by user_name.
db_session_id	The ID of the database session.
user_name	The user name of the user who started the session.
client_host	The host name of the machine on which the session was started.
start	The starting time of the session.
last_used	The last time the client session contacted the server.
session_status	The session status. Active means that the session is connected and has not timed out. Inactive means that the session has timed out.

### Permissions

Anyone can use this method.

### General notes

LIST\_SESSIONS returns a collection of query result objects whose attributes contain information about all the currently active sessions and a user-specified number of historical sessions. The historical sessions are past sessions that are not currently active.

The maximum number of historical sessions returned by LIST\_SESSIONS is determined by the parameter `history_sessions` in the server's startup file (the `server.ini` file). There is also a startup parameter, called `history_cutoff`, to define a cutoff time for the history sessions. For example, if `history_cutoff` is set to 15 minutes, then LIST\_SESSIONS will not return any historical sessions older than 15 minutes.

For a description of the `server.ini` file and the parameters you can set in it, refer to [The server.ini file, page 88](#), in the *Content Server Administrator's Guide*.

If you have a large number of active sessions, use [SHOW\\_SESSIONS, page 319](#), instead of LIST\_SESSIONS.

## Related administration methods

[SHOW\\_SESSIONS, page 319](#)

## Examples

This example executes LIST\_SESSIONS with the BRIEF\_INFO argument defaulted to FALSE:

```
EXECUTE list_sessions
```

This example executes LIST\_SESSIONS with the BRIEF\_INFO argument set to TRUE:

```
dmAPIGet ("apply, c, NULL, LIST_SESSIONS, BRIEF_INFO, B, T")
```

# LIST\_TARGETS

**Purpose** Lists the connection brokers to which the server is currently projecting.

## Syntax

```
EXECUTE list_targets  
dmAPIGet ("apply, session, NULL, LIST_TARGETS")
```

## Arguments

LIST\_TARGETS has no arguments.

## Return value

LIST\_TARGETS returns a collection with one query result object. The Table 2-9. The values across the attributes at one index position represent the information about one connection broker to which the server is projecting.

**Table 3-30. Query result object attributes for LIST\_TARGETS**

Attribute	Datatype	Description
projection_targets	string(80)	A repeating attribute that contains the names of the hosts on which the target connection brokers are running.
projection_proxval	integer	A repeating attribute that contains the proximity value projected to the connection broker identified in the corresponding index position in projection_targets.
projection_notes	string(80)	A repeating attribute that contains any user-defined note for the connection broker defined at the corresponding index position in projection_targets and projection_proxval.



Attribute	Datatype	Description
docbroker_status	string(64)	A repeating attribute that records whether the connection broker identified in projection_targets at the corresponding index position is available. Valid values are: Available and Unavailable.
comments	string(64)	A repeating attribute that records whether the projection target entry was taken from the server config object or the server.ini file.

## Permissions

Anyone can use this method.

## General notes

A server's connection broker targets are those connection brokers to which the server is sending checkpoint messages. Target connection brokers are defined in the server's server config object and may also be defined in the server.ini file. This method returns a list of the targets defined in the server config object.

## Related administration methods

None

## Examples

The following examples list the current connection broker targets for the server:

```
EXECUTE list_targets
dmAPIGet ("apply,c,NULL,LIST_TARGETS")
```

# LOG\_OFF

**Purpose** Turns off the RPC logging.

## Syntax

```
EXECUTE log_off  
dmAPIGet ("apply, session, NULL, LOG_OFF")
```

## Arguments

LOG\_OFF has no arguments.

## Return value

LOG\_OFF returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

None

## Related administration methods

[LOG\\_ON](#), page 251

## Example

The following example turns off RPC logging:

```
dmAPIGet ("apply, c, NULL, LOG_OFF")
```

# LOG\_ON

**Purpose** Turns on logging for RPC calls.

## Syntax

```
EXECUTE log_on [WITH detail=true|false]
dmAPIGet ("apply, session, NULL, LOG_ON
[, DETAIL, B, T|F] ")
```

## Arguments

Table 3-31. LOG\_OFF arguments

Argument	Datatype	Value	Description
DETAIL	B	T (TRUE) or F (FALSE)	Indicates whether you want information about the arguments passed with the RPC call and the results. The default is FALSE.

## Return value

LOG\_ON returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

The LOG\_ON function turns on logging for RPC calls.

If you execute the LOG\_ON function without including the DETAIL argument, the server logs information about the start and stop names and time information for the

operation. If you set `DETAIL` to `TRUE`, the server also includes information about the arguments passed with each call and the results of the call.

## Related administration methods

[LOG\\_OFF](#), page 250

## Examples

These examples turn on detailed RPC logging:

```
EXECUTE log_on WITH detail=true  
dmAPIGet ("apply, c, NULL, LOG_ON, DETAIL, B, T")
```

# MAKE\_INDEX

**Purpose** Creates an index for a persistent object type.

## Syntax

```
EXECUTE make_index
WITH type_name='object_type',attribute='attribute_name',
{attribute='attribute_name',}
[unique=true|false,]index_space='name'

dmAPIGet ("apply, session, NULL, MAKE_INDEX
TYPE_NAME, S, object_type, ATTRIBUTE, S, attribute_name,
{ATTRIBUTE, S, attribute_name,}
[UNIQUE, B, T|F, ] INDEX_SPACE, S, name")
```

## Arguments

**Table 3-32. MAKE\_INDEX arguments**

Argument	Datatype	Value	Description
TYPE_NAME	S	<i>object_type</i>	Identifies the object type for which to create an index.  This is a required argument.
ATTRIBUTE	S	<i>attribute_name</i>	Identifies the attribute or attributes on which to build the index. You can specify multiple attributes, but you cannot mix single-valued and repeating attributes in the same statement.

## Return value

MAKE\_INDEX returns a collection with one query result object. The object has one attribute, named result, that contains the object ID of the dmi\_index\_object for the new index if successful.

If the method failed because the syntax was incorrect or the attribute specified did not exist, the result attribute contains F (FALSE).

If the specified index already exists, the result attribute contains '0000000000000000'.

## Permissions

You must have Superuser privileges to use this method.

## General notes

MAKE\_INDEX creates an index for a persistent object type. You can create an index for any persistent object type.

You can specify one or more attributes. However, if you specify multiple attributes, you must specify all single-valued attributes or all repeating attributes. You cannot mix single and repeating valued attributes in the same argument list. (This is because the underlying RDBMS stores an object's single and repeating attributes in separate tables.)

If you specify multiple attributes, the sort order within the index corresponds to the order in which the attributes are specified in the statement. To include the `r_object_id` column in the index, that column must be first column in the index. This means the `r_object_id` must be the first attribute listed in the arguments to MAKE\_INDEX.

The attributes you specify must be defined for the type you specify. They cannot be inherited attributes. For example, if you want to create an index on the subject and title attributes, you would specify `dm_sysobject` in TYPE\_NAME, as these attributes are defined for `dm_sysobject`. ([Chapter 2, Object Reference](#), in the *Object Reference Manual* lists the attributes defined for each object type.)

## Related administration methods

[DROP\\_INDEX](#), page 206  
[FINISH\\_INDEX\\_MOVES](#), page 215  
[MOVE\\_INDEX](#), page 271

## Examples

This example creates an index for the `dm_sysobject` object type, indexing on the subject and title attributes:

```
EXECUTE make_index WITH type_name='dm_sysobject',  
attribute='subject', attribute='title'
```

The next example creates an index for the `dm_sysobject` type on the attribute `r_creator_name`:

```
dmAPIGet ("apply, c, NULL, MAKE_INDEX,
```

```
TYPE_NAME,S, dm_sysobject,  
ATTRIBUTE_NAME,S, r_creator_name")
```

The next example creates an index for the dm\_sysobject type on the attributes r\_creator\_name and r\_creation\_date:

```
dmAPIGet ("apply, c, NULL, MAKE_INDEX,  
TYPE_NAME,S, dm_sysobject,  
ATTRIBUTE_NAME,S, r_creator_name,  
ATTRIBUTE_NAME,S, r_creation_date")
```

# MARK\_AS\_ARCHIVED

**Purpose** Sets the `i_is_archived` attribute of an audit trail entry to TRUE.

## Syntax

```
EXECUTE mark_as_archived FOR audit_obj_id  
  
dmAPIExec("apply,session,audit_obj_id,MARK_AS_ARCHIVED")  
  
audit_obj_id is the object ID of the audit trail entry. This is a dm_audittrail,  
dm_audittrail_acl, or dm_audittrail_group object.
```

## Arguments

This method has no arguments.

## Return value

The method returns TRUE if successful or FALSE if unsuccessful.

## Permissions

You must have Superuser privileges to execute this method.

## General notes

Use this method to set an audit trail entry's `i_is_archived` attribute to TRUE after you archive the entry.

## Related administration methods

[PURGE\\_AUDIT](#), page 275

## Examples

```
EXECUTE mark_as_archived FOR 5f000012ae6217ce.
```



```
dmAPIExec ("apply, S0, 5f000012ae6217ce, MARK_AS_ARCHIVED")
```

# MARK\_FOR\_RETRY

**Purpose** Finds queue items representing events queued to the Index Agent that are in the acquired or failed state and resets them to the pending state

## Syntax

```
EXECUTE mark_for_retry  
WITH NAME = 'index_name'  
  
dmAPIGet("apply, session, NULL, MARK_FOR_RETRY,  
NAME, S, index_name")
```

## Arguments

Table 3-33. MARK\_FOR\_RETRY arguments

Argument	Datatype	Value	Description
NAME	S	<i>index_name</i>	Identifies the index that contains the objects to mark for a retry. Use the name associated with the index's fulltext index object.

## Return value

MARK\_FOR\_RETRY returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

Use MARK\_FOR\_RETRY in the recovery procedure if you encounter problems with the Index Agent. The method scans the repository to find queue items representing events

queued to the Index Agent that are in the acquired or failed state and resets them to the pending state.

The Index Agent registers itself for certain events on all SysObjects. The events serve as notice to the Index Agent that the objects need to be indexed or updated in the index. If there is a problem with the Index Agent, for example, if it crashes, some of the queue items representing those events may be left in the acquired or failed state. MARK\_FOR\_RETRY is used to reset those events to pending so that the Index Agent will pick up those events and properly process the objects after the problem is resolved.

## Related administration methods

None

## Examples

These examples find all the content objects in the index that have an update\_count value of -5 and marks them for a retry:

```
EXECUTE mark_for_retry
WITH NAME = 'storagel_index'

dmAPIGet ("apply, c, NULL, MARK_FOR_RETRY,
NAME, S, storagel_index")
```

# MIGRATE\_CONTENT

**Purpose** Moves content files from one storage area to another.

## Syntax

To migrate a single object:

```
EXECUTE migrate_content [FOR] object_id
WITH target_store='target_storage_name'
[, renditions=value] [, remove_original=T|F] [, log_file='log_file_path']
```

```
dmAPIGet ("apply, session, object_id, MIGRATE_CONTENT,
TARGET_STORE, S, target_storage_name [, RENDITIONS, S, value] [, REMOVE_ORIGINAL, B, T|F]
[, LOG_FILE, S, log_file_path]")
```

To migrate all objects in a particular storage area:

```
EXECUTE migrate_content WITH source_store='source_storage_name',
target_store='target_storage_name', log_file='log_file_path'
[, max_migrate_count=value] [, batch_size=value]
[, remove_original=value]
```

```
dmAPIGet ("apply, session, NULL, MIGRATE_CONTENT,
SOURCE_STORE, S, source_storage_name,
TARGET_STORE, S, target_storage_name, LOG_FILE, S, log_file_path
[, MAX_MIGRATE_COUNT, I, value] [, BATCH_SIZE, I, value]
[, REMOVE_ORIGINAL, B, T|F]
```

To migrate a set of objects identified by a DQL query:

```
EXECUTE migrate_content WITH target_store='target_storage_name', query=
'DQL_predicate' [, sysobject_query=T|F], log_file='log_file_path'
[, renditions=value] [, max_migrate_count=integer] [, batch_size=value]
[, remove_original=value]
```

```
dmAPIGet ("apply, session, NULL, MIGRATE_CONTENT,
TARGET_STORE, S, target_storage_name, QUERY, S, dql_predicate [, SYSOBJECT_QUERY, B, T|F]
[, LOG_FILE, S, log_file_path [, RENDITIONS, S, value] [, MAX_MIGRATE_COUNT, I, value]
[, BATCH_SIZE, I, value] [, REMOVE_ORIGINAL, B, T|F]
```

## Arguments

**Table 3-34. MIGRATE\_CONTENT arguments**

Argument	Datatype	Value	Description
<i>object_id</i>	ID	<i>content_object ID</i> or <i>SysObject ID</i>	<p>When migrating a single content file, you can identify the file using its content object ID or the object ID of the SysObject (or SysObject subtype) that contains the content file.</p> <p>Specifying a content object ID requires Superuser privileges. Specifying a SysObject object ID requires Write permission on the object.</p>
SOURCE_STORE	S	<i>source_storage_name</i>	<p>Name of the storage area whose content you wish to migrate. The storage area may be a file store, a ca store, a blobstore, or a distributed store. Use the name of the storage area's object.</p>
TARGET_STORE	S	<i>target_storage_name</i>	<p>Name of the storage area to which you are migrating the content. The storage area may be a file store, a ca store, or a distributed store. Use the name of the storage area's object.</p> <p>(Blobstores cannot be specified as a target storage area.)</p>

Argument	Datatype	Value	Description
RENDITIONS	S	all, primary, or secondary	<p>Indicates which content files associated with the SysObject or SysObjects you wish to move. This argument is used only when <i>object_id</i> references a SysObject or when SYSOBJECT_QUERY is set to T. Valid values are:</p> <p><b>all</b>, meaning move the first primary content file (page 0) and its renditions</p> <p><b>primary</b>, meaning move only the first primary content file (page 0)</p> <p><b>secondary</b>, meaning move only the renditions of the first primary file (page 0)</p> <p>The default is primary.</p>
REMOVE_ORIGINAL	B	T (TRUE) or F (FALSE)	<p>Whether to remove the content file from the source storage area. T directs Content Server to remove the original content file if possible. F directs Content Server not to remove the content.</p> <p>If set to F, you must include the LOG_FILE argument also.</p> <p>The default is T.</p>

Argument	Datatype	Value	Description
LOG_FILE	S	<i>log_file_path</i>	Identifies where to log messages generated by the method.  You must include this argument if you are moving all content from one storage area to another or if you are using a DQL predicate to select content for migration.
MAX_MIGRATE_COUNT	I	<i>value</i>	Defines the maximum number of content files to migrate.
BATCH_SIZE	I	<i>value</i>	Defines how many content files to include in a single transaction during the migration operation.  The default is 500.
QUERY	S	<i>dql_predicate</i>	Specifies a DQL predicate to be used to select content files for migration.  This must be a valid WHERE clause qualification.  The predicate is applied to either content objects or SysObjects, depending on the setting of the SYSOBJECT_QUERY argument.  For more information, refer to the Usage Notes.
SYSOBJECT_QUERY	B	T (TRUE) or F (FALSE)	T directs the method to apply the DQL predicate specified in the QUERY argument to SysObjects. F directs the method to apply the predicate to content objects.  The default is F.

## Return value

The method returns a collection with one query result object. The object has one integer attribute, named `result`, whose value is the number of contents successfully migrated.

## Permissions

Superuser permissions are required to:

- Specify a content object ID when moving a single content file.
- Move all content from one storage area to another.
- Move content using a DQL predicate that operates on content objects (SYSOBJECT\_QUERY=F)

Write permission on the object is required to:

- Specify a SysObject ID when moving the content of a single object.
- Move content using a DQL predicate that operates on a SysObject (SYSOBJECT\_QUERY=T)

## General notes

MIGRATE\_CONTENT is the preferred way to move content files from a file store, ca store, blobstore, or distributed store storage area to a file store, ca store, or distributed store storage area. You cannot move content to a blob store storage area. Additionally, you cannot use this method to move files to or from turbo or external storage. If content is stored in a retention-enabled ca store storage area, you can only move the content to another retention-enabled ca store storage area. If you move a file to a distributed storage area, the file is placed in the distributed component local to the Content Server to which you are connected when executing the method.

The method moves content files associated with immutable objects as well as changeable objects. Additionally, it does not update a SysObject's `r_modify_date` attribute when the object's content is moved. You must use this method if you want to move content files from an unencrypted storage area to an encrypted storage area. The method allows you to move one file, all files in a particular storage area, or a set of files selected by a DQL WHERE clause qualification.

The method operates on `dmr_content` objects, resetting their `storage_id` attribute to the new storage area and resetting the data ticket attribute. It also updates the `i_vstamp` attribute. The `i_vstamp` attribute is also incremented for any SysObject that contains the content file as the first primary content (page 0). (Consequently, if the affected content objects are associated with replicated SysObjects, the next replication job for those objects will refresh the objects.) Similarly, if the content objects are associated with persistently cached objects, the next time the cached objects are accessed by the client, they will be refreshed.

Here are some general guidelines for using MIGRATE\_CONTENT:



- If you want to move content to or from a distributed storage area, you must specify the `dm_distributedstore` object as the target or source storage area—do not specify the actual distributed component.
- Make sure none of the objects whose content you intend to migrate are checked out. If the method attempts to migrate content associated with a checked out object, the method fails with an error.
- If the content to be migrated is stored in a content-addressed storage area with a retention date, make sure the retention date has expired. If not, the method fails with an error.
- If the content's current storage area and the target storage area are the same content-addressed storage area and the storage area has a retention period, the method updates the content's metadata and retention period and creates a new address for the content. However, the content is not physically moved.
- Back up the repository and the file store storage areas before using the method.
- Make sure that the target storage area has enough disk space to hold the migrated content.
- If you choose not to remove the original content, that content becomes an orphaned content file and can be removed using `dmfilescan`.

### Migrating to retention-enabled content-addressed storage

If you are moving content files to a content-addressed storage area that requires a retention date, you must set the retention date for the content before migrating the content. Use either the `Setcontentattrs` method or the `SET_CONTENT_ATTRS` administration method to set the retention date.

**Note:** A content-addressed storage area requires a retention date if:

- The storage area's `a_retention_attr_name` attribute is set and the `a_default_retention_date` is not null, or
- The `a_retention_attr_name` attribute is set and the `a_retention_attr_required` attribute is set to T

### Referencing SysObjects

When you execute the method against `SysObjects`, the method acts only on the first primary content page (page 0) of each object, the renditions of the primary pages, or both. If a `SysObject` has multiple primary content pages, the primary pages other than page 0 and their renditions are not affected by the method. (If you want to move all the content, use a method syntax that references the content objects, rather than the `SysObject`.)

You can specify which of these files you want the method to act on by setting the `RENDITIONS` argument. The default for `RENDITIONS` is `primary`, which means that the method will only operate on the first primary content files of the `SysObjects`. If you set the argument to `all`, the method moves both the primary content page and its renditions. If you set the argument to `secondary`, the method moves only the renditions but not the primary content.

If you want to execute a DQL predicate against SysObjects, you must set the `SYSOBJECT_QUERY` argument to T (TRUE). The method affects content of only those objects that match the predicate and for which you have at least Write permission.

## Controlling the size of the operation

If you are moving a large number of files, you can control the operation using the `MAX_MIGRATE_COUNT` and `BATCH_SIZE` arguments. `MAX_MIGRATE_COUNT` controls how many content files are moved with each execution of the method. Use this argument if you have a large number of files to migrate and want to perform the migration in smaller steps rather than all in one operation. `BATCH_SIZE` controls how many content files are moved in a single transaction within the migration operation. Setting `BATCH_SIZE` can help protect the operation from overrunning system and database resource limits during the operation.

If you set `MAX_MIGRATE_COUNT`, make sure that each execution of the method does not attempt to move content files that have already been migrated in previous executions of the method.

## Including a query

If you include a DQL predicate, the query built from the predicate is applied to `dmr_content` or `dm_sysobject` objects, depending on the value of the `SYSOBJECT_QUERY` argument. `SYSOBJECT_QUERY` is F (FALSE) by default, meaning that the predicate is applied to content objects. You must have Superuser permissions to run the predicate against content objects.

When you run the query against content objects, it is recommended that you use a predicate that references a unique value in the content object. For example, reference the storage area ID—there is only one storage area in the repository with any given storage ID.

If you set `SYSOBJECT_QUERY` to T (TRUE), the predicate is executed against `dm_sysobjects`. When the query is executed, the method affects the content of only those objects that match the query and for which you have at least Write permission. (For more information about running `MIGRATE_CONTENT` against SysObjects, refer to [Referencing SysObjects, page 265.](#))

## Removing the original content

When you move content, you can choose whether to remove the content from the current (source) storage area or not.

If you choose not to remove the content from the current storage area and no other content objects reference that file, Content Server writes a message to the log file to identify the content file as an orphaned file that must be removed manually. (The file is now an orphan because it is no longer referenced by any content objects.)

If you choose to remove the content from the current storage area and the storage area is a file store storage area, Content Server first checks to ensure that no other content objects reference that content. If another content object is found that references the content, the content is not removed. (Multiple content objects may reference one content file if content duplication checking and prevention is enabled for the file store storage area. For information about that feature, refer to [Content duplication checking and prevention, page 224](#), in the *Content Server Administrator's Guide*.)

## Log file use

You must identify a log file location if you are migrating an entire storage area or including a DQL predicate or setting REMOVE\_ORIGINAL to false in the method arguments. Specifying a log file location is only optional if you are migrating just a single object. For all other cases, you must specify a log file location. If the file you specify already exists, the method appends to that file.

The method logs the following messages:

- A success message for each content object successfully migrated without errors.
- A failure message for each content object that was not successfully migrated. The message includes the reason for the failure.
- Messages for all database errors. (These are written to the server log and the session log file.)
- A success or failure message for each batch in the operation.

**Note:** If you are migrating a single content file, you can retrieve messages using the GetMessage method.

## Related administration methods

None

## Examples

These two examples migrate a single content file, represented by the content object 0600000272ac100d:

```
EXECUTE migrate_content FOR 0600000272ac100d
WITH target_store='filestore_02'
```

```
dmAPIGet ("apply,s0,0600000272ac100d,MIGRATE_CONTENT,  
TARGET_STORE,S,filestore_02")
```

The next two examples migrate all the content from filestore\_01 to engr\_filestore:

```
EXECUTE migrate_content  
WITH source_store='filestore_01',  
target_store='engr_filestore',batch_size=100,  
log_file='C:\temp\migration.log'
```

```
dmAPIGet ("apply,s0,NULL,MIGRATE_CONTENT,  
SOURCE_STORE,S,filestore_01,  
TARGET_STORE,S,engr_filestore,BATCH_SIZE,I,100,  
LOG_FILE,S,C:\temp\migration.log
```

The final two examples use a DQL predicate to select the content to migrate. The query uses content size to select the content.

```
EXECUTE migrate_content  
WITH target_store='engr_filestore',  
query='content_size>1000',max_migrate_count=1000,  
batch_size=100,log_file='C:\temp\migration.log'
```

```
dmAPIGet ("apply,s0,NULL,MIGRATE_CONTENT,  
TARGET_STORE,S,engr_filestore,  
QUERY,S,content_size>1000,MAX_MIGRATE_COUNT,I,1000  
BATCH_SIZE,I,100,LOG_FILE,S,C:\temp\migration.log
```

# MODIFY\_TRACE

**Purpose** Turns tracing on and off for full-text index query operations.

## Syntax

```
EXECUTE modify_trace
WITH subsystem='fulltext',value='tracing_level'

dmAPIGet ("apply, session, NULL, MODIFY_TRACE
SUBSYSTEM, S, fulltext, VALUE, S, tracing_level
```

## Arguments

**Table 3-35. MODIFY\_TRACE arguments**

Argument	Datatype	Value	Description
SUBSYSTEM	S	fulltext	The keyword fulltext indicates that you want to turn on tracing for the full-text querying operations.
VALUE	S	<i>tracing_level</i>	The tracing_level must be one of:  <i>none</i> , to turn off tracing  <i>all</i> , to log both Content Server and full-text querying messages

## Return value

MODIFY\_TRACE returns a collection with one query result object. The object has one Boolean attribute that indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

MODIFY\_TRACE turns tracing on and off for full-text index operations. When tracing is on, tracing information for all full-text queries is logged. The trace information is placed in the server's log file.

The tracing information includes informational, verbose, and debug messages. All trace messages include time stamps and process and thread ID information.

If the method returns FALSE, the level of tracing remains unchanged. For example, if tracing is currently set at all and you execute MODIFY\_TRACE to reset the level to none, but the method returned FALSE, the trace level remains at the all level.

## Related administration methods

[LOG\\_OFF](#), page 250

[LOG\\_ON](#), page 251

[SET\\_OPTIONS](#), page 312

## Examples

The following examples turn on full tracing:

```
EXECUTE modify_trace  
WITH subsystem = 'fulltext', value = 'all'  
  
dmAPIGet ("apply,c,NULL,MODIFY_TRACE,  
SUBSYSTEM,S,fulltext,VALUE,S,all")
```

# MOVE\_INDEX

**Purpose** Moves an existing object type index from one tablespace or segment to another. (This method is not supported for servers running against DB2.)

## Syntax

```
EXECUTE move_index FOR 'dmi_index_obj_id'
WITH name = 'new_home'

dmAPIGet ("apply, session, dmi_index_obj_id, MOVE_INDEX,
NAME, S, new_home")
```

## Arguments

**Table 3-36. MOVE\_INDEX arguments**

Argument	Datatype	Value	Description
NAME	S	<i>new_home</i>	Identifies the new tablespace or segment for the index. Use the name of the tablespace or segment.

## Return value

MOVE\_INDEX returns a collection with one query result object. The object contains one Boolean attribute that indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

MOVE\_INDEX moves an existing object type index from one tablespace or segment to another. You can obtain the index's object ID from the index's index object. Refer to [Chapter 2, Object Reference](#), of the *Object Reference Manual* for a list of the attributes defined for the index type.

MOVE\_INDEX only operates on object type indexes that are represented in the repository by dmi\_index objects.

**Note:** Those indexes that do not have a dmi\_index object (DMI\_OBJECT\_TPE\_UNIQUE and DM\_FED\_LOG\_INDEX for example) may be moved manually. However, be sure to shut down Content Server before moving them.

## Related administration methods

[DROP\\_INDEX](#), page 206  
[FINISH\\_INDEX\\_MOVES](#), page 215  
[MAKE\\_INDEX](#), page 253

## Examples

These examples move the index represented by 1f0000012643124c to the index2 tablespace:

```
EXECUTE move_index FOR '1f0000012643124c'  
WITH name = 'index2'  
  
dmAPIGet ("apply,s0,1f0000012643124c,MOVE_INDEX,  
NAME,S,index2")
```



# PING

**Purpose** Determines if the client still has an active server connection.

## Syntax

```
dmAPIGet ("apply, c, NULL, PING[, RETRY_ON_TIMEOUT, B, T|F] ")
```

You cannot use the EXECUTE statement to invoke PING.

