

# Full-Text Search Specialty Data Store User's Guide

Full-Text Search SDS Version 11.9.2

Document ID: 36521-01-1192-01

Last Revised: September 1, 1998

Principal author: Lori Johnson

Contributing authors: Martin Ash, Pam Gilpatrick, Vic Mesenzeff, Bill Seiger

Document ID: 36521-01-1192

This publication pertains to Full-Text Search SDS Version 11.9.2 of the Sybase database management software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

## Document Orders

---

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor.

Upgrades are provided only at regularly scheduled software release dates.

Copyright © 1989–1998 by Sybase, Inc. All rights reserved.

No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

## Sybase Trademarks

---

Sybase, the Sybase logo, APT-FORMS, Certified SYBASE Professional, Column Design, Data Workbench, First Impression, InfoMaker, ObjectCycle, PowerBuilder, PowerDesigner, Powersoft, Replication Server, S-Designor, SQL Advantage, SQL Debug, SQL SMART, Transact-SQL, Visual Components, VisualWriter, and VQL are registered trademarks of Sybase, Inc. Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise Monitor, Adaptive Server IQ, Adaptive Warehouse, ADA Workbench, AnswerBase, Application Manager, AppModeler, APT-Build, APT-Edit, APT-Execute, APT-Library, APT-Translator, APT Workbench, Backup Server, BayCam, Bit-Wise, ClearConnect, Client-Library, Client Services, CodeBank, Connection Manager, DataArchitect, Database Analyzer, DataExpress, Data Pipeline, DataWindow, DB-Library, dbQueue, Developers Workbench, DirectConnect, Distribution Agent, Distribution Director, Embedded SQL, EMS, Enterprise Client/Server, Enterprise Connect, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, EWA, Formula One, Gateway Manager, GeoPoint, ImpactNow, InformationConnect, InstaHelp, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Logical Memory Manager, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI

Database Gateway, media.splash, MetaBridge, MetaWorks, MethodSet, Net-Gateway, NetImpact, Net-Library, Next Generation Learning, ObjectConnect, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT-Execute, PC DB-Net, PC Net Library, Power++, Power AMC, PowerBuilt, PowerBuilt with PowerBuilder, Power Dynamo, Power J, PowerScript, PowerSite, PowerSocket, Powersoft Portfolio, PowerStudio, Power Through Knowledge, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, *QuickStart* DataMart, *QuickStart* MediaMart, *QuickStart* ReportSmart, Replication Agent, Replication Driver, Replication Server Manager, Report-Execute, Report Workbench, Resource Manager, RW-DisplayLib, RW-Library, SAFE, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, SQL Code Checker, SQL Edit, SQL Edit/TPU, SQL Modeler, SQL Remote, SQL Server, SQL Server/CFT, SQL Server/DBM, SQL Server Manager, SQL Server SNMP SubAgent, SQL Station, SQL Toolset, Sybase Central, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Gateways, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase Synergy Program, Sybase Virtual Server Architecture, Sybase User Workbench, SybaseWare, SyBooks, System 10, System 11, the System XI logo, SystemTools, Tabular Data Stream, The Enterprise Client/Server Company, The Future is Wide Open, The Learning Connection, The Model for Client/Server Solutions, The Online Information Center, Translation Toolkit, Turning Imagination Into Reality, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, VisualSpeller, WarehouseArchitect, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server, and XP Server are trademarks of Sybase, Inc. 1/98

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

## Restricted Rights

---

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., 6475 Christie Avenue, Emeryville, CA 94608.



# Table of Contents

## About This Book

Audience . . . . .	xv
How to Use This Book . . . . .	xv
Adaptive Server Enterprise Documents . . . . .	xvi
Other Sources of Information . . . . .	xviii
Conventions . . . . .	xviii
Directory Paths . . . . .	xviii
Formatting SQL Statements . . . . .	xix
SQL Syntax Conventions . . . . .	xix
Case . . . . .	xx
Obligatory Options {You Must Choose At Least One} . . . . .	xx
Optional Options [You Don't Have to Choose Any]. . . . .	xx
Ellipsis: Do It Again (and Again)... . . . .	xxi
If You Need Help . . . . .	xxi

## 1. Introduction

What Is the Full-Text Search Specialty Data Store? . . . . .	1-1
Capabilities of the Full-Text Search Engine . . . . .	1-1
Capabilities of the Enhanced Full-Text Search Engine . . . . .	1-2

## 2. Understanding the Full-Text Search Engine

Components of the Full-Text Search Engine . . . . .	2-1
Filters . . . . .	2-1
The Source Table . . . . .	2-1
The Verity Collections . . . . .	2-2
The <i>text_db</i> Database . . . . .	2-2
The <i>vesaux</i> Table . . . . .	2-3
The <i>vesauxcol</i> Table . . . . .	2-3
The Index Table . . . . .	2-3
The <i>text_events</i> Table . . . . .	2-4
Relationships Between the Components . . . . .	2-5
How a Full-Text Search Works . . . . .	2-6

## 3. Configuring Adaptive Server for Full-Text Searches

Configuring Adaptive Server for a Full-Text Search Engine . . . . .	3-1
---	-----

Enabling Configuration Parameters .....	3-1
Running the <i>installtextserver</i> Script .....	3-1
Editing the <i>installtextserver</i> Script .....	3-2
Running the <i>installtextserver</i> Script .....	3-3
Running the <i>installmessages</i> Script .....	3-3
Running the <i>installevent</i> Script .....	3-4
Editing the <i>installevent</i> Script .....	3-4
Running the <i>installevent</i> Script .....	3-5
Creating and Maintaining the Text Indexes .....	3-5
Setting Up Source Tables for Indexing .....	3-6
Adding an IDENTITY Column to a Source Table .....	3-6
Adding a Unique Index to an IDENTITY Column .....	3-7
Creating the Text Index and Index Table .....	3-7
Specifying Multiple Columns When Creating a Text Index .....	3-8
Bringing the Database Online for Full-Text Searches .....	3-9
Propagating Changes to the Text Index .....	3-9
Replicating Text Indexes .....	3-10
Example: Enabling a New Database for Text Searches .....	3-11
Step 1. Verify That the <i>text_events</i> Table Exists .....	3-11
Step 2. Check for an IDENTITY Column .....	3-12
Step 3. Create a Unique Index on the IDENTITY Column .....	3-12
Step 4. Create the Text Index and Index Table .....	3-12
Step 5. Bring the Database Online for a Full-Text Search .....	3-12

#### 4. Setting Up Verity Functions

Enabling Query-By-Example, Summarization, and Clustering .....	4-1
Editing the Master <i>style.prm</i> File .....	4-2
Editing Individual <i>style.prm</i> Files .....	4-3
Setting Up a Column to Use As a Sort Specification .....	4-4
Using Filters on Text That Contains Tags .....	4-6
Creating a Custom Thesaurus (Enhanced Version Only) .....	4-7
Examining the Default Thesaurus (Optional) .....	4-8
Creating the Control File .....	4-9
Control File Syntax .....	4-9
Creating the Thesaurus .....	4-10
Replacing the Default Thesaurus with the Custom Thesaurus .....	4-10
Creating Topics (Enhanced Version Only) .....	4-11
Creating an Outline File .....	4-12
Creating a Topic Set Directory .....	4-13
Creating a Knowledge Base Map .....	4-14

Defining the Location of the Knowledge Base Map .....	4-14
Executing Queries Against Defined Topics .....	4-14
Troubleshooting Topics .....	4-15

## 5. Writing Full-Text Search Queries

Components of a Full-Text Search Query .....	5-1
Pseudo Columns in the Index Table .....	5-2
Using the <i>score</i> Column to Relevance-Rank Search Results .....	5-3
Using the <i>sort_by</i> Column to Specify a Sort Order .....	5-4
Using the <i>summary</i> Column to Summarize Documents .....	5-6
Using Pseudo Columns to Request Clustered Result Sets .....	5-6
Preparing to Use Clustering .....	5-7
Writing Queries Requesting a Clustered Result Set .....	5-7
Full-Text Search Operators .....	5-8
Considerations When Using Verity Operators .....	5-9
Using the Verity Operators .....	5-10
<i>accrue</i> .....	5-11
<i>and, or</i> .....	5-11
<i>complement</i> .....	5-11
<i>in</i> .....	5-11
<i>like</i> .....	5-12
<i>near, near/n</i> .....	5-13
<i>or</i> .....	5-13
<i>phrase</i> .....	5-13
<i>paragraph</i> .....	5-14
<i>product</i> .....	5-14
<i>sentence</i> .....	5-14
<i>stem</i> .....	5-15
<i>sum</i> .....	5-15
<i>thesaurus</i> .....	5-15
<i>topic</i> (Enhanced Version Only) .....	5-16
<i>wildcard</i> .....	5-17
<i>word</i> .....	5-18
<i>yesno</i> .....	5-18
Operator Modifiers .....	5-19

## 6. System Administration

Starting the Full-Text Search Engine on UNIX .....	6-1
Creating the Runserver File .....	6-1
Starting the Full-Text Search Engine on Windows NT .....	6-2

Starting the Full-Text Search Engine As a Service . . . . .	6-3
<b>Shutting Down the Full-Text Search Engine</b> . . . . .	6-4
<b>Modifying the Configuration Parameters</b> . . . . .	6-4
Modifying Values in the Standard Version . . . . .	6-6
Modifying Values in the Enhanced Version . . . . .	6-7
Setting the Default Language . . . . .	6-7
Setting the Default Character Set . . . . .	6-8
Setting the Default Sort Order . . . . .	6-9
Setting Trace Flags . . . . .	6-10
Setting Open Server Trace Flags . . . . .	6-12
Setting Case Sensitivity . . . . .	6-12
<b>Backup and Recovery for the Standard Full-Text Search Engine</b> . . . . .	6-13
Backing Up Verity Collections . . . . .	6-14
Restoring Verity Collections and Text Indexes from Backup . . . . .	6-15
<b>Backup and Recovery for the Enhanced Full-Text Search Engine</b> . . . . .	6-16
Backing Up Verity Collections . . . . .	6-16
Restoring Collections and Text Indexes from Backup . . . . .	6-17