## Arguments

**Table 3-37. PING arguments**

Argument	Datatype	Value	Description
RETRY_ON_TIMEOUT	B	T (TRUE) or F (FALSE)	T (TRUE) forces a connection attempt between the client and the server. The default is F.

## Return value

PING returns a collection with one query result object. The object has one attribute, called result, that is TRUE if the connection is alive. If the connection has timed out, a null collection is returned.

## Permissions

Anyone can use this method.

## General notes

None

## Related administration methods

None

## Example

```
dmAPIGet ("apply, c, NULL, PING, RETRY_ON_TIMEOUT, B, T")
```

# PURGE\_AUDIT

**Purpose** Removes an audit trail entry from the repository.

## Syntax

```
EXECUTE purge_audit WITH delete_mode='mode'
[,date_start='start_date',date_end='end_date']
[,id_start='start_id',id_end='end_id']
[,object_id='object_id'][,dql_predicate='predicate']
[,purge_non_archived=T|F] [,purge_signed=T|F] [,commit_size=value]

dmAPIGet ("apply, session, NULL, PURGE_AUDIT, DELETE_MODE, S, mode [, DATE_
START, S, start_date, DATE_END, S, end_date] [, ID_START, S, start_
id, ID_END, S, end_id] [, OBJECT_ID, S, object_id] [, DQL_
PREDICATE, S, predicate] [, PURGE_NON_ARCHIVED, B, T|F] [, PURGE_
SIGNED, B, T|F] [, COMMIT_SIZE, I, value]
```

## Arguments

**Table 3-38. PURGE\_AUDIT arguments**

Argument	Datatype	Value	Description
DELETE_MODE	S	<i>mode</i>	<p>Defines how the entries to be deleted are chosen. Valid values are:</p> <ul style="list-style-type: none"> <li>• <b>DATE_RANGE</b> Deletes all audit trail entries generated within a specified date range. If you include this, include the DATE_START and DATE_END arguments.</li> <li>• <b>ID_RANGE</b> Deletes all audit trail entries whose object IDs fall within a specified range. If you include this, include the ID_START and ID_END arguments.</li> </ul>

Argument	Datatype	Value	Description
			<ul style="list-style-type: none"> <li>• <b>ALL_VERSIONS</b> Deletes all audit trail entries whose audited_obj_id value corresponds to a specified object ID. ALL_VERSIONS is valid only for audit trail entries whose audited_obj_id identifies a SysObject or SysObject subtype. If you include ALL_VERSIONS, include the OBJECT_ID argument.</li> <li>• <b>SINGLE_VERSION</b> Deletes all audit trail entries whose audited_obj_id value corresponds to a specified object ID. If you include SINGLE_VERSION, include the OBJECT_ID argument.</li> <li>• <b>AUDIT_RECORD</b> Deletes the audit trail entry whose object ID corresponds to a specified object ID. If you include AUDIT_RECORD, include the OBJECT_ID argument.</li> </ul>
DELETE_MODE	S	<i>mode</i>	<ul style="list-style-type: none"> <li>• <b>PREDICATE</b> Deletes all audit trail entries that satisfy a DQL predicate. If you include PREDICATE, include the DQL_PREDICATE argument.</li> </ul>

Argument	Datatype	Value	Description
DATE_START	S	<i>start_date</i>	<p>The starting date of the entries to be deleted. The date is compared to the value in the <code>time_stamp</code> attribute in the audit trail entry. (<code>time_stamp</code> records the local time at which the entry was generated.)</p> <p>The date format can be any acceptable format that does not require a pattern specification. (Refer to <a href="#">Date literals, page 15</a>, for a list of valid formats.</p> <p>If you include <code>DATE_START</code>, you must include <code>END_DATE</code> also.</p> <p>Include a start and end date only when <code>DELETE_MODE</code> is <code>DATE_RANGE</code>.</p>
DATE_END	S	<i>end_date</i>	<p>The ending date of the entries to be deleted. The date is compared to the value in the <code>time_stamp</code> attribute in the audit trail entry. (<code>time_stamp</code> records the local time at which the entry was generated.)</p> <p>The date format can be any acceptable format that does not require a pattern specification. (Refer to <a href="#">Date literals, page 15</a>, for a list of valid formats.</p> <p>If you include <code>END_DATE</code>, you must also <code>START_DATE</code>. The end date must be greater than the start date.</p> <p>Include a start and end date only when <code>DELETE_MODE</code> is <code>DATE_RANGE</code>.</p>
ID_START	S	<i>start_id</i>	<p>The object ID of an audit trail entry.</p>

Argument	Datatype	Value	Description
ID_END	S	<i>end_id</i>	<p>Include ID_START if the DELETE_MODE is ID_RANGE. You must also include ID_END.</p> <p>The object ID of an audit trail entry.</p>
OBJECT_ID	S	<i>object_id</i>	<p>Include ID_END if the DELETE_MODE is ID_RANGE. You must also include ID_START. The ID_END object ID must be larger than the object ID identified in ID_START</p> <p>Include this argument if DELETE_MODE is ALL_VERSIONS, SINGLE_VERSION, SINGLE_ID, or AUDIT_RECORD.</p> <p>For all modes except AUDIT_RECORD, <i>object_id</i> is the object ID recorded in the audited_obj_id attribute of the audit trail entries.</p> <p>For AUDIT_RECORD mode, <i>object_id</i> is the object ID of an audit trail entry.</p>
PURGE_NON_ARCHIVED	B	T (TRUE) or F (FALSE)	<p>Determines whether unarchived audit trial entries are deleted. Setting this argument to T (TRUE) directs Content Server to delete audit trail entries whose i_is_archived attribute is set to F (FALSE).</p> <p>The default for this argument is F (FALSE), meaning that only entries whose i_is_archived attribute is set to T (TRUE) are removed.</p>

Argument	Datatype	Value	Description
PURGE_SIGNED	B	T (TRUE) or F (FALSE)	<p>You cannot set this to T if the DQL_PREDICATE argument is included.</p> <p>Determines whether audit trail entries for dm_addsignature and dm_addesignature events are considered for deletion.</p> <p>Setting this argument to T (TRUE) means entries for dm_addsignature and dm_addesignature events are considered for deletion.</p> <p>The default is F (FALSE), meaning the dm_addsignature and dm_addesignature entries are not considered for deletion.</p>
DQL_PREDICATE	S	<i>predicate</i>	<p>You cannot set this to T if the DQL_PREDICATE argument is included.</p> <p>Defines a DQL predicate to choose the audit trail entries to be deleted. A valid predicate is that portion of a DQL SELECT statement that occurs after the FROM keyword.</p>
COMMIT_SIZE	I	<i>value</i>	<p>If you include this argument, you may not include PURGE_NON_ARCHIVED or PURGE_SIGNED.</p> <p>Defines how many audit trail entries to delete in each transaction within the overall operation. The default is 1000.</p> <p>Setting this to 0 means that all entries are deleted in one transaction.</p>

## Return value

The PURGE\_AUDIT method returns a collection with one query result object. The query result object has two attributes, result and deleted\_objects. The result attribute is set to T if the method completed successfully and F if the method did not complete successfully. deleted\_objects records the number of audit trail entries that were deleted.

## Permissions

You must have Purge Audit privileges to execute this method.

## General notes

Executing the PURGE\_AUDIT method always generates at least one audit trail entry with the event name dm\_purgeaudit. The operation generates one audit trail entry for each transaction within the operation. For example, if you set COMMIT\_SIZE to 100 and the operation deletes 700 audit trail entries, the operation generates 7 audit trail entries, one for each transaction.

The entry for each transaction has the event name dm\_purgeaudit. The optional attributes record the following information for the transaction:

- string\_1 stores the time\_stamp value of the first audit trail entry deleted in the transaction
- string\_2 stores the time\_stamp value of the last audit trail entry deleted in the transaction
- string\_3 stores the actual number of audit trail entries deleted by the transaction
- string\_5 stores the entire list of arguments from the method's command line
- id\_1 stores the object ID of the first audit trail object deleted in the transaction
- id\_2 stores the object ID of the last audit trail object deleted in the transaction

## Related administration mMethods

None

## Examples

This example deletes all archived audit trail entries generated from January 1, 2003 to January 1, 2004:

```
EXECUTE purge_audit WITH DELETE_MODE='DATE_RANGE',  
date_start='01/01/2003 00:00:00 AM',  
date_end='01/01/2004 00:00:00 AM'  
  
dmAPIGet ("apply,c,NULL,PURGE_AUDIT,  
DELETE_MODE,S,DATE_RANGE,
```



```
DATE_START,S,01/01/2003 00:00:00 AM,
DATE_END,S,01/01/2004 00:00:00 AM")
```

This example deletes all audit trail entries generated from January 1, 2003 to January 1, 2004, including unarchived entries. The number of entries deleted in each transaction is set to 500:

```
EXECUTE purge_audit WITH delete_mode='DATE_RANGE',
date_start='01/01/2003 00:00:00 AM',
date_end='01/01/2004 00:00:00 AM'
purge_non_archived=T,commit_size=500
```

```
dmAPIGet ("apply,c,NULL,PURGE_AUDIT,
DELETE_MODE,S,DATE_RANGE,
DATE_START,S,01/01/2003 00:00:00 AM,
DATE_END,S,01/01/2004 00:00:00 AM,
PURGE_NON_ARCHIVED,B,T,COMMIT_SIZE,I,500")
```

This example deletes all archived audit trail entries that identify the document 09000003ac5794ef as the audited object:

```
EXECUTE purge_audit WITH delete_mode='ALL_VERSIONS',
object_id='09000003ac5794ef'
```

```
dmAPIGet ("apply,c,NULL,PURGE_AUDIT,
DELETE_MODE,S,ALL_VERSIONS,
OBJECT_ID,S,09000003ac5794ef")
```

This example deletes the single audit trail entry whose object ID is 5f0000021372ac6f:

```
EXECUTE purge_audit WITH delete_mode='AUDIT_RECORD',
object_id='5f0000021372ac6f'
```

```
dmAPIGet ("apply,c,NULL,PURGE_AUDIT,
DELETE_MODE,S,AUDIT_RECORD,
OBJECT_ID,S,5f0000021372ac6f")
```

This example deletes all audit trail entries whose object IDs range from 5f1e9a8b003a901 to 5f1e9a8b003a925, including unarchived entries:

```
EXECUTE purge_audit WITH delete_mode='ID_RANGE',
id_start='5f1e9a8b003a901',id_end='5f1e9a8b003a925',
purge_non_archived=T
```

```
dmAPIGet ("apply,c,NULL,PURGE_AUDIT,
DELETE_MODE,S,ID_RANGE,ID_START,S,5f1e9a8b003a901,
ID_END,S,5f1e9a8b003a925,PURGE_NON_ARCHIVED,B,T")
```

This example deletes all audit trail entries that satisfy the specified DQL predicate:

```
EXECUTE purge_audit WITH delete_mode='PREDICATE',
dql_predicate=
'dm_audittrail where event_name like 'dcm%'and
r_gen_source=0'
```

If you need to include single-quotes in the predicate string, (for example, 'dcm%'), escape the single-quotes with single-quotes. This is illustrated in the example above.

```
dmAPIGet ("apply,c,NULL,PURGE_AUDIT,
```

```
DELETE_MODE,S,PREDICATE,DQL_PREDICATE,S,dm_audittrail  
where event_name like 'dcm%' and r_gen_source=0")
```

---

# PURGE\_CONTENT

**Purpose** Sets a content file off line and deletes the file from its storage area.

## Syntax

```
EXECUTE purge_content FOR 'content_object_id'  
dmAPIGet ("apply, c, content_obj_id, PURGE_CONTENT")
```

## Arguments

PURGE\_CONTENT has no arguments.

## Return value

PURGE\_CONTENT returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser user privileges to use this method.

## General notes

PURGE\_CONTENT marks a content file as off-line. Additionally, if the file is referenced by only one content object, it deletes the file from the storage area. If the file is referenced by other content objects, the file is not deleted. (A file may be referenced by multiple content objects if the storage area is a file store storage area configured to use content duplication checking and prevention. For information about that feature, refer to [Content duplication checking and prevention, page 224](#), in the *Content Server Administrator's Guide*.)

Do not use this method unless you have previously moved the file to an archival or back up storage area. PURGE\_CONTENT does not back up the file. The method only deletes the file from the storage area and sets the `is_offline` attribute of the file's content object to TRUE.

To use PURGE\_CONTENT, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use PURGE\_CONTENT in the session.

## Examples

These examples delete the content file associated with the content object represented by 060000018745231c and set the is\_offline attribute in the associated content object to TRUE:

```
EXECUTE purge_content FOR '060000018745231c'  
dmAPIGet ("apply,c,060000018745231c,PURGE_CONTENT")
```

# PUSH\_CONTENT\_ATTRS

**Purpose** Updates content metadata in a content addressable systems.

## Syntax

```
EXECUTE push_content_attrs FOR object_id
WITH format=format_name [,page=number]
[,page_modifier=value]

dmAPIGet ("apply,session,object_id,PUSH_CONTENT_ATTRS,
FORMAT,S,format_name[,PAGE,I,number]
[,PAGE_MODIFIER,S,value]")
```

*object\_id* is the ID of a document or other SysObject.

## Arguments

**Table 3-39. PUSH\_CONTENT\_ATTRS arguments**

Argument	Datatype	Value	Description
FORMAT	S	<i>name</i>	The file format of the content file. Use the name of the format object that defines the format.
PAGE	I	<i>number</i>	The page number of the content file in the document. If the content file is a rendition, this is the page number of the primary content with which the rendition is associated.  The default value is 0.
PAGE_MODIFIER	S	<i>value</i>	This identifies a rendition. It is the page modifier defined for the rendition when the rendition was added to the document. The value is stored in the page_modifier attribute in the content object.

## Return value

PUSH\_CONTENT\_ATTRS returns a collection with one query result object. The object has one Boolean attribute that contains TRUE if the method was successful and FALSE if unsuccessful.

## Permissions

To execute this method you must be a superuser and you must have at least Write permission on the object identified in *object\_id*.

## General notes

PUSH\_CONTENT\_ATTRS updates the metadata fields in a content-addressed storage system for a particular content file. First, the method reads the values in the following content-related attributes in the file's associated content object:

- content\_attr\_name
- content\_attr\_value
- content\_attr\_data\_type
- content\_attr\_num\_value
- content\_attr\_date\_value

content\_attr\_name identifies the metadata fields to be set in the storage system. The content\_attr\_data\_type attribute identifies the data type of the value users must provide for each metadata field in the corresponding index position in content\_attr\_name. The actual value is defined in one of the remaining attributes, depending on the field's datatype. For example, if the field name is "title" and its datatype is character string, the content\_attr\_value attribute contains the actual title value. (The remaining attributes, content\_attr\_num\_value and content\_attr\_date\_value are set to the default NULLINT and NULLDATE values.)

After the method reads the content object attributes, it invokes the content-addressed plugin library to update the storage system metadata fields.

**Note:** To be updated, a field must be defined in both the content object's content\_attr\_name attribute and in the storage object's a\_content\_attr\_name attribute. If a field is named in the content object's content\_attr\_name attribute but not defined in the storage object's a\_content\_attr\_name object, it is not set in the storage area. Similarly, if a field is named in the storage object's a\_content\_attr\_name but not in the content object's content\_attr\_name attribute, it is not set in the storage area.

If PUSH\_CONTENT\_ATTRS completes successfully, it generates a new content address for the content. The new address is appended to the i\_contents attribute of the subcontent object that records content addresses for the content file.

For complete information about how to save a document to content-addressed storage, refer to [Creating SysObjects, page 132](#), in *Content Server Fundamentals*.

## Related administration methods

[SET\\_CONTENT\\_ATTRS, page 307](#)

There is also a corresponding DMCL API method, `Setcontentattrs`.

## Examples

```
EXECUTE push_content_attrs FOR 090000026158a4fc
WITH format=txt,page=1

dmAPIGet ("apply,s0,090000026158a4fc,PUSH_CONTENT_ATTRS,
FORMAT,S,txt,PAGE,I,1")
```

# RECOVER\_AUTO\_TASKS

**Purpose** Recovers work items that have been claimed, but not yet processed, by a workflow agent associated with a failed Content Server.

## Syntax

```
EXECUTE recover_auto_tasks  
WITH server_config_name=name  
  
dmAPIGet ("apply, session, NULL, METHOD, S, RECOVER_AUTO_TASKS  
SERVER_CONFIG_NAME, S, name")
```

## Argument

Table 3-40. RECOVER\_AUTO\_TASKS argument

Argument	Datatype	Value	Description
SERVER_CONFIG_NAME	I	<i>name of the server config object</i>	Use the name of the server config object representing the Content Server that crashed or failed.

## Return Value

The method returns T (TRUE) if it completes successfully and F (FALSE) if not.

## Permissions

This method must be executed by a user with Sysadmin or Superuser privileges.

## General Notes

If a Content Server fails, its workflow agent is stopped also. When the server is restarted, the workflow agent will recognize and process any work items it had claimed but not processed before the failure. However, if you cannot restart the Content Server, you must recover those work items already claimed by its associated workflow agent so



that another workflow agent can process them. RECOVER\_AUTO\_TASKS performs that recovery.

Running RECOVER\_AUTO\_TASKS executes a query to update all work items claimed but unprocessed by the failed server's workflow agent. The query resets the a\_wq\_name attribute in the work items to empty. This allows other workflow agents running against the repository to claim those work items for processing.

Before executing this method, make sure that the specified Content Server is not running.

## Related Administration Methods

None

## Example

```
EXECUTE recover_auto_tasks
WITH server_config_name='DevRepository_1'

dmAPIGet ("apply, s0, NULL, METHOD, S, RECOVER_AUTO_TASKS
SERVER_CONFIG_NAME, S, DevRepository_1")
```

# REGISTER\_ASSET

**Purpose** Queues a request to the Media Server to generate a thumbnail, proxies, and metadata for a rich media content file.

## Syntax

```
EXECUTE register_asset FOR object_id
WITH [page=page_number][,priority=priority_level]
dmAPIGet ("apply, session, object_id, REGISTER_ASSET
[, PAGE, I, page_number][, PRIORITY, I, priority_level"]")
```

*object\_id* is the object ID of the document that contains the content file.

## Arguments

**Table 3-41. REGISTER\_ASSET arguments**

Argument	Datatype	Value	Description
PAGE	I	<i>page_number</i>	Page number of the content file. If unspecified, the default is 0.
PRIORITY	I	<i>priority_level</i>	Identifies the priority of the request. A value of zero indicates no priority. Priority rises as the value rises. The Media Server processes higher priority requests before lower priority requests.  If unspecified, the default is 5.

## Return value

REGISTER\_ASSET returns a collection with one query result object. The object has attributes, *result* and *result\_id*. The *result* attribute is a Boolean attribute whose value indicates success (TRUE) or failure (FALSE) of the operation. The *result\_id* attribute records the object ID of the newly created queue item object.

## Permissions

You must have at least Version permission on the object or Sysadmin privileges to use this method.

## General notes

REGISTER\_ASSET is called by Content Server whenever an object with rich media content is saved or checked in to the repository. The method generates an event, `dm_register_event`, that is queued to the Media Server. When the Media Server processes the event, the server creates a thumbnail, proxies, and metadata for the content.

The `dmi_queue_item` that represents the event is queued to the `dm_mediaserver` user, who represents the Media Server. (`dm_mediaserver` is created when Documentum Media Transformation Services is installed.) REGISTER\_ASSET sets the following attributes in the queue item:

**Table 3-42. Queue item attribute values set by REGISTER\_ASSET**

Attribute	Set to
<code>event</code>	<code>dm_register_event</code>
<code>item_id</code>	object ID of the SysObject that triggered the request
<code>instruction_page</code>	page number of the content in the SysObject
<code>date_sent</code>	date the request was made
<code>name</code>	<code>dm_mediaserver</code>
<code>sent_by</code>	session user
<code>priority</code>	priority level identified in the REGISTER_ASSET call

## Related administration methods

[SET\\_CONTENT\\_ATTRS](#), page 307

[TRANSCODE\\_CONTENT](#), page 325

## Example

```
EXECUTE register_asset FOR 090002410063ec21
WITH page=0,priority=8

dmAPIGet ("apply, s0, 090002410063ec21, REGISTER_ASSET,
PAGE, I, 0, PRIORITY, I, 8")
```

# REORGANIZE\_TABLE

**Purpose** Reorganizes a database table for query performance.

## Syntax

```
EXECUTE reorganize_table WITH table_name='name',
[,index='index_name']
```

```
dmAPIGet ("apply, session, NULL, REORGANIZE_TABLE,
TABLE_NAME, S, name[, INDEX, S, index_name]
```

## Arguments

**Table 3-43. REORGANIZE\_TABLE arguments**

Argument	Datatype	Value	Description
TABLE_NAME	S	<i>name</i>	Name of the RDBMS table to be reorganized.
INDEX	S	<i>index_name</i>	Name of an index on the table identified in TABLE_NAME.  This argument is required for an Oracle database table. It is optional for a SQL Server or DB2 database table. It is not used for Sybase database table.  With the exception of Sybase, which doesn't use this argument, the argument is used differently on each database. Refer to the General Notes for details.

## Return value

REORGANIZE\_TABLE returns a collection with one query result object. The object has one attribute, named result, that contains T if the method was successful or F if the method was unsuccessful.

## Permissions

You must be a superuser to execute this method.



**Caution:** Do not use this method unless directed to do so by Technical Support.

## General notes

This method is used by the UpdateStatistics administration tool, in conjunction with UDPATE\_STATISTICS, when the tool is run on a DB2 database. However, the method can be used on any supported database.

### The INDEX argument

On Oracle, you must include the INDEX argument. The method rebuilds the specified index. The index must be an index on the table identified in TABLE\_NAME.

For SQL Server, the argument is optional. If you include it, you must specify an index on the table identified in TABLE\_NAME. The method rebuilds that index. If you do not include the argument, the method rebuilds all indexes on the specified database table.

For DB2, the argument is optional. If you include it, you must specify an index on the table identified in TABLE\_NAME. The method rebuilds the database table based on the ordering defined in the index.



**Caution:** Using the INDEX argument on a DB2 database table is a powerful feature. The choice of index to use as a basis for re-ordering the table can greatly affect query performance against that table— for better or worse, depending on the index choice. It is recommended that if you use this argument, specify the index on the r\_object\_id attribute.

**Note:** Indexes created by Content Server are named using the following format: *D\_index\_obj\_id*, where *index\_obj\_id* is the object ID of the dmi\_index object for the index.

The INDEX argument is not used when the method is run against a Sybase table.

## Related administration methods

[UPDATE\\_STATISTICS](#), page 329

## Examples

These two examples reorganize the `dm_sysobject_s` table. (Note that these two examples do not work on Oracle because the required `INDEX` argument is not included.)

```
EXECUTE reorganize_table
WITH table_name='dm_sysobject_s'

dmAPIGet ("apply,S0,NULL,REORGANIZE_TABLE,
TABLE_NAME,S,dm_sysobject_s")
```

This example rebuilds an index associated with the `dm_sysobject_s` table.

```
EXECUTE reorganize_table
WITH table_name='dm_sysobject_s,
index='D_1f0C9fda80000108'

dmAPIGet ("apply,S0,NULL,REORGANIZE_TABLE,
TABLE_NAME,S,dm_sysobject_s,
INDEX,S,D_1f0C9fda80000108")
```

# REPLICATE

**Purpose** Copies content files in distributed storage areas.

## Syntax

```
EXECUTE replicate
WITH query='value',store='value'
[,type='value']

dmAPIGet ("apply,session,NULL,REPLICATE,
QUERY,S,value,STORE,S,name[,TYPE,S,type_name]")
```

## Arguments

**Table 3-44. REPLICATE arguments**

Argument	Datatype	Value	Description
QUERY	S	<i>dql_predicate expression</i>	This argument is required. The predicate expression is used to build a DQL query that selects the objects whose content you want to copy. The expression can be any expression that would be a valid WHERE clause qualification (refer to <a href="#">The WHERE clause, page 142</a> ).
STORE	S	<i>name</i>	This argument is required. It identifies where to put the new content copy or copies. The name must be the name of a component of a distributed storage area.
TYPE	S	<i>type_name</i>	This argument is optional. It identifies the type of objects whose content you are replicating. <i>type_name</i> must be a direct or indirect subtype of <code>dm_sysobject</code> .

Argument	Datatype	Value	Description
			The default is dm_sysobject.

### Return value

REPLICATE returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

### Permissions

You must have Sysadmin or Superuser privileges to use REPLICATE.

### General notes

REPLICATE copies content files from one component of a distributed storage area to another. Typically, replication of content files is performed by the ContentReplication or surrogate get. Use REPLICATE as a manual backup for these tools. (Refer to [Content Replication, page 477](#), in the *Content Server Administrator's Guide* for information about the ContentReplication and to [Using the Surrogate Get Feature, page 77](#), in the *Distributed Configuration Guide* for information about surrogate get.)

REPLICATE checks the contents stored in the storage area at a specified site and copies back to the local storage area any contents that meet the conditions defined in the function. (Refer to the *Distributed Configuration Guide* for more information.)

To use REPLICATE, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use REPLICATE in the session.

### Related administration methods

- [DELETE\\_REPLICA](#) , page 193
- [IMPORT\\_REPLICA](#), page 234

### Examples

These examples replicate all content for documents of the subtype proposals owned by jenniferk:

```
EXECUTE replicate WITH query='owner_name=jennyk',
store='distcomp_1',type='proposals'
```



```
dmAPIGet ("apply, c, NULL, REPLICATE,  
QUERY, S, owner_name='jennyk',  
STORE, S, distcomp_1, TYPE, S, proposals")
```

The next example replicates the content of all objects whose content is authored by Jane:

```
dmAPIGet ("apply, s0, NULL, REPLICATE,  
QUERY, S, any authors in ('Jane'),  
STORE, S, diststorage3")
```

# RESET\_TICKET\_KEY

**Purpose** Generates a login ticket key and stores it in the repository.

## Syntax

```
EXECUTE reset_ticket_key  
dmAPIGet ("apply, session, RESET_TICKET_KEY")
```

## Arguments

None

## Return value

The method returns a collection with one query result object. The object has one attribute, *result*, that contains T (TRUE) if the method was successful or F (FALSE) if the method was unsuccessful.

## Permissions

You must have Superuser privileges to execute this method.

## General notes

Use this method when you need to replace a login ticket key in a repository with a new key. Any login tickets generated by the repository before the LTK was reset cannot be used within the repository.

## Related administration methods

[EXPORT\\_TICKET\\_KEY](#), page 213

[IMPORT\\_TICKET\\_KEY](#), page 236

IDfSession.resetTicketKey()

**Example**

```
dmAPIGet ("apply, s0, RESET_TICKET_KEY")
```

# RESTORE\_CONTENT

**Purpose** Restores an off-line content file to its original storage area.

## Syntax

```
EXECUTE restore_content FOR 'content_object_id'
WITH file = 'path_name' [,other file = 'other_path_name']

dmAPIGet("apply,session,content_object_id,RESTORE_CONTENT,
FILE,S,file_path[,OTHER_FILE,S,other_path_name]")
```

*content\_obj\_id* is the object ID of the content object associated with the specified file.

## Arguments

**Table 3-45. RESTORE\_CONTENT arguments**

Argument	Datatype	Value	Description
FILE	S	<i>file_path</i>	Specifies the current location of the content file. Use a full path specification.
OTHER_FILE	S	<i>other_path_name</i>	Specifies the current location of the resource fork for the content file. Use a full path specification. Include this argument only if the file is a Macintosh file.

## Return value

RESTORE\_CONTENT returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser user privileges to use this method.

## General notes

RESTORE\_CONTENT restores an off-line content file to its original storage area. This method only operates on one file at a time. To restore multiple files, use the API Restore method.

The method places the file in the storage area indicated by the storage\_id attribute of the file's content object and sets the content object's is\_offline attribute to FALSE.

To use RESTORE\_CONTENT, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use RESTORE\_CONTENT in the session.

## Examples

The following examples restore the file named Myproposal using an EXECUTE statement:

```
EXECUTE restore_content  
WITH file='c:\archive1\Myproposal'
```

or, on UNIX:

```
EXECUTE restore_content  
WITH file='u02/archive1/Myproposal'
```

The next two examples perform the same operation using Apply methods:

```
dmAPIGet ("apply,c,NULL,RESTORE_CONTENT,  
FILE,S,c:\archive1\Myproposal")
```

or, on UNIX:

```
dmAPIGet ("apply,c,NULL,RESTORE_CONTENT,  
FILE,S,u02/archive1/Myproposal")
```

# ROLES\_FOR\_USER

**Purpose** Retrieves the roles assigned to the user in a particular client domain.

## Syntax

```
EXECUTE roles_for_user [[FOR] 'dm_user_id']
WITH [USER_NAME=value][,DOMAIN=domain_name]
```

With Execute, do not include the FOR clause if you include the NAME argument.

```
dmAPIGet("apply,session,dm_user_id,ROLES_FOR_USER,
[USER_NAME,S,value][,DOMAIN,S,domain_name]")
```

With Apply, specify the user ID as NULL if you include the NAME argument.

## Arguments

**Table 3-46. ROLES\_FOR\_USER arguments**

Argument	Datatype	Value	Description
USER_NAME	S	<i>user_name</i>	User whose roles you are retrieving. Identify the user by the user's user name.
DOMAIN	S	<i>domain_name</i>	Domain of a client application.  If a domain is included, the method returns the user's roles within the specified domain.  If a domain is not included, the method returns all roles for the user, regardless of domain.

## Return value

This returns a collection of one query result object that has three attributes: *user\_name*, *domain*, and *roles*. *user\_name* contains the name of the user being queried. *domain*

lists the domains searched for user roles. roles is a repeating attribute that contains the roles for the user.

## Permissions

Any one can use this method.

## General notes

This method is used by client applications to determine the roles assigned to users who use the applications. For example, when a user starts a session with Desktop Client, DTC executes this method to query the domain group associated with Desktop Client, to determine what roles the user has in Desktop Client.

(For more information about groups, roles, and domains, refer to [Chapter 4, Security Services](#), of *Content Server Fundamentals*.)

## Related administration methods

None

## Examples

```
EXECUTE roles_for_user
WITH user_name=MaryJean, domain="HR_app"

dmAPIGet ("apply, c, NULL, ROLES_FOR_USER,
USER_NAME, S, MaryJean, DOMAIN, S, HR_app")
```

# SET\_APIDEADLOCK

**Purpose** Sets a deadlock trigger on a particular API method.

## Syntax

```
EXECUTE set_apideadlock
WITH API=api_name,VALUE=T|F{,API=api_name,VALUE=T|F}
dmAPIGet ("apply,session,NULL,SET_APIDEADLOCK,
API,S,api_name,VALUE,B,T|F{,API,S,api_name,VALUE,B,T|F}")
```

## Arguments

Table 3-47. SET\_APIDEADLOCK arguments

Argument	Datatype	Value	Description
API	S	<i>api_name</i>	Name of the API method or operation on which to set the deadlock trigger. <a href="#">Table 3-48, page 305</a> lists the methods and operations on which you can set a trigger.
VALUE	BOOLEAN	T (TRUE) or F (FALSE)	TRUE sets the trigger. FALSE removes the trigger.

## Return value

SET\_APIDEADLOCK returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

Anyone can use this administration method.



## General notes

Use SET\_APIDEADLOCK to test deadlock retry code in client applications that execute in explicit transactions. Operations that occur in explicit transactions are not protected by Content Server's internal deadlock retry functionality. Consequently, applications that execute in an explicit transaction may include their own deadlock retry functionality.

To test the deadlock retry functionality, use SET\_APIDEADLOCK to place a trigger on one or more methods or operations executed in the application. When the application is tested, Content Server simulates a deadlock when one of the methods or operations executes, allowing you to test the deadlock retry code in the application. The simulated deadlock sets the `_isdeadlocked` computed attribute and issues a rollback to the database.

The deadlock trigger is removed automatically from the method or operation which triggered the simulated deadlock.

[Table 3-48, page 305](#), lists the operations and methods on which you can place a deadlock trigger.

**Table 3-48. Valid operation names for SET\_APIDEADLOCK**

Method or Operation	Representing
Acquire	Acquire method
bp_transition	The completion of the bp_transition method
Branch	Branch method
Checkin	Checkin method
Complete	The completion of a work item
Demote	Demote method
Dequeue	Dequeue method
Destroy	Destroy method
exec	Any RPC EXEC call (on behalf of the Query, Readquery, Execquery, or Cachequery methods)
Next	Next method
Promote	Promote method
Queue	Queue method
Resume	Resume method
Revert	Revert method

Method or Operation	Representing
Save	Save on any object  Use this operation to put a deadlock trigger on a Saveasnew method for a SysObject.
save_content	A content save operation during a Save, Checkin, Saveasnew, or Branch operation
save_parts	A containment save operation during a Save, Checkin, Saveasnew, or Branch operation
Suspend	Suspend method

## Related administration methods

None

## Examples

This example sets a deadlock trigger on the Checkin method:

```
EXECUTE set_apideadlock  
WITH api=checkin,value=T
```

This example sets a deadlock trigger on the Promote and Revert methods:

```
dmAPIGet ("apply,S0,NULL,SET_APIDEADLOCK,  
API,S,promote,VALUE,B,T,API,S,revert,VALUE,B,T")
```

This example removes the deadlock trigger from the Checkin method:

```
EXECUTE set_apideadlock  
WITH api=checkin,value=F
```

# SET\_CONTENT\_ATTRS

**Purpose** Sets the content-related attributes of a content object.

## Syntax

```
EXECUTE set_content_attrs FOR object_id
WITH format='format_name', [page=page_number,]
[page_modifier='value',]
parameters='name="value" {, name="value"}' [, truncate=T|F]

dmAPIGet ("apply, session, object_id, SET_CONTENT_ATTRS,
FORMAT, S, format_name, [PAGE, I, page_number,]
[PAGE_MODIFIER, S, value,]
PARAMETERS, S, 'name="value" {, name="value"}' [, TRUNCATE, B, T|F] ")
```

*object\_id* is the object ID of the document that contains the content file.

## Arguments

**Table 3-49. SET\_CONTENT\_ATTRS arguments**

Argument	Datatype	Value	Description
FORMAT	S	<i>format_name</i>	Name of the content format. This is the name of the format object.
PAGE	I	<i>page_number</i>	Page number of the content in the document's set of content files. If unspecified, the default is 0.
PAGE_MODIFIER	S	<i>value</i>	Identifies the rendition. Refer to the General Notes for a detailed description of the purpose of the page modifier.

Argument	Datatype	Value	Description
PARAMETERS	S	<i>name = value</i> pairs	Comma separated list of attribute names and values. <i>value</i> is one of:  " <i>character_string</i> " FLOAT( <i>number</i> ) DATE( <i>date_value</i> )  Refer to the General Notes for examples.
TRUNCATE	B	T (TRUE) or F (FALSE)	Whether to remove existing values in the content attributes before setting the new values. The default value is F.

## Return value

SET\_CONTENT\_ATTRS returns a collection with one query result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have at least Write permission on the object or Sysadmin privileges to use this method.

## General notes

SET\_CONTENT\_ATTRS sets five attributes of a content object:

- content\_attr\_name
- content\_attr\_value
- content\_attr\_num\_value
- content\_attr\_date\_value
- content\_attr\_data\_type

The Media Server uses this method to store the metadata it generates for a content file in the file's content object.

## Using the PARAMETERS argument

The metadata is specified in SET\_CONTENT\_ATTRS as a comma-separated list of name and value pairs in the PARAMETERS argument. The format is:

```
'name=value{,name=value}'
```

Because the PARAMETERS argument is a string argument, enclose the full string in single quotes. Additionally, if *value* is a character string, enclose the individual value in double quotes. For example:

```
EXECUTE set_content_attrs FOR 0900038000ab4d13
WITH format='jpeg',page=0,
PARAMETERS='name="photo_closeup"'
```

If the metadata value is a numeric value, use the FLOAT function to specify the value:

```
EXECUTE set_content_attrs FOR 0900038000ab4d13
WITH format='jpeg',page=0,
PARAMETERS='width=FLOAT(12.5)'
```

If the metadata value is a date, use the DATE function to specify the value:

```
EXECUTE set_content_attrs FOR 0900038000ab4d13
WITH format='jpeg',page=0,
PARAMETERS='take_down_date=DATE(09/30/2002)'
```

The method sets the content object attributes with values specified in the PARAMETERS argument beginning at the zero index position. The values are set in the order in which they are included in the PARAMETERS argument.

The content\_attr\_name and content\_attr\_data\_type values are set to the name of the property and its datatype. If the property's datatype is string, the value is stored in content\_attr\_value and the remaining two attributes (content\_attr\_date\_value and content\_attr\_num\_value) are set to the default NULL values (NULLDATE and NULLINT). If the property's datatype is numeric, the value is stored in content\_attr\_num\_value and the remaining two attributes (content\_attr\_value and content\_attr\_date\_value) are set to the default NULL values (NULLSTRING and NULLDATE). If the property's datatype is date, the value is stored in content\_attr\_date\_value and the remaining two attributes (content\_attr\_num\_value and content\_attr\_value) are set to the default NULL values (NULLINT and NULLSTRING).

For example, suppose an application executes the following statement:

```
EXECUTE set_content_attrs FOR 0900038000ab4d13
WITH format='jpeg',page=0,
PARAMETERS='name="photo_closeup",width=FLOAT(12.5),take_down_date=DATE(09/30/2002)'
```

Table 3–50, page 310, shows the resulting values in the attributes in the content object.

**Table 3-50. Example settings for content metadata attributes in content objects**

Attribute	Index position		
	[0]	[1]	[2]
content_attr_name	name	width	take_down_date
content_attr_data_type	2	4	5
content_attr_value	photo_closeup	NULLSTRING	NULLSTRING
content_attr_num_value	NULLINT	12.5	NULLINT
content_attr_date_value	NULLDATE	NULLDATE	09/30/2002

The TRUNCATE argument controls how existing values in the content attributes are handled. If TRUNCATE is set to T (TRUE), existing values in the attributes are removed before the attributes are set to the new names and values. If TRUNCATE is F (FALSE), existing names and values are not removed. If the PARAMETERS argument includes a name that already present in the attributes, the name's value is overwritten. Names specified in the PARAMETERS argument that are not currently present in the attributes are appended to the attributes. The default for TRUNCATE is F.

### Using PAGE\_MODIFIER

Use the PAGE\_MODIFIER argument to define an identifier that distinguishes a rendition from any other rendition in the same format associated with a particular content page.

There are no constraints on the number of renditions that you can create for a document. Additionally, you can create multiple renditions in the same format for a particular document. To allow users or applications to distinguish between multiple renditions in the same format for a particular document, define a page modifier.

Including the PAGE\_MODIFIER argument in SET\_CONTENT\_ATTRS sets the page\_modifier attribute of the content object. This attribute, along with three others (parent\_id, page, i\_format) uniquely identifies a rendition. Applications that query renditions can use the modifier to ensure that they return the correct renditions.

## Related administration methods

[PUSH\\_CONTENT\\_ATTRS](#), page 285

## Examples

```
EXECUTE set_content_attrs FOR 090002134529cb2e
```

```
WITH format='jpeg',page=0,  
parameters='name="garage_band",sales=FLOAT(100.2),  
release_date=DATE(01/02/2002) '  
  
dmAPIGet ("apply,s0,090002134529cb2e,  
SET_CONTENT_ATTRS,FORMAT,S,jpeg,PAGE,I,0  
PARAMETERS,S,'name="garage_band",sales=FLOAT(100.2),  
release_date=DATE(01/02/2002)'" )
```

# SET\_OPTIONS

**Purpose** Turns tracing options off or on.

## Syntax

```
EXECUTE set_options  
WITH option='option_name', "value"=true|false  
  
dmAPIGet ("apply, session, NULL, SET_OPTIONS,  
OPTION, S, option_name, VALUE, B, T|F")
```

## Arguments

**Table 3-51. SET\_OPTIONS arguments**

Argument	Datatype	Value	Description
OPTION	S	<i>option_name</i>	Identifies the tracing option you want to turn on or off. Refer to <a href="#">Table 3-52, page 313</a>
VALUE	B	T (TRUE) or F (FALSE)	TRUE turns tracing on. FALSE turns tracing off.

## Return value

SET\_OPTIONS returns a collection with one result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

The tracing results for the options representing server operations are printed to the server log file. The trace\_method\_server option traces method server operations. The



tracing results generated by setting `trace_method_server` are printed to the method server log file.

Table 3–52, page 313, lists the values for the OPTION argument:

**Table 3-52. Trace options for SET\_OPTIONS**

Value for the OPTION argument	Meaning
<code>ca_store_trace</code>	Enables tracing for operations in a content-addressed storage system. The trace messages, up to 2047 characters, are placed in the server log file.
<code>clean</code>	Removes the files from the server common area.
<code>debug</code>	Traces session shutdown, change check, launch and fork information
<code>docbroker_trace</code>	Traces connection broker information
<code>file_store_trace</code>	Traces digital shredding of content files.
<code>i18n_trace</code>	Set this to T to turn on tracing of digital shredding and to F to turn off tracing. Traces client session locale and code page. An entry is logged identifying the session locale and client code page whenever a session is started. An entry is also logged if the locale or code page is changed during the session.
<code>last_sql_trace</code>	Traces the SQL translation of the last DQL statement issued before access violation and exception errors. If an error occurs, the <code>last_sql_trace</code> option causes the server to log the last SQL statement that was issued prior to the error. This tracing option is enabled by default. It is strongly recommended that you do not turn off this option. It provides valuable information to Technical Support if it ever necessary to contact them.
<code>lock_trace</code>	Traces Windows locking information

Value for the OPTION argument	Meaning
net_ip_addr	Traces the IP addresses of client and server for authentication
nettrace	Turns on RPC tracing. Traces Netwise calls, connection ID, client host address, and client host name
retention_trace	Turns on tracing for the following retention-related operations: <ul style="list-style-type: none"><li>• Assigning a retention policy to a document</li><li>• Recording a retention period applied by a retention policy at the storage area level</li><li>• Removing a retention policy from a document</li></ul>
sqltrace	Traces SQL commands sent to the underlying RDBMS for subsequent sessions
ticket_trace	Traces import and export operations for the login ticket key and operations the use a single-use login ticket.
trace_authentication	Traces detailed authentication information
trace_complete_launch	Traces Unix process launch information
trace_method_server	Traces the operations of the method server.
trace_workflow_agent	Traces operations of the workflow agent. Messages are recorded in the server log file.

## Related administration methods

[LOG\\_OFF](#), page 250  
[LOG\\_ON](#), page 251  
[MODIFY\\_TRACE](#), page 269

## Examples

This example turns on SQL tracing:

```
EXECUTE set_options  
WITH option='sqltrace', "value"=true
```

The following example turns on logon authentication tracing:

```
dmapiGet ("apply,c,NULL,SET_OPTIONS,
```

```
OPTION,S,trace_authentication,VALUE,B,T")
```

This example turns off tracing for logon authentication:

```
dmapiGet ("apply,c,NULL,SET_OPTIONS,  
OPTION,S,trace_authentication,VALUE,B,F")
```

This example enables tracing for a content-addressed storage system:

```
dmAPIGet ("apply,c,NULL,SET_OPTIONS,OPTION,S,ca_store_trace,VALUE,B,T")
```

## SET\_STORAGE\_STATE

**Purpose** Takes a storage area offline, places an off-line storage area back online, or makes a storage area read-only.

### Syntax

```
EXECUTE set_storage_state [[FOR] 'storage_area_obj_id']
[WITH store=storage_area_name{, argument=value}]
```

With EXECUTE, if you include the STORE argument, do not include the FOR clause.

```
dmAPIGet ("apply, session, storage_area_obj_id, SET_STORAGE_STATE
[, STORE, S, storage_area_name] [, argument, datatype, value]")
```

With Apply, specify the storage area object ID as NULL if you include the STORE argument.

### Arguments

**Table 3-53. SET\_STORAGE\_STATE arguments**

Argument	Datatype	Value	Description
OFFLINE	B	T (TRUE) or F (FALSE)	If set to TRUE, the storage area is taken off-line.
READONLY	B	T (TRUE) or F (FALSE)	If set to TRUE, the storage area is read-only.
STORE	S	<i>storage_area_name</i>	Identifies the storage area whose state you are changing. Use the name of the area's storage area object.

### Return value

SET\_STORAGE\_STATE returns a collection with one result object. The object has one Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation.

## Permissions

You must have Sysadmin or Superuser privileges to use this method.

## General notes

A storage area has three possible states:

- Online

Users can read from and write to an online storage area.

- Offline

Users can neither read from nor write to an offline storage area.

- Readonly

Users can only read from a readonly storage area. Users cannot write to a readonly storage area.

To set a storage area's state to offline, issue SET\_STORAGE\_STATE with the OFFLINE argument set to T (TRUE); for example:

```
EXECUTE set_storage_state WITH store=filestore_33,OFFLINE=true
```

To set a storage area's state to readonly, issue SET\_STORAGE\_STATE with the READONLY argument set to T (TRUE); for example:

```
EXECUTE set_storage_state WITH store=filestore_33,READONLY=true
```

To change either an offline or readonly storage area back to the online state (users can read and write the storage area), issue the SET\_STORAGE\_STATE method without specifying a state argument; for example:

```
EXECUTE set_storage_state WITH store=filestore_33
```

Setting either OFFLINE or READONLY to FALSE directly has no effect.

To use SET\_STORAGE\_STATE, you must be using one server for both data and content requests. If the configuration is set up for content servers, you must issue a Connect method that bypasses the content server to use SET\_STORAGE\_STATE in the session.

## Related administration methods

None

## Examples

The following example moves the storage area called manfred offline:

```
EXECUTE set_storage_state
```

```
WITH store = 'manfred', offline = T
```

The next example sets the storage1\_engr file store storage area offline:

```
dmAPIGet ("apply,c,NULL,SET_STORAGE_STATE,OFFLINE,B,T,  
STORE,S,storage1_engr")
```

---

# SHOW\_SESSIONS

**Purpose** Returns information about all the currently active repository sessions.

## Syntax

```
EXECUTE show_sessions  
dmAPIGet ("apply, session, NULL, SHOW_SESSIONS")
```

## Arguments

SHOW\_SESSIONS has no arguments.

## Return value

SHOW\_SESSIONS returns a collection. Each query result object in the collection represents one currently active repository session. The attributes of the objects contain information about the sessions. The attributes returned by SHOW\_SESSIONS are the same as those for LIST\_SESSIONS when LIST\_SESSIONS is run to return a full set of attributes. For a description of the attributes, refer to [Table 3–28, page 244](#).

## Permissions

Anyone can use this method.

## General notes

SHOW\_SESSIONS does not return information about historical (timed out) sessions. It only returns information about currently active sessions. The information for each session is identical to the information returned by LIST\_SESSIONS for active sessions if that method is run with the BRIEF\_INFO argument set to FALSE.

**Note:** You can use SHOW\_SESSIONS to obtain the process ID if you need to shutdown or kill a session or server.

## Related administration methods

[LIST\\_SESSIONS](#), page 244

## Examples

Refer to the syntax description for examples.



# SYNC\_REPLICA\_RECORDS

**Purpose** Synchronizes the `dmi_replica_record` objects with the current definition of a distributed storage area. (This method is not available through Documentum Administrator.)

## Syntax

```
EXECUTE sync_replica_records [[FOR] 'storage_obj_id']
[WITH argument=value {, argument=value}]
```

With Execute, do not include the FOR clause if you include the STORE argument.

```
dmAPIGet ("apply, session, storage_obj_id, SYNC_REPLICA_RECORDS
{, argument, datatype, value}")
```

With Apply, specify the storage object ID as NULL if you include the STORE argument.

## Arguments

Table 3-54. SYNC\_REPLICA\_RECORDS arguments

Argument	Datatype	Value	Description
STORE	S	<i>storage_area_name</i>	Identifies the distributed storage area for which you want to synchronize the records. Use the name associated with the area's <code>dm_distributedstore</code> object. Optionally, omit the STORE argument and specify the object ID of the <code>dm_distributedstore</code> object.
COMMIT_INTERVAL	I	<i>integer</i>	Specifies the number of replica record objects to process before each commit to the repository. The default is 10,000.

Argument	Datatype	Value	Description
START_ID	S	<i>object_id</i>	Identifies a replica record object that serves as the starting point for the operation. Replica record objects that have an object ID greater than the specified object ID are processed.  Use this in conjunction with END_ID.
END_ID	S	<i>object_id</i>	Identifies a replica record object that is the stopping point for the synchronization process. Record replica objects that have an object ID greater than this are not processed.  Use this in conjunction with START_ID.
CONTINUE_ON_ERROR	B	T (TRUE) or F (FALSE)	Indicates whether the method continues in the event of an error. The default is FALSE.
REMOVE_NULL_ENTRIES	B	T (TRUE) or F (FALSE)	Indicates whether to remove null component entries. The default is T.
UPDATE_EPOCH	B	T (TRUE) or F (FALSE)	Indicates whether to bring the epoch value up to date even if the record does not need to be upgraded. The default is T.
TRACE_LEVEL	I	<i>integer</i>	Sets the tracing level of the method. Valid trace levels are:  0, meaning no tracing 1, meaning informational messages only 2, meaning trace all updated objects 5, meaning trace all objects checked for update  The default is 0.

## Return value

SYNC\_REPLICA\_RECORDS returns a collection with one query result object. The object has one Boolean attribute, named `result`, whose value indicates the success (TRUE) or failure (FALSE) of the operation.

If `CONTINUE_ON_ERROR` is set to TRUE, the method returns FALSE only when the method encounters an error from which it cannot recover.

## Permissions

You must have Sysadmin or Superuser privileges to run this method.

## General notes

Use SYNC\_REPLICA\_RECORDS after you upgrade from DocPage Server 3.x to Documentum 4.x. The method updates the `dmi_replica_record` objects to synchronize them with the Documentum 4.x architecture by removing unneeded null component entries. As a side effect, the method also removes entries for components that are no longer part of the distributed storage area.

Using SYNC\_REPLICA\_RECORDS is optional. Running the method may improve performance, but Content Server and the distributed features function correctly if you do not run it.

In a Documentum 4.x installation, you can run this method after you delete a distributed storage area component. However, it isn't required. Content Server will dynamically remove deleted component entries as the documents are referenced.

## Sizing the commit interval

The method processes objects in chunks. That is, it processes a defined number of objects and then commits the work to the repository before processing more objects. The default number of objects in a chunk is 10,000. Use the `COMMIT_INTERVAL` argument to change the size of an operational chunk.

## Batching the operation

The `START_ID` and `END_ID` arguments allow you to process a specified number of objects in each execution of the operation. If you do not include these arguments, the method processes all the replica record objects associated with the specified storage area.

## Related administration methods

None

## Examples

These examples synchronize the replica records for the diststore1 distributed storage area:

```
EXECUTE sync_replica_records  
WITH store='diststore1'  
  
dmAPIGet ("apply,c,NULL,SYNC_REPLICA_RECORDS,  
STORE,S,diststore1")
```

This example specifies a commit interval of 5,000 objects:

```
EXECUTE sync_replica_records  
WITH store='diststore1',commit_interval=5000
```

This example sets the tracing level to 5, to trace all objects checked for updating:

```
dmAPIGet ("apply,c,NULL,SYNC_REPLICA_RECORDS,  
STORE,S,diststore1,TRACE_LEVEL,I,5")
```

# TRANSCODE\_CONTENT

**Purpose** Queues a transformation request for a content file to the Media Server.

## Syntax

```
EXECUTE transcode_content FOR object_id
WITH [page=page_number,] [priority=priority_level,]
message='message',source_format='format_name',
target_format='format_name'

dmAPIGet ("apply,session,object_id,TRANSCODE_CONTENT,
[PAGE,I,page_number,][PRIORITY,I,priority_level,]
MESSAGE,S,'message',TARGET_FORMAT,S,format_name,
SOURCE_FORMAT,S,format_name")
```

*object\_id* is the object ID of the document that contains the content file.

## Arguments

**Table 3-55. TRANSCODE\_CONTENT arguments**

Argument	Datatype	Value	Description
PAGE	I	<i>page_number</i>	Page number of the content file. If unspecified, the default is 0.
PRIORITY	I	<i>priority_level</i>	Identifies the priority of the request. A value of zero indicates no priority. Priority rises as the value rises. The Media Server processes higher priority requests before lower priority requests.  If unspecified, the default is 5.

Argument	Datatype	Value	Description
MESSAGE	S	-profile_id <i>object_id</i>	Identifies the profile that contains the definition of the transformation to be performed on the content. The profile is identified by its object ID.  Refer to the General Notes for a more detailed description.
TARGET_FORMAT	S	<i>format_name</i>	Identifies the format to which to transform the content. Use the name of the format as it is defined in the format's format object.
SOURCE_FORMAT	S	<i>format_name</i>	Identifies the content file's starting format. Use the name of the format as it is defined in the format's format object.

## Return value

TRANSCODE\_CONTENT returns a collection with one query result object. The object has two attributes, result and result\_id. The result attribute is a Boolean attribute whose value indicates the success (TRUE) or failure (FALSE) of the operation. The result\_id attribute records the object ID of the queue item that is created as by the method.

## Permissions

You must have at least Write permission on the object or Sysadmin privileges to use this method.

## General notes

TRANSCODE\_CONTENT is issued by WebPublisher™ to request a transformation operation on a content file by the Media Server. The method generates an event, dm\_transcode\_content, that is queued to the Media Server. When the Media Server processes the event, the server performs the transformation on the content as defined in the specified profile.