## 7. Performance and Tuning

Updating Existing Indexes . . . . .	7-1
Increasing Query Performance . . . . .	7-2
Limiting the Number of Rows . . . . .	7-2
Ensuring the Correct Join Order for Queries . . . . .	7-2
Reconfiguring Adaptive Server . . . . .	7-3
<i>cursor rows</i> . . . . .	7-3
<i>packet size</i> . . . . .	7-3
Reconfiguring the Full-Text Search Engine . . . . .	7-4
<i>batch_size</i> . . . . .	7-4
<i>min_sessions</i> and <i>max_sessions</i> . . . . .	7-4
Using <i>sp_text_notify</i> . . . . .	7-5
Configuring Multiple Full-Text Search Engines . . . . .	7-5
Creating Multiple Full-Text Search Engines at Start-Up . . . . .	7-5
Adding Full-Text Search Engines . . . . .	7-6

## A. System Procedures

<i>sp_clean_text_events</i> . . . . .	A-2
<i>sp_clean_text_indexes</i> . . . . .	A-3
<i>sp_create_text_index</i> . . . . .	A-4
<i>sp_drop_text_index</i> . . . . .	A-7



<i>sp_help_text_index</i> .....	A-9
<i>sp_optimize_text_index</i> .....	A-10
<i>sp_redo_text_events</i> .....	A-12
<i>sp_refresh_text_index</i> .....	A-14
<i>sp_show_text_online</i> .....	A-16
<i>sp_text_cluster</i> .....	A-18
<i>sp_text_configure</i> .....	A-21
<i>sp_text_dump_database</i> .....	A-23
<i>sp_text_kill</i> .....	A-26
<i>sp_text_load_index</i> .....	A-28
<i>sp_text_notify</i> .....	A-30
<i>sp_text_online</i> .....	A-31

## B. Sample Files

Default <i>textsvr.cfg</i> Configuration File .....	B-1
The <i>sample_text_main.sql</i> Script .....	B-4
Sample Files Illustrating Full-Text Search Engine Features .....	B-5
Custom Thesaurus .....	B-5
Topics .....	B-5
Clustering, Summarization, and Query-by-Example .....	B-6
<i>getsend</i> Sample Program .....	B-6

## C. Unicode Support

## Index



# List of Figures

Figure 2-1:	Components of the Full-Text Search engine .....	2-6
Figure 2-2:	Processing a full-text search query.....	2-8



# List of Tables

Table 1:	Syntax statement conventions .....	xix
Table 2-1:	Columns in the vesaux table .....	2-3
Table 2-2:	Columns in the vesauxcol table .....	2-3
Table 2-3:	Columns in the text_events table .....	2-4
Table 5-1:	Full-Text Search engine pseudo columns.....	5-2
Table 5-2:	Values for the sort_by pseudo column .....	5-5
Table 5-3:	Verity search operators .....	5-8
Table 5-4:	Alternative Verity syntax.....	5-10
Table 5-5:	Full-Text Search engine wildcard characters .....	5-17
Table 5-6:	Verity operator modifiers .....	5-19
Table 6-1:	Definition of flags in the runserver file.....	6-1
Table 6-2:	Configuration parameters .....	6-4
Table 6-3:	Configuration parameters for Enhanced version only.....	6-6
Table 6-4:	vdkLanguage configuration parameters.....	6-8
Table 6-5:	Verity character sets.....	6-9
Table 6-6:	Sort order values for the configuration file.....	6-9
Table 6-7:	Full-Text Search engine trace flags.....	6-11
Table 6-8:	Open Server trace flags .....	6-12
Table A-1:	System procedures.....	A-1
Table A-2:	Clustering configuration parameters.....	A-18
Table A-3:	Values for backupdbs .....	A-23



# About This Book

This book explains how to use the Full-Text Search Specialty Data Store product with Sybase® Adaptive Server™ Enterprise. Although this book refers to Adaptive Server throughout, the instructions for using it with OmniConnect™ are the same.

There are two versions of the Full-Text Search Specialty Data Store:

- The Standard version is included with your purchase of Adaptive Server Enterprise
- The Enhanced version is purchased separately and has additional capabilities

This book describes the features and functionality of both versions.

## Audience

---

This book is for System Administrators who are configuring Adaptive Server for a Full-Text Search Specialty Data Store and for users who are performing full-text searches on Adaptive Server data.

## How to Use This Book

---

This book includes the following chapters:

- Chapter 1, “Introduction,” provides an overview of Full-Text Search Specialty Data Store.
- Chapter 2, “Understanding the Full-Text Search Engine,” describes the components of the Full-Text Search Specialty Data Store and how it works.
- Chapter 3, “Configuring Adaptive Server for Full-Text Searches,” describes how to configure Adaptive Server so that Full-Text Search Specialty Data Store can perform full-text searches on the Adaptive Server databases.
- Chapter 4, “Setting Up Verity Functions,” describes the setup you need to do before issuing full-text search queries.
- Chapter 5, “Writing Full-Text Search Queries,” describes the components you use to write full-text search queries.
- Chapter 6, “System Administration,” provides information about system administration issues.

- Chapter 7, “Performance and Tuning,” provides information about performance and tuning issues.
- Appendix A, “System Procedures,” describes Full-Text Search Specialty Data Store system procedures.
- Appendix B, “Sample Files,” contains the text of the *textsvr.cfg* file, describes the sample files included with Full-Text Search Specialty Data Store, and discusses issues regarding the *sample\_text\_main.sql* script.
- Appendix C, “Unicode Support,” describes how to configure Full-Text Search Specialty Data Store to use Unicode.

## Adaptive Server Enterprise Documents

---

The following documents comprise the Sybase Adaptive Server Enterprise documentation:

- The *Installation and Release Bulletin* for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the *Installation and Release Bulletin* may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use Sybase Technical Library on the Web.

- The Adaptive Server installation documentation for your platform – describes installation and upgrade procedures for all Adaptive Server and related Sybase products.
- The Adaptive Server configuration documentation for your platform – describes configuring a server, creating network connections, configuring for optional functionality, such as auditing, installing most optional system databases, and performing operating system administration tasks.
- *What’s New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server release 11.5, the system changes added to support those features, and the changes that may affect your existing applications.
- *Navigating the Documentation for Adaptive Server* – an electronic interface for using Adaptive Server. This online document provides links to the concepts and syntax in the documentation that are relevant to each task.



- *Transact-SQL User's Guide* – documents Transact-SQL®, Sybase's enhanced version of the relational database language. This manual serves as a textbook for beginning users of the database management system. This manual also contains descriptions of the *pubs2* and *pubs3* sample databases.
- *System Administration Guide* – provides in-depth information about administering servers and databases. This manual includes instructions and guidelines for managing physical resources and user and system databases, and specifying character conversion, international language, and sort order settings.
- *Adaptive Server Reference Manual* – contains detailed information about all Transact-SQL commands, functions, procedures, and datatypes. This manual also contains a list of the Transact-SQL reserved words and definitions of system tables.
- *Performance and Tuning Guide* – explains how to tune Adaptive Server for maximum performance. This manual includes information about database design issues that affect performance, query optimization, how to tune Adaptive Server for very large databases, disk and cache issues, and the effects of locking and cursors on performance.
- The *Utility Programs* manual for your platform – documents the Adaptive Server utility programs, such as `isql` and `bcp`, which are executed at the operating system level.
- *Security Administration Guide* – explains how to use the security features provided by Adaptive Server to control user access to data. This manual includes information about how to add users to Adaptive Server, administer both system and user-defined roles, grant database access to users, and manage remote Adaptive Servers.
- *Security Features User's Guide* – provides instructions and guidelines for using the security options provided in Adaptive Server from the perspective of the non-administrative user.
- *Error Messages and Troubleshooting Guide* – explains how to resolve frequently occurring error messages and describes solutions to system problems frequently encountered by users.
- *Component Integration Services User's Guide for Adaptive Server Enterprise and OmniConnect* – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.

- *Adaptive Server Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Master Index for Adaptive Server Publications* – combines the indexes of the *Adaptive Server Reference Manual*, *Component Integration Services User's Guide*, *Performance and Tuning Guide*, *Security Administration Guide*, *Security Features User's Guide*, *System Administration Guide*, and *Transact-SQL User's Guide*.

### Other Sources of Information

---

Use the Sybase Technical Library CD and the Technical Library Web site to learn more about your product:

- Technical Library CD contains product manuals and technical documents and is included with your software. The DynaText browser (included on the Technical Library CD) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting Technical Library.

- Technical Library Web site is an HTML version of the Technical Library CD that you can access using a standard Web browser.

To use the Technical Library Web site, go to [www.sybase.com](http://www.sybase.com) and choose Documentation, choose Technical Library, then choose Product Manuals.

## Conventions

---

### Directory Paths

---

For readability, directory paths in this manual are in UNIX format. On Windows NT, substitute `$$SYBASE` with `%SYBASE%` and replace slashes (/) with backslashes (\). For example, replace this user input:

```
$$SYBASE/sds/text/scripts
```

with:

```
%SYBASE%\sds\text\scripts
```

## Formatting SQL Statements

---

SQL is a free-form language: there are no rules about the number of words you can put on a line or where you must break a line. However, for readability, all examples and syntax statements in this manual are formatted so that each clause of a statement begins on a new line. Clauses that have more than one part extend to additional lines, which are indented.

## SQL Syntax Conventions

---

The conventions for syntax statements in this manual are as follows:

Table 1: Syntax statement conventions

Key	Definition
<b>command</b>	Command names, command option names, utility names, utility flags, and other keywords are in <b>bold Courier</b> in syntax statements and in <b>bold Helvetica</b> in paragraph text.
<i>variable</i>	Variables, or words that stand for values that you fill in, are in <i>italics</i> .
{ }	Curly braces indicate that you choose at least one of the enclosed options. Do not include braces in your option.
[ ]	Brackets mean choosing one or more of the enclosed options is optional. Do not include brackets in your option.
( )	Parentheses are to be typed as part of the command.
	The vertical bar means you may select only one of the options shown.
,	The comma means you may choose as many of the options shown as you like, separating your choices with commas to be typed as part of the command.