The `dm_queue_item` that represents the event is queued to the `dm_mediaserver` user, who represents the Media Server. (`dm_mediaserver` is created when Documentum Media Transformation Services is installed.) `TRANSCODE_CONTENT` sets the following attributes in the queue item:

**Table 3-56. Queue item attributes set by `TRANSCODE_CONTENT`**

Attribute	Set to
<code>event</code>	<code>dm_transcode_content</code>
<code>item_id</code>	object ID of the SysObject that triggered the request
<code>instruction_page</code>	page number of the content in the SysObject
<code>content_type</code>	starting format of the SysObject
<code>message</code>	value specified in the MESSAGE argument
<code>item_type</code>	format to which to transform the content
<code>date_sent</code>	date the request was made
<code>name</code>	<code>dm_mediaserver</code>
<code>sent_by</code>	session user
<code>priority</code>	priority level identified in the REGISTER_ASSET call

The MESSAGE argument specifies the transformation operation to perform by specifying a profile. A profile is an XML document, stored in the repository, that defines transformation operations. The value for a MESSAGE argument has the following format:

```
'-profile_id object_id'
```

*object\_id* is the object ID of the profile document.

For example, suppose that `090000132400c2e198` is the object ID of a document that you want to transform and that `09000013240053ae1b` is the object ID of a profile containing the definition of the desired transformation. The following DQL EXECUTE statement generates a request to the Media Server to perform the transformation:

```
EXECUTE transcode_content FOR 090000132400c2e198
WITH page=0,message='-profile_id 09000013240053ae1b',
source_format='jpeg',target_format='gif'
```

## Related administration methods

[REGISTER\\_ASSET, page 290](#)

**Example**

```
EXECUTE transcode_content FOR 090000017456ae1c
WITH page=0,priority=5,
message='-profile_id 090000018125ec2b',
source_format='jpeg',
target_format='jpeg_lres'

dmAPIGet ("apply,s0,090000017456ae1c,TRANSCODE_CONTENT,
PAGE,I,0,PRIORITY,I,5,
MESSAGE,S,'-profile_id 090000018125ec2b',
SOURCE_FORMAT,S,jpeg,
TARGET_FORMAT,S,jpeg_lres")
```



# UPDATE\_STATISTICS

**Purpose** Updates statistics for RDBMS tables in the repository.

## Syntax

```
EXECUTE update_statistics WITH table_name='name'
[,count=integer] [,extra_data=value]
```

```
dmAPIGet ("apply, session, NULL, UPDATE_STATISTICS,
TABLE_NAME, S, name[, COUNT, I, integer] [, EXTRA_DATA, S, value]
```

## Arguments

**Table 3-57. UPDATE\_STATISTICS arguments**

Argument	Datatype	Value	Description
TABLE_NAME	S	<i>name</i>	Name of the database table whose statistics you want to update.
COUNT	I	<i>integer</i>	This argument is available only if the database is Oracle, SQL Server, or Sybase. It is not available on DB2.  The argument is used differently for each database. Refer to the General Notes for details.
EXTRA_DATA	S	<i>value</i>	The argument is used differently for each database. Refer to the General Notes for details.

## Return value

UPDATE\_STATISTICS returns a collection with one query result object. The object has one attribute, called result, that contains T if the method was successful or F if the method was unsuccessful.

## Permissions

You must be a Superuser to execute this method.



**Caution:** Do not use this method unless directed to do so by Technical Support.

## General notes

UPDATE\_STATISTICS is used by the UpdateStatistics administration tool when the tool is run on a DB2 database. However, the method can be used on any supported database.

### The COUNT argument

This is an optional argument. It is used differently in Oracle, SQL Server, and Sybase. It is not used if the database is a DB2 database.

For Oracle, it defines the number of buckets to use when calculating the histogram statistics. The valid range for COUNT on an Oracle database is 1 to 254. The default is 75. Setting COUNT to 0 deletes all histogram statistics for the table.

For SQL Server, this argument defines what percentage of table rows to use when calculating the statistics. For example, if you set this to 50, the method uses one half (50%) of the table rows to calculate the statistics. The default is 100, meaning all rows are used to calculate the statistics.

For Sybase, this argument defines how many steps to use when calculating the statistics. The default is 20.

### The EXTRA\_DATA argument

The EXTRA\_DATA argument is an optional argument. Its use differs depending on the database.

For Oracle and Sybase, this identifies a specific set of table columns to be analyzed for statistics. The columns must exist in the table. The columns are specified as a comma-separated list of column names. Enclose the entire list in single quotes. For example:

```
dmAPIGet ("apply, c, UPDATE_STATISTICS, TABLE_NAME, S, dm_sysobject_s,  
EXTRA_DATA, 'keywords, authors'")
```

If you do not include EXTRA\_DATA, all columns in the table are used.

Including this argument for Sybase databases can be expensive because the RDBMS must scan the table and perform a sort operation on the table.

For SQL Server and DB2, the argument identifies an index for which you want to generate statistics. The index must be an index on the table identified in TABLE\_NAME.

If you do not include the argument, the database table identified in `TABLE_NAME` and all of its indexes are analyzed.

## Related administration methods

[REORGANIZE\\_TABLE](#), page 292

## Examples

This example updates statistics for the `dm_sysobject_s` table:

```
EXECUTE update_statistics
WITH table_name='dm_sysobject_s'

dmAPIGet ("apply,S0,NULL,UPDATE_STATISTICS,
TABLE_NAME,S,dm_sysobject_s")
```

This example updates the statistics for the `dm_user_s` table, specifying that only the `user_name`, `user_os_name` and `user_address` columns be used for the analysis:

```
EXECUTE update_statistics
WITH table_name='dm_user_s',
extra_data='user_name,user_os_name,user_address'

dmAPIGet ("apply,S0,NULL,UPDATE_STATISTICS,
TABLE_NAME,S,dm_user_s,
EXTRA_DATA,S,'user_name,user_os_name,user_address'")
```

# WEBCACHE\_PUBLISH

**Purpose** Invokes the `dm_webcache_publish` method to publish documents to a Web site. (This administration method cannot be run using the EXECUTE statement.)

## Syntax

```
dmAPIGet ("apply, session, webc_config_obj_id, WEBCACHE_PUBLISH
[, ARGUMENTS, S, argument_list]")
```

## Arguments

**Table 3-58. WEBCACHE\_PUBLISH arguments**

Argument	Datatype	Value	Description
ARGUMENTS	S	<i>argument_list</i>	Identifies the arguments that you want to pass to the <code>dm_webcache_publish</code> method. The valid entries for <code>argument_list</code> are described in <a href="#">Table 3-59, page 332</a> .

[Table 3-59, page 332](#), lists the valid entries for the ARGUMENTS argument.

**Table 3-59. Valid arguments for ARGUMENTS**

Argument	Value	Description
-source_object_id	<i>object_id</i>	Identifies the object or source folder to be refreshed. If this identifies an object, the object must reside under the source folder identified in the <code>webc config</code> object specified in the command line. If this identifies a folder, the folder must be the source folder or reside under the source folder.

Argument	Value	Description
		<p>If you specify a folder, all objects in and underneath the folder are refreshed.</p> <p>If you don't include this argument, the entire source data set is refreshed.</p> <p>This option is ignored if <code>-full_refresh</code> is set to TRUE.</p>
<code>-full_refresh</code>	TRUE or FALSE	<p>When this is set to TRUE, the content and attribute data at the target WebCache are deleted and republished.</p> <p>You must have Superuser privileges to set this option to TRUE. The default is FALSE.</p>
<code>-force_refresh</code>	TRUE or FALSE	<p>When this is set to TRUE, documents are refreshed even if their modification dates do not indicate a need for a refresh.</p> <p>The default for this option is FALSE.</p>
<code>-method_trace_level</code>	<i>trace_level</i>	<p>Turns on tracing for the method at the indicated level. Valid trace levels correspond to the levels available for the Trace method.</p>

Argument	Value	Description
-recreate_property_schema	TRUE or FALSE	Performs a full refresh and destroys and recreates the propdb tables.  You must have Superuser privileges to set this option to TRUE. The default is FALSE.
—resync_state_table	TRUE or FALSE	TRUE deletes state information in the probdb_tablename_m table and recreates the information based on the current configuration object.  The default for this flag is FALSE.
-store_log	TRUE or FALSE	Creates a log file object in the repository for each publish operation.  By default, this flag is TRUE for all full refresh and incrementals publishing operations and FALSE when the method is issued for a single item.

### Return value

The WEBCACHE\_PUBLISH administration method returns a collection identifier for a collection with one attribute. The attribute, method\_return\_val, contains 0 if the operation was successful.

### Permissions

You must have Sysadmin or Superuser privileges to use this administration method.

## Generating a log file

When the operations generate a log file, the document is stored in the repository in `/System/Sysadmin/Reports/Webcache`. By default, the method does not generate a log file if you are publishing only a single item. To force it to generate a log file for single-item operations you can include the `-store_log` argument, set to `TRUE`

## Example

```
dmAPIGet ("apply,s0,080015b38001a43c,WEBCACHE_PUBLISH,  
ARGUMENTS,S,-source_object 090015b38007bf19")
```





## Using DQL

This chapter introduces DQL and then provides usability information for DQL. The chapter includes the following topics:

- [Introducing DQL, page 337](#), introduces the Document Query Language and its basic query statements.
- [Quoting object type and attribute names, page 340](#), describes a recommended best practice for referencing object type and attribute names in queries.
- [NULLs, default values, and DQL, page 340](#), describes how DQL handles NULLs and default values.
- [Repeating attributes in queries, page 343](#), describes how to reference repeating attributes in queries.
- [Querying virtual documents, page 347](#), describes how to query virtual documents.
- [Full-text searching and virtual documents, page 348](#), describes how you can conduct a fulltext search on an assembled virtual document.
- [Querying registered tables, page 349](#), describes how to reference registered tables in queries.
- [Caching queries, page 350](#), describes how to cache query results.
- [Using execquery in an application, page 351](#), provides a basic outline for executing and processing queries in an application.
- [Privileges, permissions, and queries, page 354](#), describes how privileges and permissions affect query results.

For information about using DQL and XDQL to query XML documents, refer to the *XML Application Development Guide*.

## Introducing DQL

DQL is the acronym for Document Query Language, the SQL-like language that you use to query the objects in a repository. Using DQL, you can retrieve, update, and

delete objects and create new objects. DQL also allows you to search indexed content and metadata (attribute values).

You can also use DQL to access registered tables—tables in the underlying RDBMS that are known to Content Server but that are not part of the repository. ([Repositories](#), page 60, describes the database tables that make up a repository.)

The basic DQL query statements retrieve information about the objects in a repository and manipulate those objects. [Table 4-1, page 338](#), describes the basic query statements.

**Table 4-1. DQL basic query statements**

Statement	Description
SELECT	<p>The SELECT statement is the primary query statement. SELECT statements return a wide variety of information from the repository. SELECT statements are not only stand-alone statements but are also used in other DQL statements.</p> <p>There are two variants of the SELECT statement. One is the standard SELECT statement, and the other is called an FTDQL SELECT statement. FTDQL queries are run entirely against the fulltext index, which provide performance benefits. However, the syntax for an FTDQL query is a subset of that for a standard query. For more information, refer to <a href="#">Select</a>, page 115 or to the summary of the syntax in <a href="#">Appendix C, DQL Quick Reference</a></p>
UPDATE	<p>The UPDATE statement changes the RDBMS table that underlies a registered table.</p>

Statement	Description
UPDATE...OBJECT	<p>The UPDATE...OBJECT[S] statement modifies objects in the repository. Using this statement, you can:</p> <ul style="list-style-type: none"> <li>• Set the value of a single-valued attribute</li> <li>• Append or insert values in a repeating attribute</li> <li>• Remove a value from a repeating attribute</li> <li>• Truncate a repeating attribute (remove all values)</li> <li>• Link an object to a folder or cabinet or unlink an object</li> <li>• Move an object to a new folder or cabinet</li> </ul>
CHANGE...OBJECT	<p>The CHANGE...OBJECT[S] statement moves one or more objects from one object type to another. With some constraints, you can change an object of a particular type to any type that is a subtype or supertype of the current type. For example, you can change an object of type <code>dm_document</code> to a user-defined document subtype.</p>
INSERT	<p>The INSERT statement adds a row to the RDBMS table that underlies a registered table.</p>
DELETE	<p>The DELETE statement destroys rows in the RDBMS table that underlies a registered table.</p>
DELETE...OBJECT	<p>The DELETE...OBJECT[S] statement deletes objects from the repository.</p>

When you issue one of these statements, only those objects for which you have the appropriate permissions are affected. For example, if you issue an UPDATE...OBJECT statement to update document objects, only those documents for which you have at least Write permission are considered for updating. In addition, the statements have optional clauses that let you identify specifically which objects are the target of the operation.

For a complete description of query statement syntax and usage, refer to the statement descriptions in [Chapter 2, DQL Statements](#), in the *Content Server DQL Reference Manual*.

## Quoting object type and attribute names

Place double quotes around all references to object type names and attribute names in DQL queries. While not required, this best practice ensures that the names will not conflict with words reserved by DQL or the underlying RDBMS.

To encourage this practice, object type and attribute names are double quoted in all Content Server documentation.

## NULLs, default values, and DQL

Documentum does not allow you to assign NULLs as an attribute value because NULLs are used to terminate the lists of values in repeating attributes.

When you create an object, any attributes that you do not set are assigned default values by the server. If a default value is defined in the data dictionary, the server assigns that value. Otherwise, the default value depends on the data type of the attribute. [Table 4-2, page 340](#), lists the default values by data type.

**Table 4-2. Default attribute values by datatype**

Datatype	Default value
String	A single blank
Numeric data type	0
Date	The NULLDATE value: <ul style="list-style-type: none"> <li>• For Oracle and DB2, it is 01/01/0001.</li> <li>• For Sybase, the value is 1/1/1753.</li> <li>• For MS SQL Server, it is 1/1/1753.</li> </ul>

When you query an attribute containing a default value using DQL, the default values are returned in the following manner:

- The default character string value (a single blank) is returned as an empty string.
- DQL also returns a single blank found in a registered table column as an empty string.
- The default numeric value (zero) is returned as zero.
- NULLDATE values are returned as the word NULLDATE.

If you are using SQL (for example, in a user-written report writer), the default values are returned in the following manner:

- The default character string value (a single blank) is returned as a single blank.
- SQL also returns a single blank found in a registered table column as a single blank.

- The default numeric value (zero) is returned as zero.
- NULLDATE values are returned as the actual date (for example, 1/1/1 in Oracle).

## Testing for default and NULL values

Documentum provides predicates for use in WHERE clauses to test for the presence of default values or NULLs. [Table 4–3, page 341](#), briefly describes these predicates. For a complete list of predicates, refer to [Predicates, page 29](#), of the *Content Server DQL Reference Manual*.

**Table 4-3. Predicates that test for NULL and default values**

Predicate	Description
IS [NOT]NULL	Used only for columns in registered tables. Tests for the presence of NULL.
ANY...IS [NOT]NULL	Tests for the presence of the NULL terminator in a repeating attribute. This is primarily useful if you are testing values in two or more repeating attributes at corresponding index levels.
[ANY]...IS [NOT]NULLDATE	Tests for the presence of the NULLDATE value. Use the ANY option if the attribute is a repeating attribute.
[ANY]...IS [NOT]NULLSTRING	Tests for the presence of the default string value (a single blank) in an attribute. Use the ANY option if the attribute is a repeating attribute.
[ANY]...IS [NOT]NULLINT	Tests for the presence of the default numeric value (a zero) in an attribute. Use the ANY option if the attribute is a repeating attribute.

For more information about querying repeating attributes, refer to [Repeating attributes in queries, page 343](#).

## Default values and aggregate functions

Aggregate functions are functions that operate on a group of values and return a single value. For example, count is an aggregate function. It counts the values in an attribute and returns the number.

Aggregate functions do not ignore the default values assigned to an attribute. These values are included in any calculations performed by the function. For example, suppose a repeating attribute value can range from 1 to 10. If an object has three actual values (4, 7, and 6) and one default value (one zero) for the attribute, the AVG function adds 4, 7, 6, and 0, then divides the total by 4 to obtain the average:

$$(4 + 7 + 6 + 0) / 4 = 4.25$$

If the function ignores the default value, the average is:

$$(4 + 7 + 6) / 3 = 5.66$$

As another example, look at the following SELECT statement:

```
SELECT COUNT(DISTINCT ratings) FROM "recipe"  
WHERE "object_name" = 'lemon cake'
```

This statement (assuming the attribute values for ratings described for the previous example) returns the number 4 because the default value is counted.

You can use the WHERE clause to exclude the default value when you are selecting aggregate functions. For example, the following statement returns the correct number of ratings for the lemon cake recipe:

```
SELECT COUNT(DISTINCT rating) FROM "recipe"  
WHERE "object_name" = 'lemon cake'  
AND "rating" IS NOT NULL
```

## Sorting and nulls

Different databases handle NULLs differently when they are sorting values in a table. Here is how each repository handles NULLs if the sort order is ascending:

- In Oracle and DB2, NULLs are sorted to the bottom of the list.
- In Sybase, NULLs are sorted to the top of the list.
- In MS SQL Server, NULLs are sorted to the top of the list.

## Repeating attributes in queries

Repeating attributes store multiple values. In Documentum, many object types have repeating attributes. For example, because a document can have multiple authors, the authors attribute of the dm\_document type is a repeating attribute.

You can reference repeating attributes in:

- The selected values list in a SELECT statement
- WHERE clause qualifications for SELECT, UPDATE...OBJECT, CHANGE...OBJECT, and DELETE...OBJECT statements
- The update clauses of the CREATE...OBJECT, UPDATE...OBJECT, and CHANGE...OBJECT statements

For information about using a repeating attribute in SELECT statement, as a selected value or in the WHERE clause qualification, refer to [Select](#), page 115.

## Modifying repeating attributes

You can use the UPDATE...OBJECT statement to add, delete, and modify individual values for repeating attributes. In most of these operations, you must specify the *index position* of the value you want to change. The index indicates the value's position in the list of values assigned to the repeating attribute. Index numbers begin with zero for the first value and increment by 1 for each additional value.

For example, suppose the a recipe object type has a repeating attribute called ingredients. One recipe has the following ingredients: eggs, sugar, cream cheese, and amaretto. Assuming the ingredients were added to the ingredients attribute for a recipe object in the order listed, they would have the following index values:

```
ingredient[0] (eggs)
ingredient[1] (sugar)
ingredient[2] (cream cheese)
ingredient[3] (amaretto)
```

Index positions are always specified inside square brackets after the name of the attribute.

## Adding new values

To add a value to a repeating attribute, you can:

- Insert the new value, which lets you choose where to put the new value in the attribute's value list

- Append the value, which automatically places the value at the end of the attribute's value list

## Inserting values

When you want to control where the new value is placed in the repeating attribute, use the insert option as the update clause in the UPDATE...OBJECT statement. The syntax is:

```
UPDATE object_type OBJECTS
INSERT attribute_name[x] = value
...
```

For example, suppose that the authors of the Espresso Cheesecake recipe forgot to include one of the ingredients for the crust. The ingredients for the crust are listed ahead of the ingredients for the filling in the document. Consequently, the authors want to insert the forgotten ingredient ahead of the filling ingredients so that it appears with the other crust ingredients. The ingredients for the crust occupy positions 0 through 3 in the ingredients attribute, and the filling ingredients begin at position 4. The following statement inserts the forgotten ingredient at position 4.

```
UPDATE "recipe" OBJECTS
INSERT ingredients[4] = '2 T ground bittersweet chocolate'
WHERE "title" = 'Espresso Cheesecake'
```

The server inserts the requested value and rennumbers all other values from position 4 on. The filling ingredients now begin at position 5.

Inserting values into a repeating attribute never overwrites a current value. Any value currently at the specified insertion point and any following values are always renumbered.

If you do not specify an insertion point, the server automatically inserts the new value in the first position (*attribute\_name*[0]).

You can identify the value you want to insert either as literal value or by using a subquery. If you use a subquery, it can return only one value. If it returns multiple values, the INSERT statement fails with an error.

## Appending values

When you append values, the server automatically adds the new value to the end of the list of values in the attribute. You do not have to specify an index value when you append. To illustrate, suppose the authors want to add another ingredient to the Espresso Cheesecake recipe: strawberries to be used as a garnish. To append this ingredient, they use the following statement:

```
UPDATE "recipe" OBJECT
APPEND "ingredients" = '1 pint strawberries'
```



```
WHERE "title" = 'Espresso Cheesecake'
```

You can identify the appended value as a literal value or by using a subquery. For example, perhaps your company is reorganizing and employees currently working for Mr. Rico are being moved to a group called frontline. The following statement uses a subquery to find those employees and append them to the list of users in the frontline group:

```
UPDATE "dm_group" OBJECTS
APPEND "users_names"=(SELECT "user_name"
FROM "dm_user", "employees" e
WHERE e."manager"='rico' AND "r_is_group"=F)
WHERE "group_name" = 'frontline'
```

When using a subquery, you can specify the maximum number of values that you want to append. The syntax in the update clause is:

```
APPEND n attribute_name = subquery
```

where *n* represents the maximum number of appended values.

## Updating values

To update a value, use the set option as the update clause.

For example, the authors of Heavenly Cheesecakes have decided to update one of the recipes in the book. They want to replace the cream cheese requirement with a lower-fat substitute—ricotta cheese. Knowing that cream cheese is the third ingredient in the ingredients list, they use the following statement to make the substitution:

```
UPDATE "recipe" OBJECT
SET ingredients[2] = '1 lb ricotta cheese'
WHERE "title" = 'mocha cheesecake'
```

Note that the index value for cream cheese is [2]. This is because index values are counted from zero, so the first two ingredients (eggs and sugar) have index values of zero and one ([0] and [1]), respectively.

## Deleting values

To delete a value in a repeating attribute, use the remove option as the update clause and specify the index value associated with the value. To illustrate, the following example removes the fifth ingredient in a list of ingredients for brownies:

```
UPDATE "recipe" OBJECT
REMOVE ingredients[4]
WHERE "object_name" = 'Mandarin Brownies'
```

If you do not specify which value to remove, the system automatically removes the first ([0]) value.

## Forcing index correspondence in query results

**Note:** The query syntax described in this section may not be used in an FTDQL SELECT statement. (For information about FTDQL SELECT statements, refer to [Select, page 115.](#))

In some queries, you may only want objects for which the repeating attribute values are in the same relative positions. For example, perhaps you want only those recipes for which a particular author and keyword occupy the same index position.

To require index correspondence for expressions referencing repeating attributes, use the following syntax:

```
ANY ([NOT] predicate AND [NOT] predicate
    {AND [NOT] predicate})
```

To force index correspondence, the predicates must be ANDed inside the parentheses. Using the OR operator to link the predicates inside the parentheses increases the query's performance but does not force index correspondence. The query returns any object that contains one of the ORed values.

To illustrate, assume that the keywords untried, tested, approved, and rejected are assigned to recipes in various stages of acceptance and that these keywords are always placed in the first position in the keywords list, to correspond to the implicit version label (the numeric label) associated with the recipe. The following statement finds only original recipes (version 1.0) that were rejected by the testers:

```
SELECT "r_object_id", "object_name" FROM "recipe"
WHERE ANY ("r_version_label" = '1.0'
AND "keywords" = 'rejected')
```

In the above example, for the recipes returned, `r_version_label[0] = 1.0` and `keywords[0] = rejected`.

You can't specify which position is searched when you force index correspondence. You can only specify that the values be found in the same position. In the above example, it was easy to know that the values are in the first position for all pairs because the implicit version label is always stored in the first position. However, for other repeating attributes this will not be true. For example, look at the following statement:

```
UPDATE "recipe" OBJECTS
SET "subject" = 'lowfat meals'
WHERE ANY ("ingredients" IN ('skim milk', 'margarine')
AND "keywords" = 'fat free')
```

The statement updates the subject attribute for any recipe for which skim milk and fat free or margarine and fat free occupy the same respective positions within their individual attributes. For one recipe, the positions might be the fourth: `ingredients[3] = margarine` and `keywords[3] = fat free`. For another, the values might be found in the

seventh position: ingredients[6] = skim milk and keywords[6] = fat free. In all instances, the values are at the same index position within their attributes.

It is possible to combine both forms of the syntax. For example:

```
SELECT "r_object_id", "title" FROM "recipe"
WHERE ANY "ingredients" IN ('cream cheese', 'chocolate')
AND ANY("authors" = 'daphne' AND "keywords" IN ('cheesecake', 'brownies', 'mousse'))
```

This statement returns all recipes that have cream cheese or chocolate as an ingredient and also have the author daphne paired with the keyword cheesecake, brownies, or mousse.

## Performance note for Sybase or MS SQL Server users

DQL turns repeating attribute predicates (and the FOLDER predicate) into ORed subselects internally. Queries that contain ORed subselects run slowly against Sybase and MS SQL Server, and performance degrades in direct proportion to the number of SysObjects in your repository.

To improve performance, here are some suggested alternative ways to formulate queries:

Instead of:

```
SELECT "r_object_id" FROM "dm_sysobject"
WHERE ANY "authors" = 'a' OR ANY authors = 'b'
```

Use:

```
SELECT "r_object_id" FROM "dm_sysobject"
WHERE ANY ("authors" = 'a' OR "authors" = 'b')
```

Instead of:

```
SELECT "r_object_id" FROM "dm_sysobject"
WHERE FOLDER ('/Temp) OR FOLDER ('/System')
```

Use:

```
SELECT "r_object_id" FROM "dm_sysobject"
WHERE FOLDER ('/Temp', '/System')
```

## Querying virtual documents

Virtual documents are documents that contain other documents. The documents they contain can be simple documents or other virtual documents. Nesting virtual documents inside other virtual documents creates a hierarchy of components within the top-level virtual document. Documentum allows you to nest virtual documents to any depth you choose.

To identify the virtual document you want to query, you can use either the IN DOCUMENT clause or the IN ASSEMBLY clause in a SELECT statement. The IN DOCUMENT clause allows you to choose which set of the document's components to query at runtime. The IN ASSEMBLY clause directs the query to the virtual document's snapshot, which is a previously selected set of the document's components.

If you are querying a virtual document in an UPDATE...OBJECT or DELETE...OBJECT statement, you must use the IN ASSEMBLY clause.

Both the IN DOCUMENT and IN ASSEMBLY clauses include the optional keyword DESCEND. This keyword directs the server to search any components contained by components that are themselves virtual documents. The only exception is if the component is a reference link. In such cases, the server can't search for the component's contained components. The search returns only the reference link.

Documentum also allows you to specify one particular component of a virtual document as the target of a query. This feature can make it easy for users who are working on a virtual document to pull out a specific part of the document for work. The NODE option of the IN ASSEMBLY clause implements this feature.

## Full-text searching and virtual documents

Like simple documents, virtual documents can have associated content files. Marking a virtual document for indexing directs Content Server to index the content files associated with the virtual document—*not* the components of the virtual document.

In SELECT statements, if you include a SEARCH clause and an IN DOCUMENT clause, the server first searches to see whether *the content files of the specified virtual document* meet the full-text search criteria. Then, as the server assembles the virtual document, the criteria in the SEARCH clause are applied to *the content files* associated with each component.