- Syntax statements (displaying the syntax and all options for a command) are printed like this:

```
sp_dropdevice [device_name]
```

or, for a command with more options:

```
select column_name
       from table_name
       where search_conditions
```

In syntax statements, keywords (commands) are in normal font and identifiers are in lowercase: normal font for keywords, italics for user-supplied words.

- Examples showing the use of Transact-SQL commands are printed like this:

```
select * from publishers
```

- Examples of output from the computer are printed like this:

```
pub_id  pub_name                city      state
-----  -
0736    New Age Books            Boston    MA
0877    Binnet & Hardley         Washington DC
1389    Algodata Infosystems    Berkeley  CA
```

```
(3 rows affected)
```

## Case

---

In this manual, most of the examples are in lowercase. However, you can disregard case when typing Transact-SQL keywords. For example, **SELECT**, **Select**, and **select** are the same.

Adaptive Server's sensitivity to the case of database objects, such as table names, depends on the sort order installed on Adaptive Server. You can change case sensitivity for single-byte character sets by reconfiguring the Adaptive Server sort order. See "Changing the Default Character Set, Sort Order, or Language" in Chapter 19 of the *System Administration Guide* for more information.

## Obligatory Options {You Must Choose At Least One}

---

- **Curly Braces and Vertical Bars:** Choose **one and only one** option.

```
{die_on_your_feet | live_on_your_knees |
live_on_your_feet}
```

- **Curly Braces and Commas:** Choose one or more options. If you choose more than one, separate your choices with commas.

```
{cash, check, credit}
```

## Optional Options [You Don't Have to Choose Any]

---

- **One Item in Square Brackets:** You don't have to choose it.

```
[anchovies]
```

- **Square Brackets and Vertical Bars:** Choose **none or only one**.  
`[beans | rice | sweet_potatoes]`
- **Square Brackets and Commas:** Choose **none, one, or more than one** option. If you choose more than one, separate your choices with commas.  
`[extra_cheese, avocados, sour_cream]`

#### Ellipsis: Do It Again (and Again)...

---

An ellipsis (...) means that you can **repeat** the last unit as many times as you like. In this syntax statement, **buy** is a required keyword:

```
buy thing = price [cash | check | credit]
[, thing = price [cash | check | credit]]...
```

You must buy at least one thing and give its price. You may choose a method of payment: one of the items enclosed in square brackets. You may also choose to buy additional things: as many of them as you like. For each thing you buy, give its name, its price, and (optionally) a method of payment.

#### If You Need Help

---

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.



# 1

## Introduction

### What Is the Full-Text Search Specialty Data Store?

---

Full-Text Search Specialty Data Store (referred to in this book as the Full-Text Search engine) is an Open Server™ application built on Verity Search '97. Adaptive Server connects to the Full-Text Search engine through Component Integration Services (CIS), allowing queries written in the Verity query language to perform full-text searches on Adaptive Server data.

There are two versions of the Full-Text Search Specialty Data Store:

- The Standard version is included with your purchase of Adaptive Server Enterprise
- The Enhanced version is purchased separately and has additional capabilities

This book describes the features and functionality of both versions. For more information about the Verity product and the Verity operators used to perform full-text searches, see the Verity Web site: <http://www.verity.com>

### Capabilities of the Full-Text Search Engine

---

The Full-Text Search Specialty Data Store product performs powerful, full-text searches on Adaptive Server data. In Adaptive Server, without the Full-Text Search engine, you can search text columns only for data that matches what you specify in a select statement. For example, if a table contains documents about dog breeds, and you perform a search on the words “Saint Bernard,” the query produces only the rows that include “Saint Bernard” in the text column.

With the Full-Text Search engine, you can expand queries on text columns to do the following:

- Rank the results by order of how often a searched item appears in the selected document. For example, you can obtain a list of document titles that reference the words “Saint Bernard” five or more times.
- Select documents in which the words you search for appear within *n* number of words of each other. For example, you can search only for the documents that include the words “Saint

Bernard” and “Swiss Alps” and that appear within 10 words of each other.

- Select documents that include all the search elements you specify within a single paragraph or sentence. For example, you can query the documents that include the words “Saint Bernard” in the same paragraph or sentence as the words “Swiss Alps.”
- Select documents that contain one or more synonyms of the word you specify. For example, you can select documents that discuss “dogs,” and it returns documents that contain the words “dogs,” “canine,” “pooch,” “pup,” and so on.

## Capabilities of the Enhanced Full-Text Search Engine

---

In addition to the Full-Text Search engine capabilities described above, the Enhanced Full-Text Search engine provides additional functionality that allows you to refine your search. Using Enhanced Full-Text Search engine, you can:

- Create your own custom thesaurus. For example, you can create a custom thesaurus that includes “working dogs,” “St. Bernard,” “large dogs,” and “European Breeds” as synonyms for “Saint Bernard.”
- Create topics that specify the search criteria for a query. For example, you can create a topic that returns documents that include the phrase “Saint Bernard” or “St. Bernard,” followed by documents that include the phrase “working dogs,” “large dogs,” or “European Breeds.”
- Return documents grouped in clusters to give you a sense of the major topics covered in the documents.
- Select a section of relevant text in a document and search for other, similar documents.
- Sort documents using up to 16 sort orders. The Standard Full-Text Search engine allows only a single sort order.

Enhanced Full-Text Search engine also provides additional system administration features such as:

- Integrated backup and restore capabilities
- Ability to change the value of a configuration parameter using a system procedure
- Ability to optimize indexes for text searches when your server is inactive, to enhance performance



- **Additional system management reports for viewing setup information**
- **Ability to bring databases online automatically for text searches**



# 2

## Understanding the Full-Text Search Engine

This chapter describes how a Full-Text Search engine works. Topics include:

- Components of the Full-Text Search Engine 2-1
- How a Full-Text Search Works 2-6

### Components of the Full-Text Search Engine

---

The Full-Text Search engine uses the following components to provide full-text search capabilities:

- Filters for a variety of document types
- Source table
- Verity collections
- *text\_db* database
- Index table
- *text\_events* table

#### Filters

---

The text documents in a database can be stored in a variety of document types (Microsoft Word, SGML, HTML, FrameMaker, and so on). Verity includes a series of filters that allow you to index these document types.

You do not have to configure Adaptive Server or Full-Text Search engine to use these filters; they automatically detect the document type and apply the correct filter.

#### The Source Table

---

The **source table** is a user table maintained by Adaptive Server. It contains one or more columns using the *text*, *image*, *char*, *varchar*, *datetime*, or *small datetime* datatype, which holds the data to be searched. With the Enhanced Full-Text Search engine, the source table can also have *int*, *smallint*, and *tinyint* columns, which holds the data to be searched. The source table must have an IDENTITY

column, which is used to join with the *id* column of an index table during text searches.

The source table can be a local table, which holds the actual data, or it can be a proxy table that is mapped to remote data.

### The Verity Collections

---

The Full-Text Search engine uses the Verity collections, which are located in *\$SYBASE/sds/text/collections*. When you create the text indexes, as described in “Creating the Text Index and Index Table” on page 3-7, Verity creates a **collection**, which is a directory that implements a text index. This collection is queried by the Full-Text Search engine. For more information about Verity collections, see the Verity Web site:

<http://www.verity.com>

### The *text\_db* Database

---

During the installation of the Full-Text Search engine, a database named *text\_db* is added to Adaptive Server using the installation script *installtextserver*, as described in “Running the *installtextserver* Script” on page 3-1. The database does not contain any user data, but contains two support tables: *vesaux* and *vesauxcol*. These tables contain the metadata used by the Full-Text Search engine to maintain integrity between the Adaptive Server source tables and the Verity collections.

When updating the collections after an insert, update, or delete is made to an indexed column, the Full-Text Search engine queries the *vesaux* and *vesauxcol* tables. These tables determine which collections contain the modified columns so that all affected collections are updated. The Full-Text Search engine also uses these tables when it is brought online, to make sure that all necessary collections exist.

### The *vesaux* Table

---

The columns in the *vesaux* table are described in Table 2-1.

Table 2-1: Columns in the *vesaux* table

Column Name	Description
<i>id</i>	IDENTITY column
<i>object_name</i>	Name of the source table on which the external index is being created
<i>option_string</i>	Text index creation options
<i>collection_id</i>	Name of the Verity collection
<i>key_column</i>	Name of the IDENTITY column in the source table
<i>svrid</i>	Server ID of the Full-Text Search engine maintaining the collection

### The *vesauxcol* Table

---

The columns in the *vesauxcol* table are described in Table 2-2.

Table 2-2: Columns in the *vesauxcol* table

Column Name	Description
<i>id</i>	ID of the referenced row in the <i>vesaux</i> table
<i>col_name</i>	Name of the column for which you are searching
<i>col_type</i>	Column type ( <i>text</i> , <i>image</i> , <i>char</i> , <i>varchar</i> , <i>datetime</i> , <i>smalldatetime</i> ; with the Enhanced Full-Text Search engine, also <i>int</i> , <i>smallint</i> , and <i>tinyint</i> )

### The Index Table

---

The **index table** provides a means of locating and searching documents stored in the source table. The index table is maintained by the Full-Text Search engine and has an *id* column that maps to the IDENTITY column of the corresponding source table. The IDENTITY value from the row in the source table is stored with the data in the Verity collections, which allows the source and index tables to be joined. Although the index table is stored and maintained by the Full-Text Search engine, it functions as a local

table to Adaptive Server through the Component Integration Services feature.

The index table contains special columns, called **pseudo columns**, that are used by the Full-Text Search engine to determine the parameters of the search and the location of the text data in the source table. Pseudo columns have no associated physical storage—the values of a pseudo column are valid only for the duration of the query and are removed immediately after the query finishes running.

For example, when you use the *score* pseudo column in a query, to rank each document according to how well the document matches the query, you may have to use a *score* of 15 to find references to the phrase “small Saint Bernards” in the text database. This phrase does not occur very often, and a low *score* value broadens the search to include documents that have a small number of occurrences of the search criteria. However, if you are searching for a phrase that is common, like “large Saint Bernards,” you could use a *score* of 90, which would limit the search to those documents that have many occurrences of the search criteria.

You use the *score* column and the other pseudo columns, *id*, *index\_any*, *sort\_by*, *summary*, and *max\_docs*, to specify the parameters to include in your search. For a description of the pseudo columns, see “Pseudo Columns in the Index Table” on page 5-2.

### The *text\_events* Table

Each database containing tables referenced by a text index must contain an **events table**, which logs inserts, updates, and deletes to indexed columns. The name of this table is *text\_events*. It is used to propagate updated data to the Verity collections.

The columns in the *text\_events* table are described in Table 2-3.

Table 2-3: Columns in the *text\_events* table

Column Name	Description
<i>event_id</i>	IDENTITY column.
<i>id</i>	ID of the row that was updated, inserted, or deleted.
<i>tableid</i>	Name of the table that contains the row that was updated, inserted, or deleted.

**Table 2-3: Columns in the text\_events table (continued)**

Column Name	Description
<i>columnid</i>	Name of the column that the text index was created on.
<i>event_date</i>	Date and time of the update, insert, or delete.
<i>event_type</i>	Type of update (update, insert, or delete).
<i>event_status</i>	Indicates whether the update, insert, or delete has been propagated to the collections.
<i>srv_id</i>	Server ID of the Full-Text Search engine maintaining the collection.

### Relationships Between the Components

The relationships between the Full-Text Search engine components are shown in Figure 2-1.




Figure 2-1: Components of the Full-Text Search engine

## How a Full-Text Search Works

To perform a full-text search, you enter a select statement that joins the `IDENTITY` column from the source table with the `id` column of the index table, using pseudo columns as needed to define the search. For example, the following query searches for documents in the `blurbs` table of the `pubs2` database in which the word "Greek" appears near the word "Gustibus" (the `i_blurbs` table is the index table):

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 20
and t1.max_docs = 10
and t1.index_any = "<near>(Greek, Gustibus)"
```

Adaptive Server and the Full-Text Search engine split the query processing, as follows:



1. The Full-Text Search engine processes the query:

```
select t1.score, t1.id
from i_blurbs t1
where t1.score > 20
and t1.max_docs = 10
and t1.index_any = "<near>(Greek, Gustibus)"
```

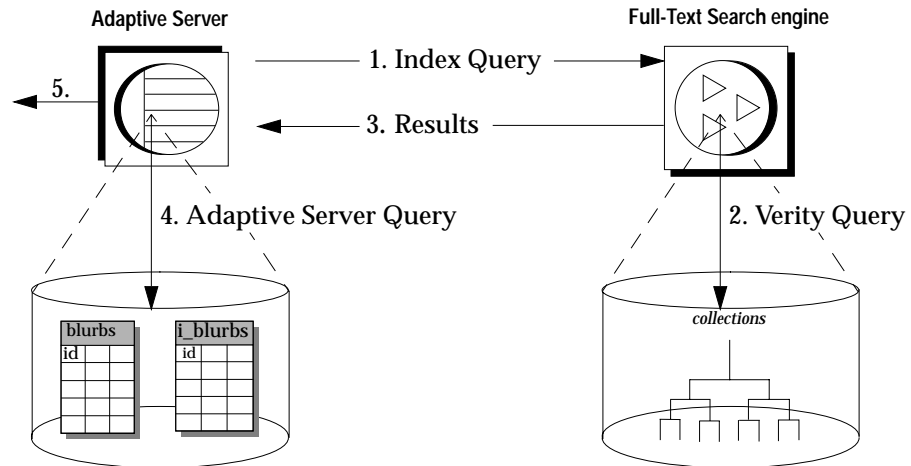
The select statement includes the Verity operator `near` and the pseudo columns `score`, `max_docs`, and `index_any`. The operator and pseudo columns provide the parameters for the search on the Verity collections—they narrow the result set from the entire `copy` column to the 10 documents in which the words “Greek” and “Gustibus” appear closest to each other.

2. Adaptive Server processes the following select statement on the result set that is returned by the Full-Text Search engine in step 1:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
```

This joins the `blurbs` and `i_blurbs` tables (the source table and the index table, respectively) on the `IDENTITY` column of the `blurbs` table and the `id` column of the `i_blurbs` table.

Figure 2-1 describes how Adaptive Server and the Full-Text Search engine process the query.



1. Adaptive Server sends the index query to the Full-Text Search engine.
2. The Full-Text Search engine processes the Verity operators in the query and produces a result set from the collections.
3. The Full-Text Search engine returns the result set to Adaptive Server.
4. Adaptive Server processes the select statement on the local table.
5. Adaptive Server displays the results of the query.

Figure 2-2: Processing a full-text search query

# 3

## Configuring Adaptive Server for Full-Text Searches

This chapter describes how to configure Adaptive Server to perform full-text searches. Topics include:

- Configuring Adaptive Server for a Full-Text Search Engine 3-1
- Creating and Maintaining the Text Indexes 3-5

### Configuring Adaptive Server for a Full-Text Search Engine

---

The Full-Text Search engine is a remote server that Adaptive Server connects to through Component Integration Services (CIS). Before you can use the Full-Text Search engine, configure Adaptive Server for a Full-Text Search engine as follows:

- Enable the `enable cis` and `rpc` configuration parameters if you have not done so
- Run the `installtextserver` script to define one or more Full-Text Search engines
- Run the `installmessages` script to install messages for the Full-Text Search engine's system procedures
- Run the `installevent` script to create the `text_events` table in the user database

### Enabling Configuration Parameters

---

To connect to the Full-Text Search engine, Adaptive Server must be running with the `enable cis` and `rpc` configuration parameters enabled. If those parameters are not enabled, log in to Adaptive Server using `isql` and use `sp_configure` to enable them. For example:

```
exec sp_configure "enable cis", 1
exec sp_configure "rpc", 1
```

Adaptive Server displays a series of messages stating that you have altered a configuration parameter and that Adaptive Server must be rebooted for the new configuration parameters to take effect.

### Running the `installtextserver` Script

---

The `installtextserver` script:

- Defines the Full-Text Search engine as a remote server of server class *sds* to Adaptive Server.
- Creates a database for storing text index metadata. For more information about this database, see “The *text\_db* Database” on page 2-2.
- Installs the system procedures required by the Full-Text Search engine.

Run the `installtextserver` script only once (see “Running the `installtextserver` Script” on page 3-3). To add Full-Text Search engines at a later time, use `sp_addserver`. See “Configuring Multiple Full-Text Search Engines” on page 7-5 for more information about `sp_addserver`.

All Full-Text Search engines use the same database for storing text index metadata. This database is referred to in this book as the *text\_db* database, the default name.

For a list and description of the system procedures added with the `installtextserver` script, see Appendix A, “System Procedures.”

#### Editing the `installtextserver` Script

The `installtextserver` script is located in the `$$SYBASE/sds/text/scripts` directory. Use a text editor (such as `vi` or `emacs`) to open the script, and make your edits. The edits you can make are as follows:

- Changing the name of the *text\_db* database. If you use a different name, replace all occurrences of *text\_db* with the appropriate name.

#### ► Note

---

If you change the name of the *text\_db* database, you must change the name in the `defaultDb` configuration parameter (see “Modifying the Configuration Parameters” on page 6-4).

---

- Changing the name of the Full-Text Search engine. By default, the `installtextserver` script defines a Full-Text Search engine named “`textsvr`.” If your Full-Text Search engine is named differently, edit this script so that it defines the correct server name.
- Adding multiple Full-Text Search engines (for information on how this can enhance performance, see “Configuring Multiple Full-Text Search Engines” on page 7-5). If you are initially defining more than one Full-Text Search engine, edit the

`installtextserver` script so that it includes all the Full-Text Search engine definitions. `installtextserver` includes the following section for naming the Full-Text Search engine you are configuring (“`textsvr`” by default):

```
/*
** Add the text server
*/
exec sp_addserver textsvr,sds,textsvr
go
```

Add an entry for each Full-Text Search engine you are configuring. For example, if you are configuring three Full-Text Search engines named `KRAZYKAT`, `OFFICAPUP`, and `MOUSE`, replace the default “`textsvr`” line with the following lines:

```
exec sp_addserver KRAZYKAT, sds, KRAZYKAT
exec sp_addserver OFFICAPUP, sds, OFFICAPUP
exec sp_addserver MOUSE, sds, MOUSE
go
```

- If you use OmniConnect to communicate with the Full-Text Search engine, change the server name specification in the `sp_addobjectdef` calls for the `vesaux` and `vesauxcol` tables to a valid remote server. For example, if your remote server is named `REMOTE`, change the lines:

```
exec sp_addobjectdef "vesaux", "SYBASE.master.dbo.vesaux", "table"
exec sp_addobjectdef "vesauxcol", "SYBASE.master.dbo.vesauxcol", "table"
```

to something similar to:

```
exec sp_addobjectdef "vesaux", "REMOTE.master.dbo.vesaux", "table"
exec sp_addobjectdef "vesauxcol", "REMOTE.master.dbo.vesauxcol", "table"
```

### Running the *installtextserver* Script

---

Use `isql` to run the `installtextserver` script. For example, to run the `installtextserver` script in an Adaptive Server named `MYSVR`, enter:

```
isql -Usa -P -SMYSVR -i
$SYBASE/sds/text/scripts/installtextserver
```

### Running the *installmessages* Script

---

The Full-Text Search engine has its own set of system procedure messages that you must install in Adaptive Server. Use the `installmessages` script to install the messages. You run the `installmessages` script only once, even if you have multiple Full-Text Search engines.