If you want to conduct fulltext searches on the assembled document components, you must create a file from the assembled document and associate it with the virtual document as a content file. Use the following procedure:

### **To index an assembled document as a whole:**

1. Assemble the document.
2. Print the document to a file.
3. Use Setfile or Setcontent to associate the resulting file with the virtual document as a content file.

Now you have a content file for the virtual document that represents the assembled document and that can be indexed.

## Querying registered tables

Registered tables are tables in the underlying RDBMS that have been registered with the repository. You must register any RDBMS table that you wish to query using DQL. Registering a table creates an object of type `dm_registered` for that table. (Refer to [Register, page 108](#), in the *Content Server DQL Reference Manual* for information about registering tables.)

To obtain information about the repository object that represents the table, specify the `dm_registered` type in the FROM clause of the SELECT statement. The server searches for the object representing the table. To query the table itself, specify the registered table's name in the FROM clause. When you query the table, the server searches the actual underlying RDBMS table—*not* the `dm_registered` object type.

## Referencing registered tables in queries

To avoid ambiguity in a query, it is often useful to reference a database table by its fully qualified name; that is, *owner\_name.table\_name*.

The owner name is the name of the person who created the table. Tables created by the system when a repository is created (for example, `dm_sysobject_r` or `dm_sysobject_s`) are owned by the repository owner. For these tables, Oracle and DB2 use the value in the repository owner's `user_db_name` attribute as the owner name. Sybase and MS SQL Server prefixes the names of all tables created by the repository owner with the alias `dbo`. This alias makes it possible to write applications that are portable across Sybase and MS SQL Server databases.

For application portability across all repositories, Documentum provides an alias, `dm_dbo`, that you can substitute for the repository owner's name in any fully qualified reference to registered tables. (For Sybase and MS SQL Server, when referencing registered tables, the aliases `dm_dbo` and `dbo` are equivalent.)

## Security controls

Access to a registered table is controlled by the object-level permissions and table permits defined for the table's `dm_registered` object. The object-level permissions must give you at least Browse access to the `dm_registered` object and the table permits must give you permission for the operation that you want to perform. For example, if you want to update a table, the table permits must grant you `DM_TABLE_UPDATE` permission. Table permits are defined for three user levels: owner, group, and world. (For complete

information about table permits and Documentum security, refer to [Chapter 11, Protecting Repository Objects](#), in the *Content Server Administrator's Guide*.)

Additionally, the user account under which Content Server is running must have the appropriate RDBMS permission to perform the requested operation on the specified table. (The actual name of this permission depends on your RDBMS.) For example, if you want to update the table, the server account must have permission in the RDBMS to update the table.

All three of these conditions must be met to gain access to a registered table. Even if a user has permission to access the underlying table through the RDBMS, the user will not have access through DQL unless the object-level permissions and table permits on the `dm_registered` object permit access. Similarly, a user might have the correct object-level permissions and table permits, but if the server's user account in the RDBMS does not have permission, then access is denied.

## Default object-level permissions and table permits

The REGISTER statement automatically sets the object-level permissions for a table's `dm_registered` object to default values. The statement also sets the default table permit to SELECT for the owner (group and world have no default table permit).

You can change the object-level permissions and table permits by setting the appropriate attributes directly. The *Content Server Administrator's Guide* provides more information about these attributes in [Chapter 11, Protecting Repository Objects](#). For information about how to set them, refer to [Set](#), page 416, in the *Content Server API Reference Manual*.

## Caching queries

Some queries return the same results every time you run the query. For example, a payroll application may ask for the names and social security numbers of the employees in the company. Although the query results may change over a long period of time, they may not change from week to week. Rather than incur the performance cost of rerunning the query that returns the users and social security numbers each time the payroll application executes, you can cache the query results on the client host. For information about persistent caching, refer to [Persistent client caches](#), page 45.

## Using execquery in an application

This section presents a model for processing the query results returned by an Execquery method in an application. The Execquery method allows you to execute queries of any length.

Because you can execute any DQL statement using the Execquery method, the model presented in this section shows you how to process queries when the application does not know whether the query returned any result objects.

**To process the results of an Execquery method (illustrated in [Figure 4–1, page 353](#)):**

1. If the dmAPIExec succeeded, use the Getlastcoll method to retrieve the return value from the Execquery method. Examine this value.
  - If the value is NULL, then the DQL query failed. Use the GetMessage or Listmessage method to look at the error messages.
  - If the value is a collection ID, then the DQL query succeeded. However, you should still issue a GetMessage or Listmessage method to retrieve any possible warnings. Even a successful query can return warnings.
2. If Getlastcoll returned a collection ID, examine the collection's `_count` attribute.
  - If `_count` is 0, then the query cannot return a result object (for example, BEGIN TRAN and COMMIT).
  - If `_count` is greater than 0, then the query can return a result object. (For example, SELECT statements can return result objects.)
  - The value of `_count` represents the number of attributes in each returned result object if any result objects are returned. For a SELECT statement, `_count` equals the number of values in the selected values list. For other statements returning results, the `_count` value is 1.
  - (Refer to individual statement descriptions in [Chapter 2, DQL Statements](#), in the *Content Server DQL Reference Manual* for a description of values returned by different statements.)
3. Use the `_count` value to loop through the attributes and retrieve their names, datatypes, and lengths to prepare the application to receive the results.
4. Execute the Next method to retrieve the results. Each execution retrieves one result object.

When Next returns a NULL, then either there are no more objects to return or there has been an error in the RDBMS. Check the collection's `_status` attribute. If this attribute's value is greater than DM\_WARNING, then Next returned a NULL due to an error rather than a lack of objects.

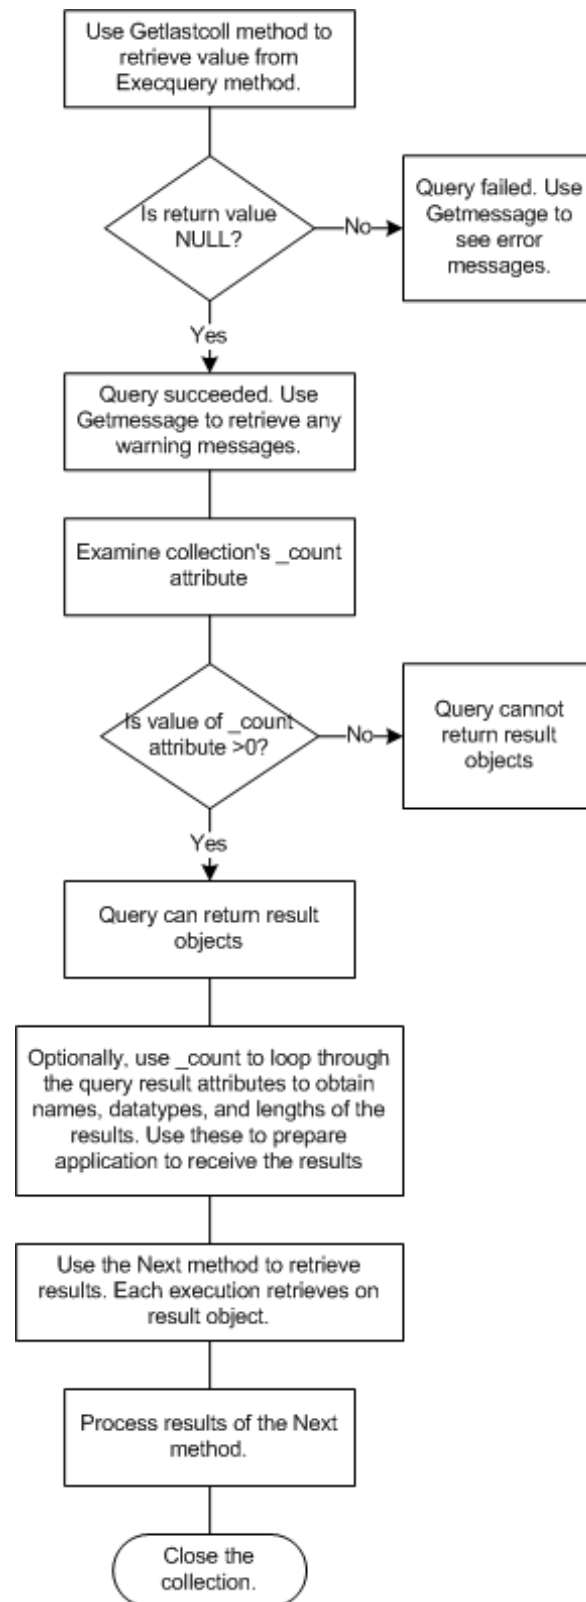
5. Process the results.

**Note:** *Important Note:* When you use the Get method to retrieve an attribute value from the collection, you must specify the attribute name in the method call in lowercase. All attribute names are returned in lowercase.

6. Repeat steps 4 and 5 until all the results are processed.
7. When you have no more results to process, close the collection.



Figure 4-1. Query processing



## Privileges, permissions, and queries

DQL query statements affect only those objects for which users have the appropriate permissions or privileges.

When a SELECT statement references a SysObject type (dm\_sysobject or any of its subtypes), by default, only those objects of the referenced type for which the user has at least Browse permission are retrieved. If the SELECT statement includes a SEARCH clause that accesses the object's content, then the user must have at least Read permission. However, the SELECT statement supports an option to allow you to specify another base permission level. For example, you might use a query that returns only objects for which a user has at least Relate permission.

If the object type specified in the SELECT statement is modified by the keyword PUBLIC, then the statement retrieves only objects that have a world permission of at least Browse (or Read for full-text searches) or objects that are owned by the user.

When a SELECT statement references a registered table, the user must have at least Browse permission on that registered table for the query to succeed. In addition, Content Server must have the SELECT privilege in the RDBMS.

The Superuser user privilege bypasses all Documentum security checks. However, Content Server must still have the SELECT privilege in the RDBMS for a superuser to query registered tables.

You must have Write permission to update an object with the UPDATE...OBJECT statement. You must have Delete permission for an object to delete it using the DELETE...OBJECT statement.

## Using DQL Hints

This appendix describes how DQL hints are implemented and how to use them in your queries.

- [General guideline for all](#), page 355
- [SQL\\_DEF\\_RESULT\\_SET N](#), page 356
- [FORCE\\_ORDER](#), page 357
- [RETURN\\_TOP N](#), page 358
- [OPTIMIZE\\_TOP N](#), page 360
- [FETCH\\_ALL\\_RESULTS N](#), page 360
- [OPTIMIZATION\\_LEVEL level\\_1 level\\_2](#), page 361
- [UNCOMMITTED\\_READ](#), page 361
- [IN and EXISTS](#), page 362
- [ROW\\_BASED](#), page 363
- [FTDQL and NOFTDQL](#), page 364
- [Including multiple hints limiting rows returned](#), page 365
- [Passthrough hints](#), page 365

### General guideline for all

Database hints affect how a query is executed. When deciding whether to use a hint, it is recommended that you execute the query with and without the hint. Compare the generated SQL queries and compare the response time and the resource use to determine whether the hint is helpful.

## SQL\_DEF\_RESULT\_SET N

The SQL\_DEF\_RESULT\_SET N is most useful on a SQL Server database. On a SQL Server database, including SQL\_DEF\_RESULT\_SET in a query causes the RDBMS to use default result sets to return the query results rather than a cursor. Using default result sets is the way that Microsoft recommends accessing a SQL Server database.

On other databases, it only controls how many rows are returned, and is identical to the RETURN\_TOP hint.

On SQL Server, using default result sets allows the RDBMS server to perform fewer logical read operations. [Table A-1, page 356](#) shows some test examples:

**Table A-1. Comparison of logical reads with and without default result sets**

Query	Number of returned rows	Number of logical reads without DRS	Number of logical reads with DRS
SELECT object_name FROM dm_document WHERE r_object_id= '0900014c8000210b'	1	36	5
SELECT object_name FROM dm_document WHERE r_creator_name= 'user2'	100	1600	700
SELECT object_name FROM dm_document WHERE r_creator_name= 'user1'	2000	22500	500

However, the reduction in I/O may be offset by higher CPU costs. Additionally, because using default result sets requires Content Server to buffer the results of a query in memory, higher memory use is a result. (When using default result sets, two queries cannot be open simultaneously. If a query is issued, before another can be issued from the same session, all results from the first must be processed. To work around this limitation, Content Server internally buffers the results of queries that are executed using default result sets.)

To determine whether using SQL\_DEF\_RESULT\_SETS is useful for a query, run the query with and without the hint and trace the results using the SQL Server Profiler.

Setting N to 0 causes all results to be returned. Setting N to an integer value limits the number of rows returned to that value.

Using SQL\_DEF\_RESULT\_SETS is recommended if:

- You are running on a SQL Server database.
- The query returns limited result sets.
- The query is a commonly executed query.
- Tests show that using the hint greatly improves query performance.

## FORCE\_ORDER

The `FORCE_ORDER` hint controls the order in which the tables referenced in the query's `FROM` clause are joined. The tables may be RDBMS tables or object type tables.

Oracle and SQL Server implement this hint at the statement level. For example, suppose you issue the following DQL:

```
SELECT object_name
FROM DM_SYSOBJECT ENABLE (FORCE_ORDER)
```

The generated SQL for Oracle is:

```
select /*+ ORDERED */ object_name from dm_sysobject_s;
```

The generated SQL for SQL Server is:

```
select object_name from dm_sysobject_s
OPTIONS (FORCE_ORDER)
```

Sybase implements the hint at the session level. If you include the hint, Content Server uses the `ct_options` function to set the `CS_OPT_FORCEPLAN` variable to true and unsets it after the query executes.

Using `FORCE_ORDER` may not ensure that you obtain a particular join order. When you examine the SQL query generated by a DQL query, you may find there are more tables in the `FROM` clause than are present in the DQL query. This is because Content Server often uses views to generate the SQL for many DQL queries and may also add ACL tables to the queries for security checks.

If you are considering using this hint, run the query without the hint and obtain the generated SQL statement and the execution plan. If the join order in the generated query appears incorrect and you believe that joining the tables in the order they appear in the DQL query will result in better performance, run the query with `FORCE_ORDER`. Compare the results to the query execution results without the hint.

If you use this hint, it is recommended that you retest the query with and without the hint occasionally to ensure that the hint is still useful. Database changes can make the plan chosen by the hint incorrect.

# RETURN\_TOP N

The RETURN\_TOP N hint limits the number of rows returned by a query.

## Database-specific implementations

The following sections describe how RETURN\_TOP is handled for individual databases.

### SQL Server

On SQL Server, using this hint results in fewer database rows being touched by a query. The internal behavior of query on SQL Server is:

1. The query executes on the database (touching whatever tables are required) and generates a list of keys or lookup IDs inside the tempdb.
2. Each time a row is fetched from the cursor, the lookup ID accesses the actual table to return the full result set with all columns.

When you include the RETURN\_TOP hint, the second step is executed only as many times as the hint directs. For example, if the hint is RETURN\_TOP 20, then only 20 rows are returned and the table is accessed only 20 times.

Including RETURN\_TOP adds the SQL Server TOP hint to the generated SQL statement. For example, suppose you issue the following DQL statement:

```
SELECT user_name FROM dm_user ENABLE (RETURN_TOP 3)
```

The generated SQL for SQL Server is:

```
select TOP 3 user_name from dm_user_s
```

### Subqueries and the hint

If the query includes a subquery, the hint is not applied to the subquery.

### DB2

On a DB2 database, including RETURN\_TOP adds the FETCH FIRST N ROWS ONLY hint to the generated SQL statement. For example, the following DQL query

```
SELECT user_name FROM dm_user ENABLE (RETURN_TOP 3)
```

is translated to the following SQL statement

```
select user_name from dm_user_s
FETCH FIRST 3 ROWS ONLY
```

It is recommended that you use RETURN\_TOP in conjunction with OPTIMIZE\_TOP.

The previous example becomes:

```
SELECT user_name FROM dm_user
ENABLE (RETURN_TOP 3, OPTIMIZE_TOP 3)
```

The statement generates the following SQL:

```
select user_name from dm_user_s
FETCH FIRST 3 ROWS ONLY OPTIMIZE FOR 3
```

## Oracle and Sybase

If you include RETURN\_TOP in a query running against Oracle or Sybase, the returned results are not limited at the database level, but by Content Server itself. Consequently, using RETURN\_TOP on Oracle or Sybase results in no database performance gains, but may reduce network traffic.

## Effects of a SEARCH clause

If the DQL query includes a SEARCH clause, to limit results to documents that are indexed, the implementation of RETURN\_TOP depends on whether the searched documents are public or not.

If all the searched documents are public, the returned results are limited at the fulltext index level. If the searched documents are not all public, then the limits are imposed when Content Server performs the security checking on the returned results.

## Recommended use

Use RETURN\_TOP when:

- Only a particular number of rows is needed from the database.
- The application accepts ad hoc queries from users and you want to limit the potential damage an unbounded query might cause.

Using RETURN\_TOP if the query sorts the results or includes the DISTINCT keyword reduces the efficiency of the hint.

## OPTIMIZE\_TOP N

The OPTIMIZE\_TOP N hint directs the database server to return the first N rows returned by a query quickly. The remaining rows are returned at the normal speed.

On SQL Server and DB2, you can include an integer value (as N) to define how many rows you want returned quickly. On Oracle, the number is ignored. The OPTIMIZE\_TOP hint is not available on Sybase.

For example, suppose you issue the following DQL query:

```
SELECT r_object_id FROM dm_sysobject
ENABLE (OPTIMIZE_TOP 4)
```

On SQL Server, the generated SQL statement is:

```
select r_object_id from dm_sysobject_s
OPTION (FAST 4)
```

On DB2, the generated SQL statement is:

```
select r_object_id from dm_sysobject_s
OPTIMIZE FOR 4 ROWS
```

On Oracle, the generated SQL statement is:

```
select /*+ OPTIMIZE_TOP */ r_object_id
from dm_sysobject_s
```

OPTIMIZE\_TOP is recommended for use:

- When you want to return all the results but want the first few rows more quickly
- With the RETURN\_TOP hint, to optimize the return of specified rows
- When the execution plan chosen for query is a bad plan and there is no obvious solution

Using OPTIMIZE\_TOP if the query sorts the results or includes the DISTINCT keyword reduces the efficiency of the hint.

## FETCH\_ALL\_RESULTS N

The FETCH\_ALL\_RESULTS N hint fetches all the results from the database immediately and closes the cursor. The hint does not affect the execution plan, but may free up database resources more quickly.

To fetch all the results, set N to 0.

On SQL Server, it is recommended that you use SQL\_DEF\_RESULT\_SETS instead of the FETCH\_ALL\_RESULTS hint. SQL\_DEF\_RESULT\_SETS provides the same benefits and is the recommended way to access SQL Server databases.



Use `FETCH_ALL_RESULTS` if you want to reduce the resources used by the database server by quickly closing cursors. On SQL Server, try `FETCH_ALL_RESULTS` if using `SQL_DEF_RESULT_SETS` did not improve query performance.

## SQL Server, the hint, and subqueries

If you use this hint on query against a SQL Server database and the query includes a subquery, the hint is not applied to the subquery.

## OPTIMIZATION\_LEVEL level\_1 level\_2

Use the `OPTIMIZATION_LEVEL` hint against a DB2 database when you want to change the optimization level for a particular query. The level you set in the `level_1` argument is used for the current query and after the query is finished, Content Server sets the optimization level to the level specified in the `level_2` argument.

Using this hint generates a SQL trace that looks like the following:

```
set current optimization level level_1
run query
set current optimization level level_2
```

This hint is ignored for all databases except DB2.

For more information or assistance on the `level_1` and `level_2` parameters, please contact your DB2 database vendor.

## UNCOMMITTED\_READ

Use the `UNCOMMITTED_READ` hint in read only queries, to ensure that the query returns quickly even if another session is holding locks on the tables queried by the read only query.

This hint is useful only on SQL Server, DB2, and Sybase databases.

## IN and EXISTS

IN and EXISTS are mutually exclusive hints that you can use in a standard WHERE clause, in a repeating attribute predicate that contains a subquery. These hints may not be used in a WHERE clause in an FTDQL query.

The syntax is:

```
WHERE ANY [IN|EXISTS] attribute_name (subquery)
```

For example:

```
SELECT "r_object_id" FROM "dm_document"
WHERE ANY IN "authors" IN
(SELECT "user_name" FROM "dm_user")
```

If you do not include either IN or EXISTS explicitly, when Content Server translates the query, it includes one or the other automatically. Which hint Content Server chooses to include is determined by the indexes present in the repository, the attributes referenced in the query, and other factors.

The queries generated by each option are different and perform differently when executed. For example, here is the generated SQL statement for the query in the previous example:

```
select all dm_sysobject.r_object_id
from dm_sysobject_sp dm_sysobject
where (dm_sysobject.r_object_id in
      (select r_object_id from dm_sysobject_r
       where authors in
        (select all dm_user.user_name
         from dm_user_sp dm_user)))
and (dm_sysobject.i_has_folder=1 and
     dm_sysobject.i_is_deleted=0)
```

If the DQL query used the EXISTS hint, the generated SQL statement would look like this:

```
select all dm_sysobject.r_object_id
from dm_sysobject_sp dm_sysobject
where (exists (select r_object_id
              from dm_sysobject_r
              where dm_sysobject.r_object_id = r_object_id
                and authors in
                 (select all dm_user.user_name
                  from dm_user_sp dm_user )))
and (dm_sysobject.i_has_folder=1 and
     dm_sysobject.i_is_deleted=0)
```

If you feel that a query that references a repeating attribute predicate that uses a subquery is performing badly, run the query and examine the generated SQL statement to determine which hint Content Server is adding to the generated SQL and note the performance time. Then, rewrite the DQL query to include the hint that the server isn't choosing. For example, if the server is choosing to add the EXISTS hint to the generated

SQL statement, rewrite the DQL query to include the IN hint. Then, rerun the query to determine if the performance improves.

## ROW\_BASED

The ROW\_BASED hint changes both the way query results are returned and the syntax rules for the query itself.

### Effects on returned results

The ROW\_BASED hint forces Content Server to return query results in a row format, as opposed to an object-based format. The difference is most readily apparent if the query selects values from a repeating attribute. In an object-based format, the server returns all selected repeating values for a particular object in one query result object. In a row-based format, the server returns each selected repeating attribute value in a separate query result object.

For example, by default the following query returns results in an object-based format:

```
SELECT "r_object_id","title","authors" FROM "dm_document"
WHERE "subject"='new_book_proposal'
```

That query returns the authors values as a list of authors in one query result object for each returned object. [Table A-2, page 363](#), shows how the results are returned. Each row in the table represents one query result object and each column is one attribute in the query result objects.

**Table A-2. Example of object-based query results**

r_object_id	title	authors
090000015973a2fc	Our Life and Times	Jennie Doe Carol Jones Hortense Smith
0900000123ac12f6	Life of an Amoeba	James Does Jules Doe

There is one query result object (one row) for each object returned by the query. Now, add the ROW\_BASED hint to the query:

```
SELECT "r_object_id","title","authors" FROM "dm_document"
WHERE "subject"='new_book_proposal'
ENABLE (ROW_BASED)
```

[Table A-3, page 364](#), illustrates how Content Server returns the results for that query.

**Table A-3. Example of row-based query results**

r_object_id	title	authors
090000015973a2fc	Our Life and Times	Jennie Doe
090000015973a2fc	Our Life and Times	Carol Jones
090000015973a2fc	Our Life and Times	Hortense Smith
0900000123ac12f6	Life of an Amoeba	James Doe
0900000123ac12f6	Life of an Amoeba	Jules Doe

There is one query result object (one row) for each repeating attribute value returned. The returned repeating attribute values for each object are not aggregated into one result object.

## Effects on query syntax rules

In addition to the changes the hint causes in the returned results, including the hint also changes some of the syntax rules for a query. The hint affects:

- Query syntax rules when a repeating attribute is a selected attribute  
[Repeating attributes, page 122](#), describes how the hint affects query syntax rules when a repeating attribute is selected.
- The use of repeating attributes in expressions in WHERE clauses  
[Using repeating attributes in qualifications, page 143](#), describes how the hint affects the use of repeating attributes in WHERE clause qualifications.
- The behavior of the asterisk as a selected value  
[The Asterisk \(\\*\) as a selected value, page 130](#), describes how the hint affects the returned values for an asterisk.

## FTDQL and NOFTDQL

The FTDQL hint supports the FTDQL functionality. If you include this hint in query, Content Server attempts to execute the query as an FTDQL query. If the remaining syntax in the query conforms to the required syntax for an FTDQL query, the query is executed as an FTDQL query. If the syntax does not conform to FTDQL query rules, an error is returned.

If you include the NOFTDQL hint, the query is run as a standard SELECT query whether the syntax satisfies the FTDQL query rules or not.

## Including multiple hints limiting rows returned

If you include more than one hint that limits the rows returned by a query, the number of rows returned is the least number of rows defined in an included hint. For example, suppose you issue the following DQL statement:

```
SELECT object_name FROM dm_document ENABLE
(FETCH_ALL_RESULTS 10, RETURN_TOP 5)
```

Content Server returns 5 rows for the query because the RETURN\_TOP hint is the most constraining hint in the list.

On SQL Server, if the list includes SQL\_DEF\_RESULT\_SET, the query is always executed using default result sets regardless of where the hint falls in the list of hints. For example, suppose you issue the following statement:

```
SELECT object_name FROM dm_document ENABLE
(SQL_DEF_RESULT_SET 10, RETURN_TOP 5)
```

The query executes using default result sets but returns only 5 rows.

## Passthrough hints

Passthrough hints are hints that are passed to the RDBMS server. They are not handled by Content Server.

SQL Server and Sybase have two kinds of hints: those that apply to individual tables and those that apply globally, to the entire statement. To accommodate this, you can include passthrough hints in either a SELECT statement's source list or at the end of the statement. The hints you include in the source list must be table-specific hints. The hints you include at the end of the statement must be global hints. For example, the following statement includes passthrough hints for Sybase at the table level and the statement level:

```
SELECT "r_object_id" FROM "dm_document"
WITH (SYBASE('NOHOLDLOCK'))
WHERE "object_name"='test' ENABLE (FORCE_PLAN)
```

For DB2 and Oracle, include passthrough hints only at the end of the SELECT statement.

## Syntax

To include a passthrough hint, you must identify the database for which the hint is valid. To identify the target database, keywords precede the hints. The valid keywords are: ORACLE, SQL\_SERVER, SYBASE, and DB2. For example, the following statement includes passthrough hints for SQL Server:

```
SELECT "r_object_id" FROM "dm_document"  
WHERE "object_name" ='test'  
ENABLE SQL_SERVER('ROBUST PLAN','FAST 4',  
  'ROBUST PLAN')
```

For portability, you can include passthrough hints for multiple databases in one statement. The entire list of hints must be enclosed in parentheses. The syntax is:

```
(database_hint_list {,database_hint_list})
```

where *database\_hint\_list* is:

```
db_keyword('hint' {,'hint'})
```

*db\_keyword* is one of the valid keywords identifying a database. *hint* is any hint valid on the specified database.