For example, to run the `installmessages` script in a server named `MYSVR`, enter:

```
isql -Usa -P -SMYSVR -i $SYBASE/sds/text/scripts/installmessages
```

### Running the *installevent* Script

---

Each database containing tables referenced by a text index must contain a `text_events` table, which logs inserts, updates, and deletes to indexed columns. It is used to propagate updated data to the Verity collections.

Run the `installevent` script, as described below, to create the `text_events` table and associated system procedures in a database. Use the `installevent` script as follows:

- If all databases require text indexes, run the `installevent` script to create a `text_events` table in the `model` database. Each newly created database will then have a `text_events` table. To add a `text_events` table to existing databases, edit the script as described below to create the `text_events` table in the existing user database.
- If not all databases have text indexes, use the `installevent` script as a sample. For each existing database and each new database that includes tables that require text indexing, run the `installevent` script. You must edit the script as described below, to create the `text_events` table in the correct user database.

► **Note**

---

If a `text_events` table does not exist in a database that includes source tables that require text indexing, changes to the source table will not be propagated to the Verity collections.

---

### Editing the *installevent* Script

---

The `installevent` script is located in the `$SYBASE/sds/text/scripts` directory. Use a text editor (such as `vi` or `emacs`) to open the script, and make the edits. The edits you can make are:

- Changing the user database name. The `installevent` script creates an events table (named `text_events`) and associated system procedures in the `model` database. The `model` database is the default database. To install the `text_events` table in an existing user

database, edit the script and replace all references to *model* with the user database name.

- Changing the *text\_db* database name. If your database for storing text index metadata is named something other than *text\_db*, replace all references to *text\_db* with the appropriate name.

► **Note**

---

The name of the *text\_db* database must be the same as the name in the `defaultDb` configuration parameter (see “Modifying the Configuration Parameters” on page 6-4).

---

### Running the *installevent* Script

---

► **Note**

---

The *text\_db* database must exist before you run the `installevent` script. If it does not exist, run the `installtextserver` script first.

---

Using `isql`, run the `installevent` script to install the *text\_events* table and related system procedures in Adaptive Server. For example, to run the `installevent` script in a server named `MYSVR`, enter:

```
isql -Usa -P -SMYSVR -i $SYBASE/sds/text/scripts/installevent
```

## Creating and Maintaining the Text Indexes

---

Before the Full-Text Search engine can process full-text searches, you must create text indexes for the source tables in the user database. After the text indexes are created, you must update them when the source data changes to keep the text indexes current. To create and maintain the text indexes:

- Set up the source table for indexing (see “Setting Up Source Tables for Indexing” on page 3-6).
- Create the text indexes and index tables (see “Creating the Text Index and Index Table” on page 3-7).
- Bring the databases online for full-text searches (see “Bringing the Database Online for Full-Text Searches” on page 3-9).
- Propagate changes in the user data to the text indexes (see “Propagating Changes to the Text Index” on page 3-9).

- If you are replicating text indexes, set up text indexing in the destination database (see “Replicating Text Indexes” on page 3-10).

For an example of setting up a text index, see the sample script `sample_text_main.sql` in the `$SYBASE/sds/text/sample/scripts` directory.

### Setting Up Source Tables for Indexing

The source table contains the data on which you perform searches (for example, the *blurbs* table in the *pubs2* database). For more information on source tables, see “The Source Table” on page 2-1.

Before you can create text indexes on a source table, you must:

- Verify that the source table has an IDENTITY column
- Create a unique index on the IDENTITY column (optional)

### Adding an IDENTITY Column to a Source Table

Every source table must contain an IDENTITY column to uniquely identify each row and provide a means of joining the index table and the source table. When you create a text index, the IDENTITY column is passed with the indexed columns to the Full-Text Search engine. The IDENTITY column value is stored in the text index and is mapped to the *id* column in the index table.

The IDENTITY column must have sufficient precision and scale to guarantee a unique IDENTITY for each row. Sybase recommends a precision of 10 and a scale of 0. You can use an existing IDENTITY column, if it is defined with sufficient precision and scale to identify each row uniquely.

For example, to create an IDENTITY column in a table named *composers*, define the table as follows:

```
create table composers (
    id          numeric(10,0)  identity,
    comp_fname  char(30)       not null,
    comp_lname  char(30)       not null,
    text_col    text
)
```

To add an IDENTITY column to an existing table, enter:

```
alter table table_name add id numeric(10,0) identity
```



### Adding a Unique Index to an IDENTITY Column

---

For optimum performance, Sybase recommends creating a unique index on the IDENTITY column. For example, to create a unique index named *comp\_id* on the IDENTITY column created above, enter:

```
create unique index comp_id
on composers(id)
```

For more information about creating a unique index, see Chapter 11, "Creating Indexes on Tables," of the *Transact-SQL User's Guide*.

### Creating the Text Index and Index Table

---

Use the `sp_create_text_index` system procedure to create the text indexes. `sp_create_text_index` does the following:

- Updates the *vesaux* and *vesauxcol* tables in the *text\_db* database
- Creates the text index (Verity collections)
- Populates the Verity collections
- Creates the index table in the user database where the source table is located

The text index can contain up to 16 columns. Columns of the following datatypes can be indexed:

#### Standard Version Datatypes

*char*, *varchar*, *nchar*, *nvarchar*, *text*, *image*, *datetime*, *smalldatetime*

#### Enhanced Version Datatypes

All Standard version datatypes, plus:

*int*, *smallint*, and *tinyint*

For example, to create a text index and an index table named *i\_blurbs* for the *copy* column in the *blurbs* table in *pubs2* on KRAZYKAT, enter:

```
sp_create_text_index "KRAZYKAT", "i_blurbs", "blurbs", " ",
"copy"
```

where:

- KRAZYKAT is the name of the Full-Text Search engine
- i\_blurbs is the name of the index table and text index you are creating

- `blurbs` is the source table on which you are creating the text indexes
- `" "` is a placeholder for text index creation options
- `copy` is the column in the `blurbs` table that you are indexing

See “`sp_create_text_index`” on page A-4 for more information.

► **Note**

---

Make sure the `text_db` database name in the configuration file (listed after the `defaultDb` parameter) matches the database name in Adaptive Server. If they do not match, the text index cannot be created. Also, verify that the `text_events` table exists in the user database. If it does not exist, run the `installevent` script for that database (refer to “Running the `installevent` Script” on page 3-4).

---

Populating the Verity collections can take a few minutes or several hours, depending on the amount of data you are indexing. You may want to perform this step when the server is not being heavily used. Increasing the `batch_size` configuration parameter will also expedite the process. See “`batch_size`” on page 7-4 for more information.

#### Specifying Multiple Columns When Creating a Text Index

When you create a text index on two or more columns, each column in the text index is placed into its own document zone. The name of the zone is the name of the column. For example, to create a text index and an index table named `i_blurbs` for both the `copy` column and the `au_id` column in the `blurbs` table in `pubs2` on `KRAZYKAT`, enter:

```
sp_create_text_index "KRAZYKAT", "i_blurbs", "blurbs", " ",
"copy", "au_id"
```

`sp_create_text_index` creates two zones in the text index named “`copy`” and “`au_id`.” When you issue a query against the `i_blurbs` text index, the search includes the `copy` and `au_id` columns. However, you can limit your search to a particular column by using the `in` operator to specify a document zone (for more information, see “`in`” on page 5-11).

## Bringing the Database Online for Full-Text Searches

---

With the Standard version of Full-Text Search engine, you must manually bring a database online before issuing full-text queries on a source table in the database. When you bring a database online, the Full-Text Search engine initializes the internal Verity structures and confirms that the Verity collections exist.

► **Note**

---

With the Enhanced Full-Text Search engine, the database is automatically brought online when the `auto_online` configuration parameter is set to 1.

---

Use the `sp_text_online` system procedure to bring a database online for full-text searches if it is not automatically brought online. For example, to bring the `pubs2` database online before issuing full-text searches on the `blurbs` table in a Full-Text Search engine named KRAZYKAT, enter:

```
sp_text_online KRAZYKAT, pubs2
```

This message appears:

```
Database 'pubs2' is now online
```

The `pubs2` database is now available for performing full-text searches.

See “`sp_text_online`” on page A-31 for more information.

## Propagating Changes to the Text Index

---

When you insert, update, or delete data in your source table, the text indexes are not updated automatically. After you update data, run the `sp_refresh_text_index` system procedure to log the changes to the `text_events` table. Then, run the `sp_text_notify` system procedure to notify the Full-Text Search engine that changes need to be processed. The Full-Text Search engine then connects to Adaptive Server, reads the entries in the `text_events` table, determines which indexes, tables, and rows are affected, and updates the appropriate collections.

See “`sp_refresh_text_index`” on page A-14 and “`sp_text_notify`” on page A-30 for more information on these system procedures.

To have `sp_refresh_text_index` run automatically after each insert, update, or delete, you can create triggers on your source tables, as follows:

- Create a trigger that runs `sp_refresh_text_index` after a delete operation.
- Create a trigger that runs `sp_refresh_text_index` after an insert operation.
- Create a trigger that runs `sp_refresh_text_index` after an update operation to an indexed column.

Triggers are not fired when you use `writetext` to update a *text* column. To have `sp_refresh_text_index` automatically run after a `writetext`:

- Set up a non-*text* column and update that column after each `writetext`.
- Create a trigger on the non-*text* column to run `sp_refresh_text_index`. Since the Full-Text Search engine reinserts the entire row when you issue `sp_text_notify`, the update to the *text* column gets propagated to the text index.

For examples of each of these triggers, see the sample script `sample_text_main.sql` in the `$SYBASE/sds/text/sample/scripts` directory.

## Replicating Text Indexes

---

To replicate tables that have text indexes, follow these guidelines:

- Create the table definition in the destination database.
- Run the `installevent` script to create the *text\_events* table in the destination database, if the *text\_events* table does not already exist (see “Running the `installevent` Script” on page 3-4).
- Run `sp_create_text_index` to create the text index on the empty table in the destination database (see “Creating the Text Index and Index Table” on page 3-7).
- Create triggers for running `sp_refresh_text_index` to insert entries into the *text\_events* table whenever you insert, update, or delete data into the table (see “Propagating Changes to the Text Index” on page 3-9).
- Create the replication definition in the Replication Server. This replicates all the data in the source table to the destination table.
- Run `sp_text_notify` to update the text index; run `sp_text_notify` periodically to process changes to the destination table (see “Propagating Changes to the Text Index” on page 3-9).

---

**► Note**

You must issue an update against a non-*text* column whenever a `writetext` command is performed. This ensures that the trigger that inserts data into the *text\_events* table is fired.

---

### Example: Enabling a New Database for Text Searches

---

This example describes the steps for creating a text index on the *plot* column of the *reviews* table in the *movies* database. This process assumes that:

- You have created a *reviews* table in a new database named *movies* on the MYSVR server
- The *reviews* table has a column named *plot* that you are going to index
- Adaptive Server and the Full-Text Search engine named MYTXTSVR have been configured to connect to each other

#### Step 1. Verify That the *text\_events* Table Exists

---

Each database containing tables referenced by a text index must contain a *text\_events* table, which logs inserts, updates, and deletes to indexed columns.

If a *text\_events* table is in your *model* database, it will be in all new databases. If a *text\_events* table is not in your *model* database, run the `installevent` script to install the *text\_events* table in the new database. For example, to install the *text\_events* table in the *movies* database:

- Save the `installevent` script as `installeventmovies`.
- Edit the script to replace all references to the word *model* with the word *movies*.
- Run the script as follows:

```
isql -Usa -P -SMYSVR -i
$SYBASE/sds/text/scripts/installeventmovies
```

See “Running the `installevent` Script” on page 3-4 for information on installing the *text\_events* table.

### Step 2. Check for an IDENTITY Column

---

Every source table must contain an IDENTITY column, which uniquely identifies each row and provides a means of joining the index table and the source table.

For example, to add an IDENTITY column to the *reviews* table, enter:

```
alter table reviews add id numeric(10,0) identity
```

See “Adding an IDENTITY Column to a Source Table” on page 3-6 for more information on creating an IDENTITY column.

### Step 3. Create a Unique Index on the IDENTITY Column

---

This step is optional. To enhance performance, Sybase recommends creating a unique index that contains only the IDENTITY column.

For example, to create a unique index named *reviews\_id* on the IDENTITY column created in step 2, issue the command:

```
create unique index reviews_id on reviews(id)
```

For more information about creating a unique index, see Chapter 11, “Creating Indexes on Tables,” of the *Transact-SQL User’s Guide*.

### Step 4. Create the Text Index and Index Table

---

The source tables in the user database need to be indexed so that you can perform full-text searches. For example, to create a text index and an index table named *reviews\_idx* for the *plot* column in the *reviews* table, enter:

```
sp_create_text_index "MYTXTSVR", "reviews_idx", "reviews", " ",  
"plot"
```

The *reviews* table is now available for running full-text searches.

See “sp\_create\_text\_index” on page A-4 for more information.

### Step 5. Bring the Database Online for a Full-Text Search

---

To bring the *movies* database online for the Full-Text Search engine named MYTXTSVR, enter:

```
sp_text_online MYTXTSVR, movies
```

**► Note**

---

Omit this step if you have Enhanced Full-Text Search engine and your `auto_online` configuration parameter is set to "1".

---

See "sp\_text\_online" on page A-31 for more information.





# 4

## Setting Up Verity Functions

This chapter describes the setup required before you can write queries with certain Verity functionality. It includes:

- Enabling Query-By-Example, Summarization, and Clustering 4-1
- Setting Up a Column to Use As a Sort Specification 4-4
- Using Filters on Text That Contains Tags 4-6
- Creating a Custom Thesaurus (Enhanced Version Only) 4-7
- Creating Topics (Enhanced Version Only) 4-11

### Enabling Query-By-Example, Summarization, and Clustering

---

The *style.prm* file specifies additional data to include in the text indexes to support the following functionality:

- Query-by-example – Retrieves documents that are similar to a phrase (see “like” on page 5-12 for more information).

► **Note**

---

The text indexes only need additional data to support phrases in the query-by-example specification of the *like* operator. If you use a document in the query-by-example specification, additional data is not required.

---

- Summarization – returns summaries of documents rather than entire documents (see “Using the summary Column to Summarize Documents” on page 5-6 for more information).
- Clustering – groups documents in result sets by subtopic (see “Using Pseudo Columns to Request Clustered Result Sets” on page 5-6 for more information). Clustering is available only with the Enhanced Full-Text Search engine.

You can enable these features for all text indexes by editing the master *style.prm* file, or you can enable them for an individual text index by editing its *style.prm* file. Both methods are describe below.

#### Query-By-Example and Clustering

To use phrases in a query-by-example specification and to use clustering, you must enable the generation of document feature

vectors at indexing time. To do this, uncomment the following line in the *style.prm* file:

```
$define DOC-FEATURES "TF"
```

### Summarization

To configure the Full-Text Search engine for summarization, uncomment one of the following lines that starts with “#\$define” in the *style.prm* file :

```
# The example below stores the best three sentences of
# the document, but not more than 255 bytes.
#$define DOC-SUMMARIES  "XS MaxSents 3 MaxBytes 255"

# The example below stores the first four sentences of
# the document, but not more than 255 bytes.
#$define DOC-SUMMARIES  "LS MaxSents 4 MaxBytes 255"

# The example below stores the first 150 bytes of
# the document, with whitespace compressed.
#$define DOC-SUMMARIES  "LB MaxBytes 150"
```

Each of those lines reflects a different level of summarization. You can specify how many bytes of data you want the Full-Text Search engine to display, by altering the numbers at the ends of these lines. For example, if you want only the first 233 bytes of data summarized, edit the script to read:

```
$define DOC-SUMMARIES  "LS MaxSents 4 MaxBytes 233"
```

The maximum number of bytes displayed is 255. Any number greater than that is truncated to 255.

### Editing the Master *style.prm* File

The master *style.prm* file is located in:

```
SSYBASE/sds/text/verity/common/style
```

It contains the default Full-Text Search engine style parameters. Edit this file to configure the Full-Text Search engine so that all tables on which you create text indexes allow clustering and literal text in your query-by-example specifications, or summarization. Uncomment the applicable lines as described above.

**► Note**


---

If you have existing text indexes, you must re-create the text index with these features enabled as described in “Editing Individual *style.prm* Files” below.

---

### Editing Individual *style.prm* Files

---

Perform the following steps to configure the Full-Text Search engine so that the individual text index allows clustering and literal text in your query-by-example specifications, or summarization:

1. Create the text index using `sp_create_text_index`. Use the word “empty” in the *option\_string* parameter so that the *style.prm* file is created for the text index, but the Verity collections are not populated with data. For example, if you are enabling clustering for the *copy* column of the *blurbs* table, use the following syntax:

```
sp_create_text_index "KRAZYKAT", "i_blurbs",
"blurbs", "empty", "copy"
```

**► Note**


---

If the text index already exists, omit this step. You do not need to create the text index again.

---

2. Use `sp_drop_text_index` to drop the text index associated with the *style.prm* file you are editing.

For example, to drop the text index created in step 1, enter:

```
sp_drop_text_index "blurbs.i_blurbs"
```

3. Edit the *style.prm* file that exists for the text index. The *style.prm* file for an existing collection is located in:

```
$$SYBASE/sds/text/collections/db.owner.index/style
```

where *db.owner.index* is the database, the database owner, and the index created with `sp_create_text_index`. For example, if you create a text index called *i\_blurbs* on the *pubs2* database, the full path to these files is:

```
$$SYBASE/sds/text/collections/pubs2.dbo.i_blurbs/style
```

4. Uncomment the applicable lines as described above.

For example, to enable clustering, uncomment the following line:

```
$define DOC-FEATURES "TF"
```

5. Re-create the text index you dropped in step 2. For example, to re-create the *i\_blurbs* text index, enter:

```
sp_create_text_index "KRAZYKAT", "i_blurbs", "blurbs", "", "copy"
```

## Setting Up a Column to Use As a Sort Specification

Before you can sort by specific columns, you must modify the *style.vgw* and *style.ufl* files. (For information on including a column in a sort specification, see “Using the *sort\_by* Column to Specify a Sort Order” on page 5-4.) Both files are in the directory:

```
SSYBASE/sds/text/collections/db.owner.index/style
```

where *db.owner.index* is the database, the database owner, and the index created using *sp\_create\_text\_index*. For example, if you created a text index called *i\_blurbs* on the *pubs2* database, the full path to those files would be similar to:

```
SSYBASE/sds/text/collections/pubs2.dbo.i_blurbs/style
```

To edit the *style.vgw* and *style.ufl* files, follow these steps:

1. Drop the text index that contains the columns for which you are adding definitions.

For example, to add definitions for the *copy* column in the *blurbs* table, use the following command to drop the text index:

```
sp_drop_text_index i_blurbs
```

2. Edit the *style.vgw* file. Following this line:

```
dda "SybaseTextServer"
```

add an entry for the column you are defining. The syntax is:

```
table: DOCUMENTS
{
    copy: fcolumn_number copy_column_number
}
```

where *column\_number* is the number of the column you are defining. Column numbers start with 0; if you want the first column to be sorted, specify “f0”; to sort the second column, specify “f1”; to sort the third column, specify “f2”, and so on.