For example:

```
SELECT object_name FROM dm_document doc dm_user u  
WHERE doc.r_creator_name = u.user_name  
ENABLE (ORACLE('RULE','PARALLEL'),  
SYBASE('AT ISOLATION READ UNCOMMITTED'),  
SQL_SERVER('LOOP JOIN','FAST 1'))
```

## Error handling and debugging

It is possible to execute a DQL statement that the server considers correct DQL but is incorrect at the RDBMS level if you incorrectly specify a passthrough hint. If this occurs, any errors returned are returned by the RDBMS server, not Content Server. Not all databases return errors on database hints. For example, Oracle does not return an error if a database hint is incorrect—it simply ignores the hint. For the other databases, the error messages may be difficult to decipher.

To debug problems with queries containing database hints, you can turn on SQL tracing. By default, Content Server runs with the last SQL trace option turned on. You can turn on full SQL tracing using a Trace method or the SET\_OPTIONS administrative function. The SET\_OPTIONS administration method is described in [SET\\_OPTIONS](#), page 312.

# Implementing Java Evaluation of Docbasic Expressions

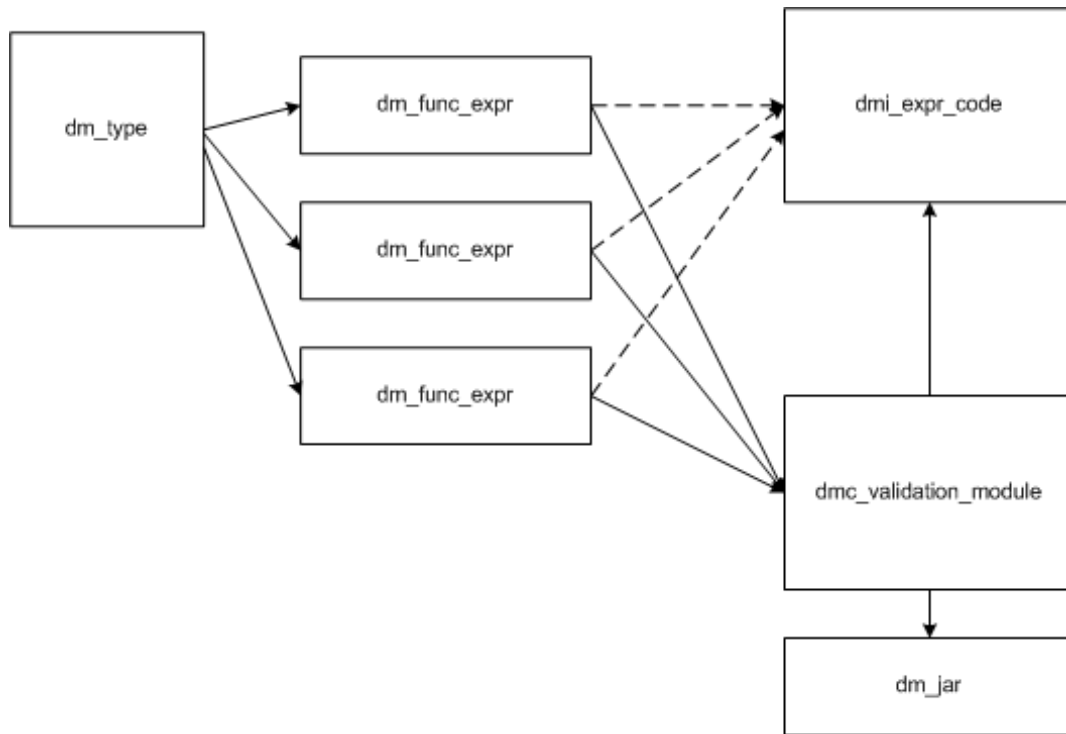
This appendix describes how to implement Java evaluation of Docbasic expressions defined for check constraints or value assistance in the data dictionary. The appendix includes the following topics:

- [Overview of Docbasic expression handling, page 367](#)
- [Migrating Docbasic expressions to Java, page 369](#)
- [Disabling or re-enabling Java evaluation, page 371](#)
- [Docbasic expression components support, page 372](#)

## Overview of Docbasic expression handling

This section describes how Content Server handles Docbasic expressions defined for check constraints or value assistance in the data dictionary. It discusses how such expressions are stored in the repository and how that implementation is enhanced if Java evaluation is implemented for the expressions.

[Figure B-1, page B-368](#), illustrates the architecture.

**Figure B-1. Architectural implementation of Docbasic expressions for object types**

For each object type that has Docbasic expressions defined as check constraints or in conditional value assistance for one or more attributes, Content Server creates one expr code object and a number of func expr objects. Each func expr objects records one Docbasic expression. The expr code object contains the compiled code and pcode that implements the expressions recorded in the func expr objects. The routine\_id attribute in the func expr objects points to the expr code object.

If you migrate an expression to Java, the migration method compiles the expression into Java code. If the expression is the first expression compiled into Java for the object type, the migration method also creates a dmc\_jar object and a dmc\_validation\_module object. The compiled code is stored with the dmc\_jar object. The jar object is related to the validation module object through a relationship named dmc\_module\_to\_jar.

The validation module object points back to the associated expr code object through an attribute and is related to the func expr objects using dmc\_validation\_relation objects. Validation relation objects are a subtype of dm\_relation. The name of the relationship represented by validation relation objects is dmc\_expr\_to\_module. Both the validation module and validation relation objects have attributes used to manage the Java evaluation.

If additional Docbasic expressions are added later for an object type, the migration method updates the validation module and jar objects.



Migrating expressions to Java is currently a manual operation. EMC Documentum provides two migration methods, described in [Migrating Docbasic expressions to Java, page 369](#). You can migrate expressions already defined in the data dictionary. You can migrate new expressions after the expressions are defined using an ALTER TYPE or CREATE TYPE statement.

## Migrating Docbasic expressions to Java

There are two methods that create compiled Java code for Docbasic expressions stored with expr code objects. These methods are:

- `dmc_MigrateDbExprsToJava`  
Use this method if you want to migrate expressions defined for multiple object types.
- `dmc_MigrateDbExprsToJavaForType`  
Use this method if you want to migrate expressions added to a single object type (or the type's attributes) using an ALTER TYPE or if you want to migrate expressions defined while creating a new type using CREATE TYPE.

Arguments for both methods let you select which Docbasic expressions to migrate to Java code. You can choose to compile only new Docbasic expressions, or those that failed compilation previously, or all expressions within the scope of the method.

The methods are executed using either a DFC script or using a DO\_METHOD administration method. The DO\_METHOD syntax is:

```
dmAPIGet ("apply, session, NULL, DO_METHOD, METHOD, S, migration_method_name, SAVE_RESULTS, B, T, ARGUMENTS, S, argument_list
```

*migration\_method\_name* is either `dmc_MigrateDbExprsToJava` or `dmc_MigrateDbExprsToJavaForType`.

Use a space to separate arguments in *argument\_list*. [Table B-1, page 369](#), lists the valid arguments for *argument\_list*.

**Table B-1. Arguments for the Docbasic expression migration methods**

Argument	Description
-docbase <i>docbase_name</i>	Name of the repository that contains the expressions you want to migrate. This must be a 5.3 (or later) repository.
-user <i>user_name</i>	Name of a user who has at least Sysadmin privileges in the specified repository.

Argument	Description
-ticket <i>login_ticket</i>	String containing a login ticket.  <b>Note:</b> For information about login tickets and generating those tickets, refer to the <i>Content Server Administrator's Guide</i> .
-select <i>selection_choice</i>	Identifies which expressions you wish to migrate. <i>selection_choice</i> is one of: <ul style="list-style-type: none"> <li>• <b>new_only</b>, includes only new expressions, that is, those for which migration has not been previously attempted</li> <li>• <b>new_and_failed</b>, includes all new expressions and those for which migration previously failed</li> <li>• <b>new_and_disabled</b>, includes all new expressions and those that do not have a corresponding enabled Java version</li> <li>• <b>all</b>, includes all expressionsAny existing Java code is replaced if currently migrated expressions are recompiled.</li> </ul> <p>For <code>dmc_MigrateDbExprsToJava</code>, expressions for all object types are considered for inclusion. For <code>dmc_MigrateDbExprsToJavaForType</code>, only expressions defined for the object type identified in the <i>type</i> argument are considered for inclusion.</p>
-maxerrs <i>value</i>	Maximum number of errors that can occur before code generation stops. By default, there is no maximum.  Setting this argument is recommended only if there are large numbers of expressions to compile.
-listExprsOnly <i>value</i>	Whether to actually perform the migration or only create a file that records the expressions selected for migration. Setting value to true directs the method to only create the file. Setting the value to false directs the method to migrate the expressions.  If you set this argument to true, you must include the <code>SAVE_RESULTS</code> argument in the <code>DO_METHOD</code> command line.
-type <i>type_name</i>	Name of the object type. Include this argument only if you are executing the <code>dmc_MigrateDbExprsToJavaForType</code> method. Use the type's internal name; for example: <code>dm_document</code> .

It is strongly recommended that you include the `SAVE_RESULTS` argument on the `DO_METHOD` command line regardless of the `-listExprsOnly` setting. The content of the generated results document varies depending on the circumstances of the method's execution:

- If `-listExprsOnly` is set to true, the results document contains a list of all expressions selected for migration.
- If `-listExprsOnly` is false and non-fatal errors occur during migration, the results document contains the errors.
- If `-listExprsOnly` is false and fatal errors occur during migration, the results document contains the exception stack trace embedded in the HTML error message.
- If `-listExprsOnly` is false and no errors occur during migration, the results document contains only the text "There were no errors."

Both methods return 0 unless a fatal error occurred during execution. If a fatal error occurs, the methods return 1. Fatal errors are reported using the standard method server mechanism.

## Disabling or re-enabling Java evaluation

It is possible to disable Java evaluation of a particular Docbasic expression. If you disable Java evaluation of an expression, the Docbasic library is used to evaluate the expression, rather than the compiled Java code.

To disable Java evaluation, execute the `dmc_SetJavaExprEnabled` method. This method has an argument that turns evaluation of a specified expression on or off. The method is executed by a DFC script or a `DO_METHOD` administration method. Here is the syntax for the `DO_METHOD` call:

```
apply, session, NULL, DO_METHOD, METHOD, S, dmc_SetJavaExprEnabled, SAVE_RESULTS, B, T, ARGUMENTS, S, argument_list
```

[Table B-2, page 371](#), lists the arguments for `dmc_SetJavaExprEnabled`. All arguments for this method are required. Use a space to separate arguments in *argument\_list*.

**Table B-2. `dmc_SetJavaExprEnabled` arguments**

Argument	Description
<code>-docbase repository_name</code>	Name of the repository that contains the Java expression implementation you want to disable or re-enable.
<code>-user user_name</code>	Name of the user performing the operation. Use the user's login name. The user must have at least Sysadmin privileges in the repository.

Argument	Description
-ticket <i>login_ticket</i>	A string containing a login ticket generated for the specified user.
-func_expr_id <i>func_expr_obj_id</i>	Object ID of the dm_func_expr object representing the Docbasic expression
-enable <i>value</i>	Set this to true to enable the Java evaluation of the Docbasic expression identified in the -func_expr_id argument. Or, set it to false, to disable Java evaluation of that expression.

## Docbasic expression components support

This section lists and briefly describes the Docbasic components, such as operators, functions, and constants, that are supported for Java migration.

### Operators

[Table B-3, page 372](#), lists the Docbasic operators for which Java evaluation is supported.

**Table B-3. Docbasic operators supported by Java evaluation**

Operator	Operation performed
&	String concatenation
*	Multiplication
+	Addition or concatenation
-	Subtraction
/	Floating point division
<	Comparison (less than)
<=	Comparison (less than or equal to)
<>	Comparison (not equal to)
=	Comparison (equal to)
>	Comparison (greater than)
>=	Comparison (greater than or equal to)
\	Integer division

Operator	Operation performed
^	Exponentiation
And	Logical or binary conjunction
Eqv	Logical or binary equivalence (not Xor)
Imp	Logical or binary implication
Is	Compares two operands and returns true if they refer to the same object; otherwise, returns false.  Because objects are not supported, using this operator always results in a type mismatch error at runtime.
Like	Compares two strings and returns true if the expression matches the given pattern; otherwise, returns false.
Mod	Returns the remainder of the <i>expression_1/expression_2</i> as whole number
Not	Returns either a logical or binary negation of an expression
Or	Logical or binary disjunction on two expressions
Xor	Exclusive Or operation

## Supported functions for Java evaluation

[Table B-4, page 373](#), lists the Docbasic functions for which Java evaluation is supported.

**Table B-4. Docbasic functions supported for Java evaluation**

Function	Description
abs(expr)	Returns the absolute value of expr
Asc(text\$)	Returns the integer containing the numeric code for the first character of the text\$ (0-255)
AscB or AscW	Returns the byte or wide-character equivalents of Asc
Atn(number)	Returns the inverse tangent of number
CBool(expr)	Converts the expr to a Boolean
CDate(expr) or CDate(expr)	Converts the expr to a Date  <b>Note:</b> If expr is a string, it is expected that the date is in canonical form unless a specific date format string was passed to the validation method.

Function	Description
CDBl(expr)	Converts expr to a Double
Choose(index, expr1, expr2,...exprN)	Returns the expression whose position in the list of expr arguments matches the value specified index. If none is found, returns null.
Chr(code) or Chr\$(code)	Returns the character whose value is code  The Java implementation returns a String variant in both functions.
CInt(expr)	Converts expr to a Docbasic integer.  The Java implementation converts expr to a Java short.
CLng(expr)	Converts expr to a Docbasic long  The Java implementation converts expr to a Java int.
Cos(angle)	Returns the cosine of angle where angle is assumed to be expressed in radians
CSng(expr)	Converts expr to a Docbasic Single  The Java implementation converts expr to a Java float.
CStr(expr)	Converts expr to a String
CVar(expr)	Converts expr to a Variant
Date	Returns the current date as a Date Variant
Date\$	Returns the current date as a String  The Java implementation will always be in canonical form or in the format passed to the validation method.
DateAdd(interval\$, increment&, date)	Increments the date field identified in interval\$ of the given date by the amount specified in interval&
DateDiff(interval\$, date1, date2)	Returns a Date variant representing the number of given time intervals between date1 and date2
DatePart(interval\$, date)	Returns an Integer representing a specific part of a date/time expression
DateSerial(year, month, day)	Returns a Date variant representing the specified date
DateValue(dateString\$)	Returns a Date variant representing the date specified by dateString\$
Day(date)	Returns the day of the month represented by date

Function	Description
Exp(val)	Returns the value of $e$ raised to the power of val
Fix(number)	Returns the integer part of number
Format or Format\$(expr, userFormat)	Returns a string formatted to user specification (supports both named and unnamed formats)
Hex[\$](number)	Returns a String containing the hexadecimal equivalent of number
Hour(time)	Returns the hour of the day encoded in time
IIf(condition, TrueExpression, FalseExpression)	Returns TrueExpression if condition is true; otherwise, returns FalseExpression
InStr([start,] search, find [,compare])	Returns the first character position of string "find" within string "search"
InStrB([start,] search, find [,compare])	Returns the first character position of string "find" within string "search"
Int(number)	Returns the integer part of number
IsDate(expr)	Returns true if expr can be legally converted to a date; otherwise, returns false
IsNull(expr)	Returns true if expr is a Variant variable that contains no valid data; otherwise, returns false
IsNumeric(expr)	Returns true if expr can be converted to a number; otherwise, returns false
Item\$(text\$, first, last [,delimiters\$])	Returns all items between first and last within the specified formatted text list
ItemCount(text\$ [,delimiters\$])	Returns an integer containing the number of items in the specified delimited list
LBound(arrayVar [,dimension])	Returns an integer containing the lower bound of the specified dimension of the specified array variable
LCase[\$](text)	Returns the lowercase equivalent of text
Left[\$](text, nChars) or LeftB[\$](text, nChars)	Returns the leftmost nChars characters (Left and Left\$) or bytes (LeftB and LeftB\$) from text  <b>Note:</b> Use LeftB with caution. Multibyte characters can be split in the middle, resulting in the construction of illegal strings.
Len(expr or lenB(expr))	Returns the number of characters in expr or the number of bytes required to store expr

Function	Description
Line\$(text\$, first [,last])	Returns a string containing a single line or a group of lines between first and last
LineCount(text\$)	Returns an integer representing the number of lines in text\$
Log(number)	Returns a double representing the natural logarithm of number
LTrim[\$](text)	Returns text with the leading spaces removed
Mid[\$](text, start [,length]) or MidB[\$](text, start [,length])	Returns a substring of the specified text, beginning with start for length number of characters or bytes  <b>Note:</b> Use MidB with caution. Multibyte characters can be split in the middle, resulting in the construction of illegal strings.
Minute(time)	Returns the minute of the day encoded in the time argument
Month(date)	Returns the month of the date
Now	Returns a date variant representing the current date and time
Oct(number) or Oct\$(number)	Returns a string containing the octal equivalent of the specified number
Random(min,max)	Returns a long value greater than or equal to min and less than or equal to max
Right[\$](text,nChars) or RightB[\$](text,nChars)	Returns the rightmost nChars characters or bytes from the text.  RightB and RightB\$ are used to return byte data from strings containing byte data.  <b>Note:</b> Use RightB with caution. Multibyte characters can be split in the middle, resulting in the construction of illegal strings.
Rnd[(number)]	Returns a random Single number between 0 and 1
RTrim[\$](text)	Returns text with the trailing spaces removed
Second(time)	Returns the second of the day encoded in the time argument
Sgn(number)	Returns an integer indicating whether a number is less than, greater than, or equal to 0
Sin(angle)	Returns a double value specifying the sine of angle



Function	Description
Space[\$](NumSpaces)	Returns a string containing the specified number of spaces
Sqr(number)	Returns a double representing the square root of number
Str[\$](number)	Returns a string representation of number
StrComp(string1, string2 [,compare])	Returns an integer indicating the result of comparing string1 and string2
String[\$](number, [CharCode   text\$])	Returns a string of length number consisting of a repetition of the specified filler character
Switch(condition1, expr1 [,condition2,expr2... [condition7,expr7]])	Returns the expression corresponding to the first true condition.
Tan(angle)	Returns a double representing the tangent of angle
Time/Time\$	Returns the system time as a string or as a Date variant
	In the Java implementation, the time is in canonical form..
Timer	Returns a Single representing the number of seconds that have elapsed since midnight
TimeSerial(hour, minute, second)	Returns a Date variant representing the given time with a date of zero
TimeValue(time_string\$)	Returns a date variant representing the time contained in time_string\$
	The date must be in canonical form or in a format that conforms to the format string passed to the validation method.
Trim[\$](text)	Returns a copy of text with leading and trailing spaces removed
TypeName(varname)	Returns the type name of the specified variable
UBound(arrayVar [,dimension])	Returns an integer containing the upper bound of the specified dimension of the specified array variable
UCase[\$](text)	Returns the uppercase equivalent of the specified string
Val(numberString)	Converts a given string expression to a number
Vartype(variable)	Returns an integer representing the type of data in variable

Function	Description
Weekday(date)	Returns an integer value representing the day of the week given by date  Sunday is 1, Monday is 2, and so forth.
Word\$(text\$, first [,last])	Returns a string containing a single word or sequence of words between first and last
WordCount(text\$)	Returns an integer representing the number of words in text\$
Year(date)	Returns the year of the date encoded in the date argument  The value returned is between 100 and 9999, inclusive.

## Unsupported functions for Java evaluation

The functions listed in [Table B-5, page 378](#) are not supported for Java evaluation. For detailed information about these functions, refer to the Docbasic documentation.

**Table B-5. Docbasic functions not supported for Java evaluation**

Names of unsupported functions			
CCur(expr)	EOF	GetSession	Pmt(Rate, NPer, Pv, Fv, Due)
Clipboard\$	Erl	GetSetting([app-name], section, key[, default])	PPmt(rate, Per, Nper, Pv, Fv, Due)
Command or Command\$	Err	Input, Input\$, InputB, and InputB\$(...)	PrinterGetOrientation
CreateObject(ClassID)	Error and Error\$(errNum)	InputBox and InputBox\$	PrintFile(filename\$)
CurDir and CurDir\$(drive)	External(pathOrObjID)	IPmt(Rate, Per, Nper, Pv, Fv, Due)	Pv(reate, NPer, Pmt, Fv, Due)
CVErr	FileAttr(fileNum, attr)	IRR(ValueArray(), Guess)	Rate(NPer, Pmt, Pv, Fv, Due, Guess)
DDB(Cosot, Salvage, Life, Period)	FileDateTime(filename)	IsError(expr)	ReadIni\$(section\$, item\$[, filename\$])

Names of unsupported functions			
DDEInitiate(application\$, topic\$)	FileExists(filename)	IsMissing(variable)	SaveFilename\$([[title\$[, extentions\$]])]
DDERequest[\$](channel, DataItem\$)	FileLen(filename)	Loc(filename)	Seek(filename)
Dir[\$]([file-spec\$[, attributes]])	FileParse(filename, op)	Lof(filename)	SetAppInfo(section\$, entry\$, string\$, filename\$)
Disk-Free&([drive\$])	FileType(filename)	MacID(text\$)	Shell(command\$[, WindowStyle])
dmAPIExec	FreeFile	MIRR(valArray, rate, reinvestRate)	ShellSync(command)
dmAPIGet	Fv(Rate, Nper, Pmt, Pv, Due)	MsgBox(msg)	Spc(numspaces)
dmAPISet	GetAppInfo(section\$, entry\$, filename\$)	NPer(rate, Pmt, Pv, Fv, Due)	SYD(cost, salvage, life, period)
dmExit	GetAttr(filename\$[, class\$])	Npv(rate, valArray)	Tab(column)
Environ or Environ\$(varName   varNum)	GetOption(name\$   id)	OpenFile-name\$([[title\$[, extensions\$]])]	TypeOf<objectVar> Is ObjectType

## Supported constants

[Table B-6, page 379](#), lists the Docbasic constants that are supported for Java evaluation.

**Table B-6. Docbasic constants supported for Java evaluation**

Constant	Description
ebBoolean	Integer representing Boolean datatype
ebCurrency	Integer representing Currency type
ebDataObject	Integer representing a data object
ebDate	Integer representing a Date type
ebEmpty	Integer representing an Empty variant
ebError	Integer representing an Error variant

Constant	Description
ebInteger	Integer representing the Integer type
ebLong	Integer representing the Long type
ebNull	Integer representing Null variant type
ebObject	Integer representing the OLE automation object type
ebSingle	Integer representing the Single type
ebString	Integer representing the String type
ebVariant	Integer representing the Variant type
Empty	Constant representing an initialized variant
False	Boolean constant
Nothing	Null Object reference
Null	Explicit Docbasic value distinguishable from Empty or Nothing. (Refer to the Docbasic documentation for complete information.)
Pi	Value of Pi
True	Boolean constant

## Unsupported constants

Table B-7, page 380, lists the constants that are not supported for Java evaluation.

**Table B-7. Constants not supported for Java evaluation**

Constant names		
ebCancel	ebNone	ebSystem
ebCritical	ebNormal	ebSystemModal
ebHidden	ebReadOnly	ebVolume

## Implicit objects

The implicit objects provided by Docbasic, such as the implicit object named Basic, are not supported for use in expressions that are compiled for Java evaluation.



## DQL Quick Reference

This appendix contains only the formal syntax descriptions of the DQL statements and the DQL reserved words. For a full description of each statement, refer to [Chapter 2, DQL Statements](#).

### The DQL statements

This section contains the formal syntax for each of the DQL statements.

#### Abort

```
ABORT [TRAN[SACTION]]
```

#### Alter Group

```
ALTER GROUP group_name ADD members
ALTER GROUP group_name DROP members
ALTER GROUP group_name SET ADDRESS email_address
```

#### Alter Type

```
ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
type_modifier_list [PUBLISH]

ALTER TYPE type_name
[FOR POLICY policy_id STATE state_name]
MODIFY (attribute_modifier_clause) [PUBLISH]
```

```
ALTER TYPE type_name
ADD attribute_def {,attribute_def} [PUBLISH]

ALTER TYPE type_name
DROP attribute_name {,attribute_name} [PUBLISH]
```

## Begin Tran

```
BEGIN TRAN[SACTION]
```

## Change...Object

```
CHANGE current_type OBJECT[S] TO new_type[update_list]
[IN ASSEMBLY document_id [VERSION version_label] [DESCEND]]
[SEARCH fulltext search condition]
[WHERE qualification]
```

## Commit

```
COMMIT [TRAN[SACTION]]
```

## Create Group

```
CREATE GROUP group_name [WITH] [ADDRESS mail_address]
[MEMBERS member_list]
```

## Create...Object

```
CREATE type_name OBJECT update_list
[,SETFILE filepath WITH CONTENT_FORMAT=format_name]
{,SETFILE filepath WITH PAGE_NO=page_number}
```

## Create Type

```
CREATE TYPE type_name [(attribute_def {,attribute_def})]
[WITH] SUPERTYPE parent_type
[type_modifier_list] [PUBLISH]
```



## Delete

```
DELETE FROM table_name WHERE qualification
```

## Delete...Object

```
DELETE type_name [(ALL)] [correlation_variable] OBJECT[S]  
[IN ASSEMBLY object_id [VERSION version_label]  
[NODE component_id] [DESCEND]]  
[SEARCH fulltext search condition]  
[WHERE qualification]
```

## Drop Group

```
DROP GROUP group_name
```

## Drop Type

```
DROP TYPE type_name
```

## Grant

```
GRANT privilege {,privilege} TO users
```

## Insert

```
INSERT INTO table_name [(column_name {,column_name})]  
VALUES (value {,value}) | dql_subselect
```

## Register

```
REGISTER TABLE [owner_name.]table_name (column_def {,column_def)  
[[WITH] KEY {column_list}]  
[SYNONYM [FOR] 'table_identification']
```

## Revoke

```
REVOKE privilege {,privilege} FROM users
```

## Select

The syntax for a standard query is:

```
SELECT [FOR base_permit_level] [ALL|DISTINCT] value [AS  
name] {,value [AS name]}  
FROM [PUBLIC] source_list  
[IN DOCUMENT clause | IN ASSEMBLY clause]  
[SEARCH [FIRST|LAST]fulltext_search_condition  
[IN FTINDEX index_name{,index_name}]  
[WHERE qualification]  
[GROUP BY value_list]  
[HAVING qualification]  
[UNION dql_subselect]  
[ORDER BY value_list]  
[ENABLE (hint_list)]
```

where the IN DOCUMENT *clause* is:

```
IN DOCUMENT object_id [VERSION version_label]  
[DESCEND] [USING ASSEMBLIES]  
[WITH binding_condition]  
[NODESORT BY attribute {,attribute} [ASC|DESC]]
```

and the IN ASSEMBLY *clause* is:

```
IN ASSEMBLY object_id [VERSION version_label]  
[NODE component_id] [DESCEND]
```

The syntax for an FTDQL SELECT statement is:

```
SELECT [FOR base_permit_level] [ALL|DISTINCT] value [AS  
name] {,value [AS name]}  
FROM [PUBLIC] source_list  
[SEARCH fulltext_search_condition  
[IN FTINDEX index_name{,index_name}]  
[WHERE qualification]  
[ORDER BY SCORE]  
[ENABLE (hint_list)]
```

Table C-1, page 387, lists the constraints imposed on the clauses in an FTDQL query.

**Table C-1. Summary of FTDQL query rules**

Applicable to	Rule
FTDQL queries in general	<p>An FTDQL query must contain one of the following:</p> <ul style="list-style-type: none"> <li>• SEARCH clause</li> <li>• One of: SCORE, SUMMARY, or TEXT as a selected value</li> <li>• ENABLE(FTDQL)</li> </ul> <p>The query may not contain:</p> <ul style="list-style-type: none"> <li>• An IN DOCUMENT or IN ASSEMBLY clause</li> <li>• The FIRST or LAST keyword in a SEARCH clause</li> <li>• A UNION, GROUP BY, or HAVING clause</li> </ul>
<i>selected values list</i>	<p>The selected values list must satisfy the following rules:</p> <ul style="list-style-type: none"> <li>• Only the following keywords are acceptable: CONTENTID, ISCURRENT, OBJTYPE, SCORE, SUMMARY, SYSOBJ_ID, TEXT, and THUMBNAIL_URL</li> <li>• Literals or expressions of any type may not be included.</li> <li>• An asterisk (*) is not allowed.</li> </ul>
AS clause	The AS clause is acceptable with no restrictions.
FROM clause	<p>The following rules apply to the FROM clause:</p> <ul style="list-style-type: none"> <li>• You may include only one object type reference, and that reference must be to dm_sysobject or a SysObject subtype.</li> <li>• You may not include references to a registered table.</li> </ul>
SEARCH clause	The SEARCH clause may not reference the FIRST or LAST keyword. There are no restrictions on the full-text search condition. However, note that the SEARCH DOCUMENT CONTAINS syntax is the preferred way to specify a SEARCH clause.

Applicable to	Rule
WHERE clause	<p>A WHERE clause in an FTDQL query is subject to the following rules:</p> <ul style="list-style-type: none"> <li>• Any repeating attributes referenced in the clause must be of type string or ID.</li> <li>• Only the DQL UPPER and LOWER functions are allowed. The functions SUBSTR, MFILE_URL, and all aggregate functions are not acceptable.</li> <li>• The following predicates are not allowed: <ul style="list-style-type: none"> <li>– BETWEEN</li> <li>– NOT LIKE</li> <li>– NOT FOLDER</li> <li>– [NOT] CABINET</li> <li>– ONLY</li> <li>– TYPE</li> </ul> </li> <li>• The IN and EXISTS keywords are not allowed.</li> <li>• Any valid form of the FOLDER predicate (except NOT FOLDER) may be used. The DESCEND option is also allowed.</li> <li>• The following rules apply to the LIKE predicate: <ul style="list-style-type: none"> <li>– The LIKE predicate may be used with pattern matching characters.</li> <li>– The LIKE predicate may not be used with an escape clause.</li> </ul> </li> <li>• The keywords TODAY, YESTERDAY, TOMORROW may not be used in the DATE function.</li> <li>• The following rules apply to all expressions in the WHERE clause: <ul style="list-style-type: none"> <li>– Expressions may not contain the ISREPLICA or USER keywords.</li> <li>– Expressions may use any comparison operator.</li> <li>– Expressions may use an arithmetic operator, but they may not be used to form a compound expression.</li> <li>– Expressions that compare one attribute to another are not allowed. For example, subject=title is invalid.</li> <li>– Expressions in the following format are not allowed:</li> </ul> </li> </ul>

Applicable to	Rule
	<pre>attribute_name operator('literal' operator 'literal')</pre> <ul style="list-style-type: none"> <li>– Expressions that force index correspondence between repeating attributes are not allowed. Such expressions AND together expressions that reference repeating attributes in the format: <pre>predicate(repeating_attr_expr AND repeating_attr_expr)</pre>           (For more information about forcing index correspondence, refer to <a href="#">Forcing index correspondence in query results, page 346.</a>)         </li> <li>• The following additional rules apply to expressions in the format <i>first_expression operator second_expression</i> <ul style="list-style-type: none"> <li>– The <i>first_expression</i> is limited to one of the following: <pre>attribute name upper(attribute name) lower(attribute name)</pre> </li> <li>– The <i>second_expression</i> is limited to one of the following: <pre>literal value upper(literal value) lower(literal value) the DATE() function</pre> </li> <li>– The operator may be any valid operator.</li> </ul> </li> </ul>
ORDER BY clause	<p>If included, this clause must be:</p> <pre>ORDER BY SCORE</pre> <p>SCORE must be referenced by name, not position, in the selected values list, also.</p>
ENABLE clause	<p>There are no restrictions on this clause.</p>

## Unregister

```
UNREGISTER [TABLE] [owner_name.]table_name
```

## Update

```
UPDATE table_name SET column_assignments
WHERE qualification
```

## Update...Object

```
UPDATE type_name [(ALL)] [correlation_variable] OBJECT[S] update_list
[,SETFILE filepath WITH CONTENT_FORMAT=format_name]
{,SETFILE filepath WITH PAGE_NO=page_number}
[IN ASSEMBLY document_id [VERSION version_label]
[NODE component_id] [DESCEND]]
[SEARCH fulltext search condition]
[WHERE qualification]
```

## DQL reserved words

If you using DQL reserved words as object or attribute names, enclose name in double quotes when using in a DQL query.

**Table C-2. DQL reserved words**

ABORT	BOOL	CONTENT_ID	DOCBASIC
ACL	BOOLEAN	COUNT	DOCUMENT
ADD	BROWSE	CREATE	DOUBLE
ADD_FTINDEX	BUSINESS	CURRENT	DROP
ADDRESS	BY	DATE	DROP_FTINDEX
ALL	CABINET	DATEADD	ELSE
ALLOW	CACHING	DATEDIFF	ELSEIF
ALTER	CHANGE	DATEFLOOR	ENABLE
AND	CHARACTER	DATETOSTRING	ENFORCE
ANY	CHARACTERS	DAY	ESCAPE
APPEND	CHAR	DEFAULT	ESTIMATE
APPLICATION	CHECK	DELETE	EXEC
AS	COMMENT	DELETED	EXECUTE
ASC	COMMIT	DEPENDENCY	EXISTS

ASSEMBLIES	COMPLETE	DEPTH	FALSE
ASSEMBLY	COMPONENTS	DESC	FIRST
ASSISTANCE	COMPOSITE	DESCEND	FLOAT
ATTR	COMPUTED	DISABLE	FOLDER
AVG	CONTAIN_ID	DISPLAY	FTINDEX
BEGIN	CONTAINS	DISTINCT	FUNCTION
BETWEEN	CONTENT_ FORMAT	DM_SESSION_DD _LOCALE	GRANT
GROUP	LINK	OF	REPORT
FOR	LOWER	ON	REVOKE
FOREIGN	MAPPING	ONLY	SCORE
FROM	MAX	OR	SEARCH
FT_OPTIMIZER	MCONTENTID	ORDER	SELECT
HAVING	MEMBERS	OWNER	SEPARATOR
HITS	MFILE_URL	PAGE_NO	SERVER
ID	MHITS	PARENT	SET
IF	MIN	PATH	SETFILE
IN	MODIFY	PERMIT	SMALLINT
INSERT	MONTH	POLICY	SOME
INT	MOVE	POSITION	STATE
INTEGER	MSCORE	PRIMARY	STORAGE
INTERNAL	NODE	PRIVATE	STRING
INTO	NODESORT	PRIVILEGES	SUBSTR
IS	NONE	PUBLIC	SUBSTRING
ISCURRENT	NOT	QRY	SUM
ISPUBLIC	NOTE	RDBMS	SUMMARY
ISREPLICA	NOW	READ	SUPERTYPE
KEY	NULL	REFERENCES	SUPERUSER
LANGUAGE	NULLDATE	REGISTER	SYNONYM
LAST	NULLINT	RELATE	SYSADMIN
LATEST	NULLSTRING	REMOVE	SYSOBJ_ID
LIST	OBJECT	REPEATING	SYSTEM

LIKE	OBJECTS	REPLACEIF	TABLE
TAG	TRUE	USER	WHERE
TEXT	TRUNCATE	USING	WITH
TIME	TYPE	VALUE	WITHIN
TO	UNION	VALUES	WITHOUT
TODAY	UNIQUE	VERSION	WORLD
TOMORROW	UNLINK	VERITY	WRITE
TOPIC	UNREGISTER	VIOLATION	YEAR
TRAN	UPDATE	WEEK	YESTERDAY
TRANSACTION	UPPER		



## Document Query Language examples

This appendix contains examples of Document Query Language (DQL) SELECT statements. The examples begin with basic SELECT statements that retrieve information about objects using the standard SELECT clauses, such as the WHERE clause and GROUP BY clause. Then the examples show SELECT statements that make use of the DQL extensions to the statement, such as the ability to use folder and virtual document containment information or registered tables to retrieve object information.

This appendix does not attempt to describe the syntax of the DQL SELECT statement in detail. For a complete description of SELECT, refer to [Select, page 115](#).

### Basic examples

This section presents basic SELECT statements. These examples demonstrate the use of standard clauses such as the WHERE, GROUP BY, and HAVING clauses, as well as how to use aggregate functions in a query. Some of these examples also demonstrate how to query repeating attributes.

### The simplest format

The basic SELECT statement has the format:

```
SELECT target_list FROM type_name
```

The *target\_list* identifies the value or values that you want to retrieve and *type\_name* identifies the types or registered tables from which you want to retrieve the requested information.

The following example returns the user names of all the users in the repository:

```
SELECT "user_name" FROM "dm_user"
```

This next example returns the names of the users and their user privileges within a repository:

```
SELECT "user_name", "user_privileges" FROM "dm_user"
```

## Using the WHERE clause

The WHERE clause restricts the information retrieved by placing a qualification on the query. Only the information that meets the criteria specified in the qualification is returned. The qualification can be simple or complex. It can contain any of the valid predicates or logical operators, such as AND, OR, or NOT.

The following example returns the user names of all users that belong to a specified group:

```
SELECT "user_name" FROM "dm_user"  
WHERE "user_group_name"='enr'
```

The next example retrieves all the types that do not have SysObject as their direct supertype and have more than 15 attributes:

```
SELECT "name" FROM "dm_type"  
WHERE "super_name" != 'dm_sysobject'  
AND "attr_count" > 15
```

## Searching repeating attributes in a WHERE Clause

To search a repeating attribute for a specific value, you use the ANY predicate in the WHERE clause qualification unless you have included the ROW\_BASED hint in the query. When you use this predicate, or if the query includes ROW\_BASED, if any of the values in the specified repeating attribute meets the search criteria, the server returns the object.

The following example searches the authors attribute to return any document that has Jim as one of its authors:

```
SELECT "object_name" FROM "dm_document"  
WHERE ANY "authors" = 'jim'
```

This next example searches the authors attribute and returns any document that has either Jim or Kendall as an author:

```
SELECT "object_name" FROM "dm_document"  
WHERE ANY "authors" IN ('jim', 'kendall')
```

The following example searches the keywords attribute and returns any document that has a key word that begins with hap:

```
SELECT "object_name" FROM "dm_document"
```

```
WHERE ANY "keywords" LIKE 'hap%'
```

## Using aggregate functions

DQL recognizes several aggregate functions. Aggregate functions are functions that operate on a set and return one value. For example, the count function is an aggregate function. It counts some number of objects and returns the total. (For a full description of all the aggregate functions that you can use in DQL queries, refer to [Aggregate functions, page 22.](#))

This example counts the number of documents owned by the user named john:

```
SELECT COUNT(*) FROM "dm_document"
WHERE "owner_name"='john'
```

The following example returns the dates of the oldest and newest documents in the repository:

```
SELECT MIN(DISTINCT "r_creation_date"),MAX(DISTINCT "r_creation_date")
FROM "dm_document"
```

This final example returns the average number of attributes in Documentum types:

```
SELECT AVG(DISTINCT "attr_count") FROM "dm_type"
```

## Using the GROUP BY clause

The GROUP BY clause provides a way to sort the returned objects on some attribute and return an aggregate value for each sorted subgroup. The basic format of the SELECT statement when you want to use the GROUP BY clause is:

```
SELECT attribute_name,aggregate_function FROM type_name
GROUP BY attribute_name
```

The attribute named in the target list must be the same as the attribute named in the GROUP BY clause.

The following example retrieves the names of all document owners and for each, provides a count of the number of documents that person owns:

```
SELECT "owner_name",count(*) FROM "dm_document"
GROUP BY "owner_name"
```

## Using the HAVING clause

The HAVING clause is used in conjunction with the GROUP BY clause. It restricts which groups are returned.

The following example returns the owner names and a count of their documents for those owners who have more than 10 documents:

```
SELECT "owner_name",count(*) FROM "dm_document"  
GROUP BY "owner_name"  
HAVING COUNT(*) > 10
```

## The ORDER BY clause

The ORDER BY clause lets you sort the values returned by the query. The clause has the format:

```
ORDER BY num [ASC|DESC] {,num [ASC|DESC] }
```

or

```
ORDER BY attribute [ASC|DESC] {,attribute [ASC|DESC]}
```

The *num* identifies one of the selected values by its position in the selected values list. The *attribute* identifies one of the selected values by its name. ASC (ascending) and DESC (descending) define the order of the sort and must be specified for each element specified in the ORDER BY clause.

The following example returns the owner's name, the object's title, and its subject for all SysObjects. The results are sorted by the owner's name and, within each owner's name group, by title:

```
SELECT "owner_name","title","subject" FROM "dm_sysobject"  
ORDER BY "owner_name", "title"
```

This next example returns the owner names and a count of their documents for those owners who have more than 10 documents. The results are ordered by the count, in descending order.

```
SELECT "owner_name",count(*) FROM "dm_document"  
GROUP BY "owner_name" HAVING COUNT(*) < 10  
ORDER BY 2 DESC
```

## Using the asterisk (\*) in queries

DQL lets you use the asterisk in the target list. If the query does not include the ROW\_BASED hint and the asterisk is the selected value, the FROM clause may specify only a type name or a registered table—it cannot specify both. If the query includes

the ROW\_BASED hint, the FROM clause may include either or both object types and registered tables. The values returned by the asterisk depend on what is specified in the FROM clause and whether the ROW\_BASED hint is included. For examples and complete details, refer to [The Asterisk \(\\*\) as a selected value, page 130](#).

## Searching cabinets and folders

In Documentum, objects of all SysObject subtypes except cabinets are stored in cabinets, and sometimes, in folders within those cabinets. It may, at times, be advantageous to restrict the search for an object or objects to a particular cabinet or folder. Or, you may want to determine what a particular cabinet or folder contains. To make these operations possible, DQL provides the CABINET and FOLDER predicates for use with the WHERE clause. The CABINET predicate lets you specify a particular cabinet to search. The FOLDER predicate lets you specify a particular folder or cabinet to search. (You can use the FOLDER predicate for cabinets as well as folders because cabinets are subtypes of folders.)

Use the cabinet or folder's folder path or its object ID to identify it in the query. A folder path has the format:

```
/cabinet_name{/folder_name}
```

To use the object ID, you must use the ID function and the object ID. The following examples illustrate the use of both folder paths and object IDs with the CABINET and FOLDER predicates.

This example returns all the folders contained in Research folder, which is identified by its object ID:

```
SELECT "object_name" FROM "dm_folder"
WHERE FOLDER (ID('0c00048400001599'))
```

The next example returns all the objects in the Marketing cabinet sorted by their type and, within each type, ordered alphabetically. The Marketing cabinet is identified by its folder path.

```
SELECT "object_name","r_object_type" FROM "dm_sysobject"
WHERE CABINET ('/Marketing')
ORDER BY "r_object_type","object_name"
```

This final example returns all the objects in the Correspondence folder in the Marketing cabinet. Again, the objects are sorted by their type and ordered alphabetically within each type group. A folder path is used to identify the correct folder.

```
SELECT "object_name","r_object_type" FROM "dm_sysobject"
WHERE FOLDER ('/Marketing/Correspondence')
ORDER BY "r_object_type","object_name"
```

The CABINET and FOLDER predicates have an optional keyword, DESCEND, that directs the server to return not only those objects directly contained within a cabinet or

folder but to also return the contents of any folders contained in that folder or cabinet, and so forth. (There is a limit of 25,000 contained folders within the specified folder when you include DESCEND in the predicate.)

The following example returns the name, object ID, and type of all objects in the Clippings folder, including the contents of any folders that are contained by the Clippings folder. The Clippings folder resides in the Marketing cabinet.

```
SELECT "object_name","r_object_id","r_object_type"  
FROM "dm_sysobject"  
WHERE FOLDER ('/Marketing/Clippings', DESCEND)
```

## Querying registered tables

A registered table is a table from your underlying RDBMS that has been registered with the repository. This registration allows you to specify the table in a DQL query, either by itself or in conjunction with a type. The following examples illustrate each possibility. (For a full discussion of the rules concerning the use of registered tables in SELECT statements, refer to [Chapter 2, DQL Statements](#).)

The first example returns all the name of all products and their owners from the registered table named ownership:

```
SELECT "product","manager" FROM "ownership"
```

This second example joins the registered table called ownership to the user-defined type called product\_info to return the names of all the objects in Collateral folder, along with the names of the managers who own the products described by the objects:


```
SELECT p."object_name", o."manager"  
FROM "product_info" p, "ownership" o  
WHERE p."product" = o."product" AND  
FOLDER ('/Marketing/Collateral', DESCEND)
```

## Querying virtual documents

DQL provides a clause that facilitates querying virtual documents. The clause is the IN DOCUMENT clause. The IN DOCUMENT clause lets you search a particular virtual document.

The examples in this section are based on the virtual document shown in [Figure D-1, page D-399](#):

Figure D-1. Virtual document model



## Determining the components

The following example returns only the directly contained components of Book1:

```
SELECT "r_object_id" FROM "dm_sysobject"
IN DOCUMENT ID('Book1_object_id')
```

This query returns Book1, Chapter 1, and Chapter 2. Book1 is included because a virtual document is always considered to be a component of itself.

To return all the components of a virtual document, including those that are contained in its components, you use the keyword `DESCEND`. For example:

```
SELECT "r_object_id" FROM "dm_sysobject"
IN DOCUMENT ID('Book1_object_id') DESCEND
```

The above example returns Book1, Chapter 1, Doc 1, Doc2, Chapter 2, and Doc 3, in that order.

**Note:** The components of a virtual document are returned in a pre-determined order that you cannot override with an `ORDER BY` clause.

You can limit the search to a particular object type by specifying the type in the `FROM` clause. For example, the following statement returns only the documents that make up Book1 (Doc 1, Doc 2, and Doc 3):

```
SELECT "r_object_id" FROM "dm_document"
IN DOCUMENT ID('Book1_object_id')
```





- \* (asterisk)
  - in queries, 396
  - wildcard in search queries, 141
- % (percent sign)
  - pattern matching character, 34
- \_ (underbar)
  - pattern matching character, 34
- [] (square brackets) for index values, 163
- (\*) asterisk
  - as selected value, 130
  - ROW\_BASED hint, effect of, 130

## A

- ABORT (statement), 42
- accent marks, in search strings, 141
- access controls for registered tables, 349
- ACLs
  - object types, assigning to, 74
  - setting default, 50
- administration methods
  - BATCH\_PROMOTE, 175
  - CAN\_FETCH, 177
  - CHECK\_CACHE\_CONFIG, 179
  - CHECK\_RETENTION\_  
EXPIRED, 182
  - CHECK\_SECURITY, 186
  - CLEAN\_LINKS, 189
  - DB\_STATS, 191
  - DELETE\_REPLICA, 193
  - DESTROY\_CONTENT, 195
  - DO\_METHOD, 197
  - DROP\_INDEX, 206
  - ESTIMATE\_SEARCH, 208
  - EXEC\_SQL, 211
  - executing
    - Apply method, 168
    - Documentum Administrator, 168
    - EXECUTE statement, 168
  - EXPORT\_TICKET\_KEY, 213
  - FINISH\_INDEX\_MOVES, 215

- GET\_INBOX, 219
  - GET\_LAST\_SQL, 223
  - GET\_PATH, 225
  - GET\_SESSION\_DD\_LOCALE, 227
  - HTTP\_POST, 228
  - IMPORT\_REPLICA, 234
  - IMPORT\_TICKET\_KEY, 236
  - LIST\_AUTH\_PLUGINS, 238
  - LIST\_RESOURCES, 240
  - LIST\_SESSIONS, 244
  - LIST\_TARGETS, 248
  - LOG\_OFF, 250
  - LOG\_ON, 251
  - MAKE\_INDEX, 253
  - MARK\_AS\_ARCHIVED, 256
  - MARK\_FOR\_RETRY, 258
  - MIGRATE\_CONTENT, 260
  - MODIFY\_TRACE, 269
  - MOVE\_INDEX, 271
  - PING, 273
  - PURGE\_AUDIT, 275
  - PURGET\_CONTENT, 283
  - PUSH\_CONTENT\_ATTRS, 285
  - RECOVER\_AUTO\_TASKS, 288
  - REGISTER\_ASSET, 290
  - REORGANIZE\_TABLE, 292
  - REPLICATE, 295
  - RESET\_TICKET\_KEY, 298
  - RESTORE\_CONTENT, 300
  - ROLES\_FOR\_USER, 302
    - scope, 168
  - SET\_APIDEADLOCK, 304
  - SET\_CONTENT\_ATTRS, 307
  - SET\_OPTIONS, 312
  - SET\_STORAGE\_STATE, 316
  - SHOW\_SESSIONS, 319
  - SYNC\_REPLICA\_RECORDS, 321
  - TRANSCODE\_CONTENT, 325
  - UPDATE\_STATISTICS, 329
  - WEBCACHE\_PUBLISH, 332
- aggregate functions

- as selected values, 124
    - AVG, 24
    - COUNT, 22
    - examples of use, 395
    - MAX, 23
    - MIN, 23
    - SUM, 24
  - aggregate functions, default values
    - and, 342
  - ALL (keyword), 91, 121
  - ALL (predicate), 31
  - ALTER GROUP (statement), 43
  - ALTER TYPE (statement), 45
  - AND (logical operator), 38
  - ANSI format for dates, 17
  - ANY keyword in WHERE clauses, 143
  - ANY...comparison operator
    - (predicate), 32
  - ANY...IN subquery (predicate), 32
  - ANY...IS NULL (predicate), 32
  - ANY...IS NULLDATE (predicate), 32
  - ANY...IS NULLINT (predicate), 32
  - ANY...IS NULLSTRING (predicate), 32
  - ANY...LIKE (predicate), 32
  - application server
    - sending request to, 228
  - Apply method
    - operations, 169
  - archived documents, restoring, 300
  - archiving
    - audit trail entries, 256
    - archiving content files, 283
  - arithmetic expressions as selected values, 125
  - arithmetic operators, 29
  - ascending sort order, 148
  - assemblies
    - IN ASSEMBLY clause, 137
    - SELECT statements and, 136
  - asterisk (\*)
    - as selected value, 130
    - in queries, 396
    - ROW\_BASED hint, effect of, 130
    - wildcard in search queries, 141
  - attribute drop clause, 57
  - attribute values, retrieving, 121
  - attributes, 343
    - See also* repeating attributes
    - adding to type definition, 58
    - assume\_user\_location and disabling
      - assume user program, 204
    - data dictionary information,
      - defining, 74
    - default values, defining, 77
    - dropping, 58
    - lengthening character string types, 58
    - mapping table information,
      - defining, 77
    - modifiers, 54
    - modifying, 58
    - names, quoting in applications, 41
    - NULL values and, 340
    - querying indexed, 139
    - referencing, 340
  - audit trail entries
    - i\_is\_archived, setting, 256
    - removing from repository, 275
  - auditing
    - audit trail entries, removing, 275
  - authentication
    - LIST\_AUTH\_PLUGINS
      - administration method, 238
    - LIST\_SESSIONS description, 245
    - tracing, 314
  - authentication plugins
    - list of, obtaining, 238
  - automatic activities
    - RECOVER\_AUTO\_TASKS
      - administration method, 288
  - AVG (aggregate function), 24
- ## B
- BATCH\_PROMOTE administration
    - method, 175
  - BEGIN TRAN (statement), 60
  - brackets for index values, 163
- ## C
- CABINET (predicate), 37
  - cabinets
    - CABINET (predicate), 37
    - querying (examples of), 397
  - cache config objects
    - consistency checking, 179
  - cached queries, 350
  - CAN\_FETCH
    - administration method, 177

- case sensitivity
  - scientific notation, 14
- case sensitivity, in search strings, 140
- CHANGE OBJECT (statement), 61
- CHANGE...OBJECT (statement), 339
- character strings
  - converting dates to, 27
  - dates within, 15
  - lengths, 15
  - literal formats, 15
  - lowercase and LOWER function, 21
  - pattern arguments, 17
  - returning substrings, 21
  - short date formats, 16
  - uppercase and UPPER function, 21
- check constraints, 80, 83
- CHECK\_CACHE\_CONFIG
  - (administration method), 179
- CHECK\_RETENTION\_EXPIRED
  - (administration method), 182
- CHECK\_SECURITY (administration method), 186
- CLEAN\_LINKS
  - administration method, 189
- client applications
  - checking for server connection, 273
  - short date format defaults, 16
- client code pages
  - tracing, 313
- client machines
  - short date format defaults, 18
  - tracing information for, 314
- client roles
  - determining, 302
- code page requirements
  - SYNONYM FOR clause, 112
- columns
  - defining, 109
  - inserting values into, 105
  - updating, 155
- comments, in DQL, 41
- COMMIT (statement), 64
- common area
  - removing files from, 313
- component specifications
  - dropping, 53
  - removing, 53
  - specifying, 52
- component\_specification, 84
- connection brokers
  - obtaining list of targets, 248
  - tracing information, 313
- connections
  - checking with PING, 273
- constraints
  - check (data validation), 80, 83
  - creating with ALTER TYPE, 56
  - creating with CREATE TYPE, 78
  - dropping with attribute\_drop\_clause, 57
  - dropping with type drop clause, 53
  - foreign key, 79
  - foreign keys, 83
  - NOT NULL, 80
  - primary key, 78, 82
  - removing, 53
  - specifying, 52
  - unique key, 79, 82
- CONTAIN\_ID (keyword), 127
- content addressable storage
  - PUSH\_CONTENT\_ATTRS
    - administration method, 285
- content files, 295
  - archiving, 283
  - deleting replicas, 193
  - fetching, 177
  - importing replicas, 234
  - MIGRATE\_CONTENT administration method, 260
  - obtaining file path, 225
  - removing, 195
  - removing from content-addressed storage, 196
  - restoring archived, 300
  - transformations, requesting, 325
- content objects
  - destroying, 195
  - SET\_CONTENT\_ATTRS
    - administration method, 307
- Content Server
  - common area, removing files from, 313
  - Docbasic expressions, handling of, 367
  - history\_cutoff startup parameter, 247
  - history\_sessions startup parameter, 247
  - listing active sessions, 319
  - listing connection broker targets, 248
  - obtaining resource information, 240