For example, to define the first column in a table, the syntax is:

```

table: DOCUMENTS
{
    copy: f0 copy_f0
}

```

Then, your *style.vgw* file will be similar to this:

```

#
#       Sybase Text Server Gateway
#
$control: 1
gateway:
{
    dda:    "SybaseTextServer"
    {
        copy: f0 copy_f0
    }
}

```

3. Edit the *style.uff* file, by adding the column definition for a data table named *fts*. The syntax is:

```

data-table:    fts
{
    fixwidth:  copy_fcolumn_number precision datatype
}

```

Column numbers start with 0; if you want the first column to be sorted, specify "f0"; to sort the second column, specify "f1"; to sort the third column, specify "f2", and so on. For example, to add a definition for the first column of a table, with a precision of 4, and a datatype of *date*, enter:

```

data-table:  fts
{
    fixwidth:  copy_f0 4    date
}

```

Similarly, to add a definition for the second column of a table with a precision of 10, and a datatype of *character*, enter:

```

data-table:  fts
{
    fixwidth:  copy_f1 10  text
}

```

4. Re-create the index, using *sp\_create\_text\_index*.

## Using Filters on Text That Contains Tags

---

To perform accurate searches on documents that contain tags (such as HTML or postscript), the text index must use a filter to strip out the tags. When you create the text index to use a filter, the data for each type of tag in the document is placed into its own document zone.

For example, if you have a tag called “chapter,” all chapter names are placed into one document zone. You can issue a query that searches the entire document, or that searches only for data in the “chapter” zone (for more information, see “in” on page 5-11).

To create a text index that uses a filter, modify the *style.dft* file for that text index. To edit the *style.dft* file, follow these steps:

1. Create the text index using `sp_create_text_index`. Use the word “empty” in the *option\_string* parameter so that the *style.dft* file is created for the text index, but the Verity collections are not populated with data. For example, to create a text index for the *copy* column of the *blurbs* table, use the following syntax:

```
sp_create_text_index "KRAZYKAT", "i_blurbs",
"blurbs", "empty", "copy"
```

---

◆ **WARNING!**

**You should specify only one column in the text index when the text index uses a filter.**

---

2. Drop the text index that you create in step 1. This drops the text index, but not the *style.dft* file. For example, use the following command to drop the *i\_blurbs* text index:

```
sp_drop_text_index i_blurbs
```

3. Edit the *style.dft* file. The *style.dft* file is in the directory:

```
$$SYBASE/sds/text/collections/db.owner.index/style
```

where *db.owner.index* is the database, the database owner, and the index created using `sp_create_text_index`. For example, if you created a text index called *i\_blurbs* on the *pubs2* database, the full path to the *style.dft* file would be similar to:

```
$$SYBASE/sds/text/collections/pubs2.dbo.i_blurbs/style
```

Following this line:

```
field: f0
```

add the following syntax to use a filter:

```
/filter="universal"
```

Then, your *style.dft* file will look like this:

```
$control: 1
dft:
{
    field: f0
        /filter="universal"
    field: f1
    field: f2
    .
    .
    field: f15
}
```

4. Re-create the index, using `sp_create_text_index`. For example:

```
sp_create_text_index "KRAZYKAT", "i_blurbs", "blurbs",  
"", "copy"
```

## Creating a Custom Thesaurus (Enhanced Version Only)

The Verity thesaurus operator expands a search to include the specified word and its synonyms (for information on using the thesaurus operator, see “thesaurus” on page 5-15). In the Enhanced version of the Full-Text Search engine, you can create a custom thesaurus that contains application-specific synonyms to use in place of the default thesaurus.

For example, the default English language thesaurus contains these words as synonyms for “money”: “cash,” “currency,” “lucre,” “wampum,” and “greenbacks.” You can create a custom thesaurus that contains a different set of synonyms for “money”; for example, such as: “bid,” “tokens,” “credit,” “asset,” and “verbal offer.”

To create a custom thesaurus, follow these steps:

1. Make a list of the synonyms that you will use with your application. It may help to examine the default thesaurus (see “Examining the Default Thesaurus (Optional)” on page 4-8).
2. Create a control file that contains the synonyms you are defining for your custom thesaurus (see “Creating the Control File” on page 4-9).

3. Create the custom thesaurus using the `mksyd` utility (see “Creating the Thesaurus” on page 4-10). This uses the control file as input.
4. Replace the default thesaurus with your custom thesaurus (see “Replacing the Default Thesaurus with the Custom Thesaurus” on page 4-10).

For more information on “Custom Thesaurus Support” and the `mksyd` utility, see the Verity Web site at:

<http://www.verity.com>

In the Enhanced version of Full-Text Search engine, two sample files illustrate how to set up and use a custom thesaurus:

- `sample_text_thesaurus.ctl` is a sample control file
- `sample_text_thesaurus.sql` issues queries against the custom thesaurus defined in the sample control file

These files are in the `$$SYBASE/sds/text/sample/scripts` directory.

### Examining the Default Thesaurus (Optional)

A control file contains all the synonym definitions for a thesaurus. To examine the default thesaurus, create its control file using the `mksyd` utility. Use the syntax:

```
mksyd -dump -syd  
$SYBASE/sds/text/verity/common/vdkLanguage/vdk20.syd -f  
work_location/control_file.ctl
```

where:

- `vdkLanguage` – is the value of the `vdkLanguage` configuration parameter (for example, “english0”)
- `work_location` – is the directory where you want to place the control file
- `control_file` – is the name of the control file you are creating from the default thesaurus

Examine the control file (`control_file.ctl`) that it creates to view the default synonym lists.



## Creating the Control File

---

Create a control file that contains the new synonyms for your custom thesaurus. The control file is an ASCII text file in a structured format. Using a text editor (such as vi or emacs), either:

- Edit the control file from the default thesaurus and add new synonyms to the existing thesaurus (see “Examining the Default Thesaurus (Optional)” on page 4-8), or
- Create a new control file that includes only your synonyms

## Control File Syntax

---

The control file contains synonym list definitions in a **synonyms:** statement. For example, the following is a control file named *colors.ctl*:

```
$control: 1
synonyms:
{
list: "red, ruby, scarlet, fuchsia,\
magenta"
list: "electric blue <or> azure"
/keys = "lapis"
}
$$
```

The **synonyms:** statement includes:

- **list:** keywords that specify the start of a synonym list. The synonyms in the list are either in query form or in a list of words or phrases separated by commas.
- Each list: can optionally have a */keys* modifier that specifies one or more keys separated by commas. In the example above, no keys are specified in the first “list”. This means the list is found when the thesaurus is queried for the word “red,” “ruby,” “scarlet,” “fuchsia,” or “magenta.” The second “list” uses the */keys* modifier to specify one key. This means the words or phrases in the list will satisfy a query only when you specify <thesaurus>lapis.

---

► **Note**

If you use `emacs` to build a synonym list and any of your lists go beyond one line, turn off `auto-fill` mode. If you separate your list into multiple lines, include a backslash (`\`) at the end of each line so that the lines are treated as one list.

---

For more complex examples of control files, see the Verity Web site.

---

### Creating the Thesaurus

The `mksyd` utility creates the custom thesaurus using a control file as input. It is located in:

`$$SYBASE/sds/text/verity/bin`

Run, or define an alias to run, `mksyd` from this `bin` directory. Create your custom thesaurus in any work directory.

The `mksyd` syntax for creating a custom thesaurus is:

```
mksyd -f control_file.ctl -syd custom_thesaurus.syd
```

where:

- `control_file` – is the name of the control file you create in “Creating the Control File” above
- `custom_thesaurus` – is the name of the custom thesaurus you are creating

For example, to execute the `mksyd` utility reading the sample control file defined above, and directing output to a work directory, use the syntax:

```
mksyd -f /usr/u/sybase/dba/thesaurus/colors.ctl  
-syd /usr/u/sybase/dba/thesaurus/custom.syd
```

---

### Replacing the Default Thesaurus with the Custom Thesaurus

The default thesaurus named `vdk20.syd` is located in:

`$$SYBASE/sds/text/verity/common/vdkLanguage`

where `vdkLanguage` is the value of the `vdkLanguage` configuration parameter (for example, the English directory is `$$SYBASE/sds/text/verity/common/english0`). Each application and user reading from this location at runtime uses this thesaurus. To replace it with your custom thesaurus, follow these steps:

1. Back up the default thesaurus before replacing it with the custom thesaurus. For example:

```
mv /sybase/sds/text/verity/common/english0/vdk20.syd default.syd
```

2. Replace the *vdk20.syd* file with your custom thesaurus. For example:

```
cp custom.syd /sybase/sds/text/verity/common/english0/vdk20.syd
```

3. Restart your Full-Text Search engine; no configuration file changes are required. The thesaurus is read from this location when the Full-Text Search engine is started, not when a query is executed.

Queries using the thesaurus operator will now use the custom thesaurus.

## Creating Topics (Enhanced Version Only)

---

A **topic** is a grouping of information related to a concept or subject area. With topic definitions in place, a user can perform searches on the topic instead of having to write queries with complex syntax.

The user creates topics which can be combinations of words and phrases, Verity operators and modifiers, and weight values. Then, any user can query the topic.

Before you create topics, determine your application requirements, and establish standards for naming conventions and for the location of the following:

- Outline files – contains the topic definitions. Each topic has its own outline file.
- Topic set directories – contains the compiled topic. Each topic has its own topic set directory.
- Knowledge base map file – contains pointers to the topic set directories.

To implement topics, perform the following steps:

1. Create one or more outline input files to define your topics (see “Creating an Outline File” on page 4-12)
2. Create and populate a topic set directory, using the **mktopics** utility (see “Creating a Topic Set Directory” on page 4-13)
3. Create a **knowledge base map**, specifying the locations of one or more topic set directories (see “Creating a Knowledge Base Map” on page 4-14)

4. Set the `knowledge_base` configuration parameter to point to the location of the knowledge base map (see “Defining the Location of the Knowledge Base Map” on page 4-14)
5. Execute queries against defined topics.

For more information about outline formats, operator precedence rules, and the `mktopics` utility, see the Verity Web site:

<http://www.verity.com>.

See also the Verity document *Search '97 Introduction to Topics*.

The following sample files illustrate the topics feature:

- `sample_text_topics.otl` is a sample outline file
- `sample_text_topics.kbm` is a sample knowledge base map
- `sample_text_topics.sql` issues queries using defined topics

These files are in the `$SYBASE/sds/text/sample/scripts` directory.