- obtaining session information, 244
  - content-addressed storage
    - objects with expired retention, finding, 182
  - content-addressed storage systems
    - content, removing, 196
    - tracing operations in, 313
  - CONTENTID (keyword), 126
  - contents
    - adding with SETFILE, 161
  - correlation names, 134
  - COUNT (aggregate function), 22
  - CREATE...GROUP (statement), 65
  - CREATE...OBJECT (statement), 67
  - CREATE...TYPE (statement), 71
  - cursors, RDBMS, 192
- D**
- data dictionary
    - attribute information, defining, 74
    - component specification, 84
    - current locale, obtaining, 227
    - default values for attributes, 340
    - defining
      - check constraints, 80, 83
      - default attribute values, 77
      - foreign key constraints, 79, 83
      - mapping table information, 77
      - NOT NULL constraints, 80
      - primary key constraints, 78, 82
      - unique key constraints, 79, 82
      - value assistance information, 75
    - defining type information, 81
    - Java evaluation of Docbasic expressions, 367
    - PUBLISH (DQL keyword), 48, 74
    - update\_modifier specification, 74, 81
  - data validation, check constraints, 80
  - database hints
    - FETCH\_ALL\_RESULTS, 360
    - FORCE\_ORDER, 357
    - FTDQL and NOFTDQL, 364
    - IN and EXISTS, 362
    - including multiple hints, 365
    - OPTIMIZATION\_LEVEL, 361
    - OPTIMIZE\_TOP, 360
    - passthrough hints, 365
    - RETURN\_TOP, 358
    - ROW\_BASED, 363
    - SQL\_DEF\_RESULT\_SET, 356
    - UNCOMMITTED\_READ, 361
    - using, 355
  - datatypes
    - lengthening character string types, 58
  - date functions
    - DATEADD, 26
    - DATEDIFF, 25
    - DATEFLOOR, 26
    - DATETOSTRING, 27
  - date literals, 15
  - DATEADD (date function), 26
  - DATEDIFF (date function), 25
  - DATEFLOOR (date function), 26
  - dates
    - ANSI format, 17
    - date literals, 15
    - earliest acceptable year, 15
    - functions, 24
    - input formats, 15
    - localized formats, 16
    - NOW (keyword), 18
    - output formats, 18
    - short date format defaults, 18
    - TODAY (keyword), 18
    - TOMORROW (keyword), 18
    - truncating, 26
    - using in queries, 15
    - YESTERDAY (keyword), 18
  - DATETOSTRING (date function), 27
  - DB\_STATS administration method, 191
  - DB2
    - RETURN\_TOP, 358
  - DB2 databases
    - OPTIMIZATION\_LEVEL (DQL hint), 361
  - dbo alias, 349
  - deadlocks, simulating, 304
  - debug tracing, 313
  - default values
    - aggregate functions and, 342
    - default attribute specification, 55
    - groups for a type, 51
    - private and public groups, 66
    - returned format, 340
    - storage area for types, 51
    - testing for, 341
  - DELETE (statement), 87, 339
  - DELETE...OBJECT (statement), 89, 339

- DELETE\_REPLICA administration
  - method, 193
- DEPTH (keyword), 127
- DESCEND (keyword), 135
- descending sort order, 148
- DESTROY\_CONTENT administration
  - method, 195
- diacritical marks, in search strings, 141
- digital shredding
  - tracing, 313
- disabling
  - assume\_user program, 204
- DISTINCT (keyword), 121
- distributed storage
  - deleting replicas, 193
  - fetching content files, 177
  - importing replicas, 234
  - replicating files, 295
- distributed storage areas
  - updating replica records, 321
- dm\_bp\_batch method, 175
- dm\_CheckCacheConfig job, 181
- dm\_dbo alias, 349
- dm\_registered object type, 349
- DM\_SESSION\_DD\_LOCALE
  - (keyword), 19
- dmc\_MigrateDbExprsToJava method, 369
- dmc\_MigrateDbExprsToJavaForType
  - method, 369
- dmc\_SetJavaExprEnabled method, 371
- DMCL (client library)
  - version number returned by
    - LIST\_SESSIONS, 245
- dmi\_audittrail\_attrs
  - SELECT constraint, 133
- DO\_METHOD administration
  - method, 197
- DO\_METHOD procedures
  - running as server, 203
- DO\_METHOD result object attributes, 199
- Docbasic, 367
  - See also* Docbasic expression migration
  - check constraints, 80, 83
  - constants
    - supported for Java
      - evaluation, 379
    - unsupported for Java
      - evaluation, 380
  - executing scripts, 197
  - expression handling, 367
  - expressions, migrating to Java, 367
  - functions
    - supported for Java
      - evaluation, 373
    - unsupported for Java
      - evaluation, 378
  - operators, supported for Java
    - evaluation, 372
  - value\_assistance\_modifier, 76
- Docbasic expression migration
  - dmc\_MigrateDbExprsToJava
    - method, 369
  - dmc\_MigrateDbExprsToJavaForType
    - method, 369
  - implicit objects, 381
  - results document, 371
  - supported expression
    - components, 372
- Docbasic expressions
  - Java evaluation, disabling, 371
- Docbasic expressions migration
  - return values, 371
- double quotes, using, 340
- DQL (Document Query Language)
  - applications, use in, 351
  - arithmetic operators, 29
  - basic query statements, 338
  - comments, in-line, 41
  - comparison operators, 30
  - date functions, 24
  - examples, 393
  - functions, 20
  - GET\_LAST\_SQL administration
    - method, 223
  - ID function, 27
  - keywords, 19
  - literals, 13
  - logical operators, 38
  - MFILE\_URL function, 28
  - predicates, 29
  - queries, caching, 350
  - query results, processing, 351
  - registered tables, referencing, 349
  - repeating attributes, referencing, 343
  - reserved words, 390
  - tracing generated SQL, 313
  - user privileges and, 354
- DQL hints
  - ROW\_BASED, 123
  - using, 149

DROP GROUP (statement), 93  
 DROP TYPE (statement), 94  
 DROP\_INDEX administration  
   method, 206  
 dump files  
   date format in, 19

## E

error messages  
   date formats in, 19  
 escape character, 35  
 ESTIMATE\_SEARCH (administration  
   method), 208  
 EXEC\_SQL administration method, 211  
 EXECUTE statement  
   described, 96  
 EXISTS (database hint), 362  
 EXISTS (predicate), 31  
 EXPORT\_TICKET\_KEY administration  
   method, 213  
 expressions  
   Java evaluation, disabling, 371  
 external applications  
   deadlocks, simulating, 304  
   executing, 197  
   query execution/processing, 351  
 external procedures  
   method objects and, 201

## F

FALSE (keyword), 19  
 FETCH\_ALL\_RESULTS, 360  
 files  
   removing from server common  
     area, 313  
 FINISH\_INDEX\_MOVES administration  
   method, 215  
 FIRST keyword, 139  
 floating point literals, 14  
 FOLDER (predicate), 36  
   nested folder limit, 37  
 folders  
   FOLDER (predicate), 36  
   querying (examples of), 397  
 FORCE\_ORDER, 357  
 foreign key constraints, 79, 83  
 formats  
   date literals, 15

  date output, 18  
   floating point literals, 14  
   folder path, 36  
   integer literals, 13  
   setting with DATETOSTRING  
     function, 27  
   value assistance, 75  
 FROM clause  
   disambiguating types and tables, 134  
   table name specification, 133  
   type name specification, 132  
 FTDQL DQL hint, 364  
 FTDQL queries, *see* SELECT (statement)  
   distinct\_query\_results server.ini key  
     and, 121  
 full-text indexes  
   MARK\_FOR\_RETRY administration  
     method, 258  
   querying, 208  
   SEARCH clause, 138  
   virtual documents and, 348  
 full-text indexing  
   tracing operations, 269  
 full-text keywords, 126  
 functions  
   aggregate, 342  
   Apply method, 169  
   date, 24  
   DATEADD, 26  
   DATEDIFF, 25  
   DATEFLOOR, 26  
   DATETOSTRING, 27  
   defined, 20  
   EXECUTE statement, 97  
   GROUP BY clause and, 146  
   ID, 27  
   LOWER, 21  
   MFILE\_URL, 28, 125  
   SUBSTR, 21  
   UPPER, 21

## G

GET\_INBOX administration method, 219  
 GET\_LAST\_SQL administration  
   method, 223  
 GET\_PATH administration method, 225  
 GET\_SESSION\_DD\_LOCALE  
   administration method, 227  
 GRANT (statement), 103

- GROUP BY clause
  - described, 146
  - examples of use, 395
- groups
  - alter group (statement), 43
  - CHECK\_SECURITY (administration method), 186
  - creating, 65
  - default for type, 51
  - private, 44
  - public, 44
  - removing, 93
- H**
- HAVING clause
  - defined, 146
  - examples of use, 396
- hierarchy, searching, 135
- history\_cutoff (startup parameter), 247
- history\_sessions (startup parameter), 247
- HTTP\_POST administration method, 228
- I**
- i\_is\_archived attribute, 256
- ID function, 27
- ID literal format, 15
- IMPORT\_REPLICA administration method, 234
- IMPORT\_TICKET\_KEY administration method, 236
- IN (database hint), 362
- IN (predicate), 31
- IN and EXISTS, 362
- IN ASSEMBLY clause
  - DELEET...OBJECT and, 91
  - SELECT and, 137
  - unioned selects and, 147
- IN DOCUMENT clause
  - described, 134
  - NODESORT BY option, 137
  - unioned selects and, 147
  - WITH option, 136
- IN FTINDEX clause, 142
- inboxes
  - GET\_INBOX administration method, 219
- index positions
  - defined, 343
  - forcing correspondence, 346
- indexed attributes
  - querying, 139
- indexes, *see* object type indexes
- input formats for dates, 15
- INSERT (statement), 105, 339
- integer literal format, 13
- internationalization
  - client code pages, tracing, 313
  - object types
    - altering and, 48
    - creating and, 74
  - session locales, tracing, 313
- IS NULL (predicate), 31
- IS NULLDATE (predicate), 31
- IS NULLINT (predicate), 31
- IS NULLSTRING (predicate), 31
- ISCURRENT (keyword), 126
- ISPUBLIC (keyword), 126
- ISREPLICA (keyword), 128
- J**
- Java
  - evaluation of Docbasic expressions, disabling, 371
- java servlets, invoking, 228
- jobs
  - dm\_CheckCacheConfig, 181
- K**
- keywords
  - as selected values, 125
  - DQL, 19
  - full-text, 126
  - miscellaneous, 128
  - PUBLISH, 48, 74
  - THUMBNAIL\_URL, 129
  - virtual document, 127
- L**
- LAST keyword, 139
- LAUNCH\_DIRECT, run\_as\_server and, 204
- lifecycles
  - batch promotion, 175
  - defining default for object types, 51
- LIKE (predicate)
  - described, 31

- pattern matching with, 34
  - link record object type, 189
  - linked storage areas
    - removing unneeded objects from, 189
  - linking objects with UPDATE...
    - OBJECT, 164
  - links
    - removing with UPDATE...
      - OBJECT, 164
  - LIST\_AUTH\_PLUGINS administration
    - method, 238
  - LIST\_RESOURCES administration
    - method, 240
  - LIST\_SESSIONS administration
    - method, 244
  - LIST\_TARGETS administration
    - method, 248
  - literals
    - character string, 15
    - date, 15
    - date keywords, 18
    - escape character and pattern
      - matching, 35
    - floating point, 14
    - ID format, 15
    - integer, 13
    - pattern matching, 34
    - types of, 13
  - locales
    - determining current, 227
    - DM\_SESSION\_DD\_LOCALE
      - keyword, 19
  - locking
    - tracing Windows, 313
  - log files
    - LOG\_OFF administration
      - method, 250
    - LOG\_ON administration method, 251
  - LOG\_OFF administration method, 250 to 251
  - LOG\_ON administration method, 251
  - logical operators
    - described, 38
    - precedence order, 39
  - login ticket key
    - EXPORT\_TICKET\_KEY
      - administration method, 213
    - IMPORT\_TICKET\_KEY
      - administration method, 236
    - RESET\_TICKET\_KEY administration
      - method, 298
  - login ticket keys
    - tracing import/export, 314
  - login tickets
    - tracing, 314
  - logs
    - tracing and, 270
  - LOWER (scalar function), 21
- ## M
- MAKE\_INDEX administration
    - method, 253
  - mapping information, dropping, 57
  - mapping table specification, 55, 77
  - MARK\_AS\_ARCHIVED administration
    - method, 256
  - MARK\_FOR\_RETRY administration
    - method, 258
  - MAX (aggregate function), 23
  - maximums
    - characters in string literals, 15
    - floating point literal values, 14
    - number of attributes in ALTER TYPE
      - statements, 58
    - values returned by MAX aggregate
      - function, 23
  - method objects
    - run\_as\_server attribute, 203
    - running procedures as server, 203
  - method server
    - tracing execution, 314
  - methods
    - dm\_bp\_batch, 175
    - dmc\_SetJavaExprEnabled, 371
    - executing with DO\_METHOD, 168
    - tracing method server, 314
  - MFILE\_URL (DQL function), 28, 125
  - MIGRATE\_CONTENT administration
    - method, 260
  - MIN (aggregate function), 23
  - minimum values
    - floating point literals, 14
    - returned by MIN aggregate
      - functions, 23
  - MODIFY\_TRACE administration
    - method, 269
  - MOVE\_INDEX administration
    - method, 271



**N**

names  
 user, 129  
 NODESORT BY, 137  
 NOFTDQL DQL hint, 364  
 NOT (logical operator), 38  
 NOT EXISTS (predicate), 31  
 NOT IN (predicate), 31  
 NOT LIKE (predicate)  
 described, 31  
 pattern matching with, 34  
 NOT NULL constraints, 80  
 NOW (keyword), 18  
 NULL values  
 in attributes, 340  
 sorting, 342  
 testing for, 341  
 NULLDATE (keyword), 19

**O**

object IDs  
 IN DOCUMENT clause and, 135  
 object type indexes  
 creating, 253  
 DROP\_INDEX administration  
 method, 206  
 FINISH\_INDEX\_MOVES  
 administration method, 215  
 moving, 271  
 object types  
 attributes  
 maximum allowed number of, 73  
 attributes, adding, 58  
 changing to another, 339  
 changing with ALTER TYPE, 45  
 component specifications,  
 removing, 53  
 component\_specification, 84  
 constraints, removing, 53  
 creating, 71  
 data dictionary information,  
 defining, 81  
 defining default lifecycle for, 51  
 deleting, 94  
 mapping table specification, 51, 81, 84  
 names, quoting in applications, 41  
 referencing, 340  
 referencing in queries, 132  
 removing indexes for, 206

object-level permissions  
 DQL statements and, 354  
 objects  
 batch promotion, 175  
 changing to another type, 339  
 changing type, 61  
 creating, 67  
 deleting, 89  
 linking with UPDATE...OBJECT, 164  
 moving to new folder/cabinet, 164  
 unlinking with UPDATE...  
 OBJECT, 164  
 Update...Object (statement), 158  
 OBJTYPE (keyword), 126  
 operators  
 arithmetic, 29  
 comparison, 30  
 logical, 38  
 OPTIMIZE\_TOP, 360  
 OR (logical operator), 38  
 Oracle  
 RETURN\_TOP, 359  
 SYNONYM FOR code page  
 requirement, 112  
 ORDER BY clause  
 examples of use, 396  
 SELECT statement, 148

**P**

PARENT (keyword), 128  
 parentheses to change precedence of  
 logical operators, 39  
 passthrough database hints, 365  
 pattern matching  
 characters, 34  
 escape character, 35  
 percent sign (%)  
 pattern matching character, 34  
 performance  
 SEARCH clauses, 209  
 searches and PUBLIC keyword, 132  
 WHERE clause, 143  
 permits, registered tables, 349  
 persistent client caching  
 CHECK\_CACHE\_CONFIG  
 (administration method), 179  
 dm\_CheckCacheConfig job, 181  
 PING administration method, 273  
 precedence

- date formats, 17
- logical operators, 39
- predicates
  - ANY...comparison operator, 32
  - ANY...IN, 32
  - ANY...IS NULL, 32
  - ANY...IS NULLDATE, 32
  - ANY...IS NULLINT, 32
  - ANY...IS NULLSTRING, 32
  - ANY...LIKE, 32
- arithmetic operators, 29
- CABINET, 37
- comparison operators, 30
- defined, 29
- FOLDER, 36
- pattern matching, 34
- repeating attribute predicates, 32
- single-valued attribute predicates, 31
- SysObject, 35
- TYPE, 36
- where clause components, 142
- WHERE clause components, 29
- primary key constraints, 78, 82
- printing, assume user feature and, 204
- private
  - groups, 44, 66
- privileges
  - granting, 103
  - revoking, 113
- programs
  - assume user, 204
- public
  - groups, 44, 66
  - objects, selecting, 126
- PUBLIC keyword
  - creating public groups, 65
  - DELETE...OBJECT statement, 91
  - SELECT statements and, 131
- PUBLISH keyword, 48, 74
- publishing, to Web site, 332
- PURGE\_AUDIT administration
  - method, 275
- PURGE\_CONTENT administration
  - method, 283
- PUSH\_CONTENT\_ATTRS administration
  - method, 285

## Q

- qualifying table names, 134

- queries
  - database hints, 355, 366
  - executing SQL, 211
  - FROM clause limitations, 132
- query optimization
  - OPTIMIZATION\_LEVEL (DQL hint), 361
- query result objects
  - ROW\_BASED DQL hint effects, 363
- query results
  - returning as object-based, 122
  - ROW\_BASED DQL hint, 123
- queue items
  - marking for indexing retry, 258

## R

- RDBMS
  - deleting table rows, 87
  - inserting table rows, 105
  - operation statistics, 191
  - passthrough hints for SELECT, 149
  - permissions for Superusers, 354
  - registered tables
    - described, 349
    - permissions for, 350
    - referencing in queries, 349
    - security, 349
  - REORGANIZE\_TABLE
    - administration method, 292
  - unregistering tables, 153
  - UPDATE\_STATISTICS administration
    - method, 329
  - updating rows, 155
- RECOVER\_AUTO\_TASKS administration
  - method, 288
- REGISTER (statement), 108
- REGISTER\_ASSET administration
  - method, 290
- registered tables
  - creating, 108
  - deleting rows, 87
  - described, 349
  - DQL predicates for, 31
  - inserting rows, 105
  - permissions, default, 350
  - querying, 124
  - querying (examples of), 398
  - referencing in queries, 132, 349
  - security, 349

- SYNONYM clause, 111
  - table names, qualifying, 134
  - unregistering, 153
  - updating rows, 155
  - removing
    - repeating attribute values, 164
  - REORGANIZE\_TABLE administration
    - method, 292
  - repeating attributes
    - adding values, 343, 345
    - comparison expressions, use in, 144
    - compound ANY predicates, use in, 143
    - deleting values, 345
    - forcing index correspondence, 346
    - index positions, 343
    - NULL values and, 340
    - predicates for, 32
    - querying (examples of), 394
    - removing multiple values, 164
    - truncating, 164
    - updating values, 345
    - values
      - appending, 163
      - inserting, 163
      - removing, 164
    - WHERE clause use, 143
  - replicas
    - deleting, 193
    - determination of, 128
    - importing, 234
  - REPLICATE administration method, 295
  - replicating, 295
  - repositories
    - audit trail entries, removing, 275
    - IMPORT\_TICKET\_KEY
      - administration method, 236
      - owner name aliases, 349
  - repository sessions
    - query caching, 350
  - reserved words
    - table of, 390
  - RESET\_TICKET\_KEY administration
    - method, 298
  - RESTORE\_CONTENT administration
    - method, 300
  - retention policies
    - tracing use, 314
  - retention\_trace tracing option, 314
  - RETURN\_TOP
    - described, 358
    - SEARCH clause and, 359
  - REVOKE (statement), 113
  - roles, *see* client roles
  - ROLES\_FOR\_USER administration
    - method, 302
  - ROW\_BASED DQL hint, 363
    - affect on asterisk as selected value, 130
  - RPC logging
    - turning off, 250
    - turning on, 251
  - RPC tracing, 314
  - run\_as\_server attribute, 203
- ## S
- scalar functions
    - as selected values, 124
    - LOWER, 21
    - SUBSTR, 21
    - UPPER, 21
  - SCORE (keyword), 126
  - SEARCH clause
    - accents and diacritical marks in, 141
    - asterisk as wildcard, 141
    - case sensitivity, 140
    - DOCUMENT CONTAINS
      - format, 139
    - IN FTINDEX option, 142
    - RETURN\_TOP and, 359
    - SELECT statement, 138
    - TOPIC format, 142
    - words and phrases, specifying, 139
  - search string
    - words and phrases, specifying, 139
  - search strings
    - accent marks in, 141
    - asterisk as wildcard, 141
    - case sensitivity, 140
    - diacritical marks in, 141
  - security
    - CHECK\_SECURITY (administration method), 186
    - DQL statements and, 354
  - SELECT (statement)
    - ALL keyword, 121, 132
    - arithmetic expressions as selected values, 125
    - AS clause, 130
    - assigning result attribute names, 130

- asterisk as selected value, 130
  - DELETED keyword, 132
  - described, 115, 338
  - disambiguating types and tables, 134
  - DISTINCT (keyword), 121
  - DQL functions as selected values, 124
  - DQL hints, 149
  - FROM clause, 131
  - FROM clause limitations, 132
  - FTDQL queries, 120
  - GROUP BY clause, 146
  - HAVING clause, 146
  - IN ASSEMBLY clause, 137
  - IN DOCUMENT clause, 134
  - keywords as selected values, 125
  - MFILE\_URL as selected value, 125
  - ORDER BY clause, 148
  - permissions and, 354
  - retrieving attribute values, 121
  - retrieving column values, 124
  - SEARCH clause, 138, 208
  - table name specification, 133
  - type name specification, 132
  - UNION clause, 147
  - WHERE clause, 142
  - where clause qualification
    - components, 142
  - WITH option, 136
  - SELECT statement
    - dmi\_audittrail\_attrs constraint, 133
  - server.ini file, session information
    - parameters, 247
  - servers. *See* Content Server, 350
  - session locales
    - tracing, 313
  - sessions
    - listing active, 319
    - listing current and historical, 244
  - SET\_APIDEADLOCK administration
    - method, 304
  - SET\_CONTENT\_ATTRS administration
    - method, 307
  - SET\_OPTIONS administration
    - method, 312
  - SET\_STORAGE\_STATE administration
    - method, 316
  - SETFILE clause, 161
  - short date formats
    - defaults, 18
    - described, 16
  - SHOW\_SESSIONS administration
    - method, 319
  - SOME predicate, 31
  - sorting (ORDER BY clause), 148
  - SQL (Structured Query Language)
    - executing queries, 211
    - tracing last statement, 313
    - tracing statements, 314
  - SQL Server
    - FETCH\_ALL\_RESULTS hint
      - constraint, 361
    - RETURN\_TOP, 358
    - RETURN\_TOP hint constraint, 358
  - SQL\_DEF\_RESULT\_SET, 356
  - square brackets for index values, 163
  - statistics
    - RDBMS, determining, 191
  - storage areas
    - archiving files, 283
    - making read-only, 316
    - moving offline/online, 316
    - restoring archived content, 300
    - setting default, 51
  - SUBSTR (function), 21
  - SUM (aggregate function), 24
  - SUMMARY (keyword), 126
  - Superuser user privilege
    - RDBMS permissions, 354
  - Sybase
    - RETURN\_TOP, 359
  - SYNC\_REPLICA\_RECORDS
    - administration method, 321
  - SYNONYM clause, 111
  - SYNONYM FOR clause
    - internationalization requirements, 112
  - syntax
    - dates, 15
  - SYSOBJ\_ID (keyword), 126
  - SysObjects
    - predicates, 35
  - system administration
    - execute (statement), 96
- ## T
- table permits, 349
  - tables
    - deleting rows, 87
    - inserting rows, 105
    - querying, 124

- registering, 108
- unregistering, 153
- updating rows, 155
- tablesaces
  - moving type indexes, 271
- TEXT (keyword), 126
- THUMBNAIL\_URL (keyword), 129
- thumbnails
  - generating, 290
- TODAY (keyword), 18
- TOMORROW (keyword), 18
- tracing
  - content-addressed storage
    - operations, 313
  - full-text index operations, 269
  - login ticket keys, 314
  - login tickets, single-use, 314
  - retention policy use, 314
  - RPC calls, 251
  - trace file storage, 312
  - turning on/off, 312
  - workflow agent operations, 314
- transactions
  - aborting explicit, 42
  - committing, 64
  - starting explicit, 60
- TRANSCODE\_CONTENT administration
  - method, 325
- transformations, requesting, 325
- TRUE (keyword), 19
- truncating repeating attributes, 164
- TYPE (predicate), 36
- types
  - alter type statement, 45
  - create...object (statement), 67
  - creating, 71
  - deleting, 94
  - dropping attributes, 58
  - modifying attributes, 58
  - setting default ACL, 50

## U

- UNCOMMITTED\_READ (database hint), 361
- underbar ( \_ )
  - pattern matching character, 34
- UNION clause, 147
- unique key constraints, 79, 82
- UNIX process launch, tracing, 314

- unlinking
  - objects with UPDATE...OBJECT, 164
- Unregister (statement), 153
- UPDATE (statement), 155, 338
- update modifier, 54
- UPDATE...OBJECT (statement), 158, 339
- update\_modifier specification, 74, 81
- UPDATE\_STATISTICS administration
  - method, 329
- UPPER (scalar function), 21
- URLs
  - MFILE\_URL (DQL function), 28
  - THUMBNAIL\_URL (keyword), 129
- use\_estimate\_search (server.ini key), 209
- USER (keyword), 19, 129
- user privileges
  - granting, 103
  - revoking, 113
- user privileges, DQL statements and, 354
- users
  - CHECK\_SECURITY (administration method), 186
  - name of current, 129
- USING ASSEMBLIES, 136

## V

- va\_clause (data dictionary), 76
- value assistance
  - dropping, 57
  - format, 75
  - specification, 54
- VERSION (keyword), 135
- virtual documents
  - full-text indexes and, 348
  - IN DOCUMENT clause, 134
  - keywords for, 127
  - querying, 347
  - querying (examples of), 398

## W

- Web sites
  - publishing to, 332
- WEBCACHE\_PUBLISH administrative
  - method, 332
- WHERE clause, 142
  - examples of use, 394
  - predicates, 29
  - valid qualifications, 142

wildcards for pattern matching, 34  
workflow agent  
  tracing, 314

## Y

YESTERDAY (keyword), 18