### Creating an Outline File

A topic outline file specifies all the combinations of words and phrases, Verity operators and modifiers, and weight values that the search engine uses when you issue a query using the topic. The outline file is an ASCII text file in a structured format.

For example, the following outline file defines the topic “saint-bernard”:

```
$control: 1
saint-bernard <accrue>
*0.80 "Saint Bernard"
*0.80 "St. Bernard"
* "working dogs"
* "large dogs"
* "European breeds"
$$
```

When you issue a query specifying the topic “saint-bernard”, the Full-Text Search engine:

- Returns documents that contain one or more of the following phrases: “Saint Bernard,” “St. Bernard,” “working dogs,” “large dogs,” and “European breeds”
- Scores documents that contain the phrase “Saint Bernard” or “St. Bernard” higher than documents that contain the phrase “working dogs,” “large dogs,” or “European breeds”

This example is a very basic topic definition. An outline can introduce more complex relationships by using:

- Multiple levels of subtopics
- Combinations of Verity operators (this example uses *accrue*)
- Verity modifiers

For complex examples of outline files, see the Verity Web site.

► **Note**

---

In Windows NT, you can use the graphical user interface of the Verity topicEDITOR product to create topic outlines. It is available from Verity. If you use topicEDITOR, it automatically creates a topic set directory, and you can go to “Creating a Knowledge Base Map” on page 4-14 to continue setting up your topics.

---

### Creating a Topic Set Directory

---

Use the `mktopics` utility to create and populate a topic set directory. It is located in:

`$$YBASE/sds/text/verity/bin`

Run, or define an alias to run, `mktopics` from this *bin* directory. You can create a topic set directory or directories in any work directory.

The `mktopics` syntax is:

```
mktopics -outline outline_file.otl -topicset topic_set_directory
```

where:

- *outline\_file* – is the name of the outline file you create in “Creating an Outline File” on page 4-12
- *topic\_set\_directory* – is the name of the topic set directory you are creating

For example, to execute the `mktopics` utility reading the *saint-bernard.otl* file defined above, and directing output to a work directory, use the syntax:

```
mktopics -outline /usr/u/sybase/topic_outlines/saint-bernard.otl
-topicset /usr/u/sybase/topic_sets/saint-bernard_topic
```

## Creating a Knowledge Base Map

---

A **knowledge base map** specifies the locations of one or more topic set directories. Create an ASCII knowledge base map file that defines the fully-qualified directory paths to your topic sets.

For example, the following knowledge base map file illustrates how you can list multiple knowledge bases in the map. The first entry identifies the topic set directory created with `mktopics` above.

```
$control:
1 kbases:
{
kb:
/kb-path = /usr/u/sybase/topic_sets/saint-bernard_topic
kb:
/kb-path = /usr/u/sybase/topic_sets/another_topic
}
```

## Defining the Location of the Knowledge Base Map

---

Set the `knowledge_base` configuration parameter to point to the location of the knowledge base map. For example:

```
sp_text_configure KRAZYKAT, 'knowledge_base',
'/usr/u/sybase/topic_sets/sample_text_topics.kbm'
```

The `knowledge_base` configuration parameter is static, and you must restart the Full-Text Search engine for the definition to take effect.

## Executing Queries Against Defined Topics

---

You can now execute queries using the defined topic instead of a complex query. For example, before you create the “saint-bernard” topic, you would have to use the following syntax:

```
...where i.index_any = "<accrue> ([80]Saint
Bernard, [80]St. Bernard, working dogs, large
dogs, European breeds)"
```

to find documents that:

- Contain one or more of the following phrases: “Saint Bernard,” “St. Bernard,” “working dogs,” “large dogs,” and “European breeds”
- Score documents containing the phrase “Saint Bernard” or “St. Bernard” higher than documents containing the phrase “working dogs,” “large dogs,” or “European breeds”

After you create the topic “saint-bernard”, you can use this syntax:

```
...where i.index_any = "<topic>saint-bernard"
```

or:

```
...where i.index_any = "saint bernard"
```

► **Note**

---

If you enter a word in a query expression, the Full-Text Search engine tries to match it with a topic name. If you enter a phrase in a query expression, the Full-Text Search engine replaces spaces with hyphens (-), and then tries to match it with a topic name. For example, the Full-Text Search engine matches “saint bernard” with the topic “saint-bernard”.

---

See the *sample\_text\_topics.sql* file for examples of using topics in queries.

### Troubleshooting Topics

---

If the `knowledge_base` configuration parameter specifies a knowledge base map file that does not exist, the Full-Text Search engine will not be able to start a session with Verity, and the server will not start. If the map file exists but contains invalid entries, Verity issues warning messages at start-up time.





# 5

## Writing Full-Text Search Queries

This chapter describes the pseudo columns, search operators, and modifiers that you can include in a full-text search. Topics include:

- Components of a Full-Text Search Query 5-1
- Pseudo Columns in the Index Table 5-2
- Full-Text Search Operators 5-8
- Operator Modifiers 5-19

### Components of a Full-Text Search Query

---

To write a full-text search query, you enter the search parameters as part of an Adaptive Server `select` statement, using the `isql` utility. Then the Full-Text Search engine processes the search. The `select` statement requires:

- A `where` clause that assigns a Verity language query to the `index_any` pseudo column
- Pseudo columns to further define the parameters of the search (optional)
- A join between the `IDENTITY` column from the source table and the `id` column from the index table

For example, to return the 10 documents from the `copy` column of the `blurbs` table that have the most occurrences of the word “software,” enter:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "<many> <word> software"
and t1.max_docs = 10
```

Adaptive Server passes the Verity query to the Full-Text Search engine to process the search. For more information on how Adaptive Server processes the query, see “How a Full-Text Search Works” on page 2-6.

## Pseudo Columns in the Index Table

**Pseudo columns** are columns in the index table that define the parameters of the search and provide access to the results data. (For more information about index tables, see “The Index Table” on page 2-3.) These columns are valid only in the context of a query; that is, the information in the columns is valid only for the duration of the query. If the query that follows contains a different set of parameters, the pseudo columns contain a different set of values.

Each pseudo column in an index table describes a different search attribute. For example, if you indicate the *score* column, the query displays only the result set that falls within the parameters you define. For example, the following query displays only the results that have a *score* value greater than 90:

```
index_table_name.score > 90
```

Other pseudo columns (like *highlight*) are used to retrieve data generated by Verity for a particular document. Table 5-1 describes the pseudo columns that are maintained by the Full-Text Search engine.

Table 5-1: Full-Text Search engine pseudo columns

Pseudo Column Name	Description	Datatype	Length (in Bytes)
<i>cluster_number</i>	Enhanced Full-Text Search engine only. Contains the cluster that the row is part of. Clusters are numbered starting with 1. You can use the <i>cluster_number</i> column only in the <i>select</i> clause of a query.	<i>int</i>	4
<i>cluster_keywords</i>	Enhanced Full-Text Search engine only. Contains the keywords that Verity uses to build the cluster. You can use <i>cluster_keywords</i> only in the <i>select</i> clause of a query.	<i>varchar</i>	255
<i>highlight</i>	Offsets within the document all words from the query. You can use <i>highlight</i> only in the <i>select</i> clause of a query.	<i>text</i>	16
<i>id</i>	Uniquely identifies a document within a collection. Used to join with the <i>IDENTITY</i> column of the source table. You can use <i>id</i> in the <i>select</i> clause or <i>where</i> clause of a query.	<i>numeric</i>	6
<i>index_any</i>	Provides a Verity language query to the Full-Text Search engine. You can use <i>index_any</i> only in a <i>where</i> clause.	<i>varchar</i>	255

Table 5-1: Full-Text Search engine pseudo columns (continued)

Pseudo Column Name	Description	Datatype	Length (in Bytes)
<i>max_docs</i>	Limits results to the first <i>n</i> documents, based on the default sort order. In a clustered result set, limits results to the first <i>n</i> documents in each cluster. You can use <i>max_docs</i> only in a <i>where</i> clause.	<i>int</i>	4
<i>score</i>	The normalized measure of correlation between search strings and indexed columns. The <i>score</i> associated with a specific document has meaning only in reference to the query used to retrieve the document. You can use <i>score</i> in a <i>select</i> clause or a <i>where</i> clause.	<i>int</i>	4
<i>sort_by</i>	Specifies the sort order in which to return the result set. <ul style="list-style-type: none"> <li>The Standard Full-Text Search engine allows a single sort specification in the <i>sort_by</i> column.</li> <li>The Enhanced Full-Text Search engine allows up to 16 sort specifications in the <i>sort_by</i> column.</li> </ul> You can use <i>sort_by</i> only in a <i>where</i> clause.	<i>varchar</i>	35
<i>summary</i>	Selects summarization data. You can use the <i>summary</i> column only in the <i>select</i> clause of a query.	<i>varchar</i>	255

The following sections describe the functionality of the pseudo columns.

### Using the *score* Column to Relevance-Rank Search Results

**Relevance ranking** is the ability of the Full-Text Search engine to assign the *score* parameter a value that indicates how well a document satisfies the query. The *score* calculation depends on the search operator used in the query (for more information, see “Using the Verity Operators” on page 5-10). The closer the document satisfies the query, the higher the *score* value is for that document.

For example, if you search for documents that contain the word “rain,” a document with 12 occurrences of “rain” receives a higher *score* value than a document with 6 occurrences of “rain.”

If you set *score* to a high value (such as 90) in your query, you limit the result set to documents that have a *score* value greater than that number.

**► Note**


---

Verity uses decimals for *score* values; Sybase uses whole numbers. For example, if Verity reports a score value of .85, Sybase reports the same value as 85.

---

For example, the following query searches for documents that contain the word “raconteur” or “Paris,” or both, and that have a *score* of 90 or greater:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 90
and t1.index_any = "<accrue>(raconteur, Paris)"
score copy
```

-----  
(0 rows affected)

The query does not find any documents that contain the word “raconteur” or “Paris” and that have a score greater than 90. However, if the *score* value in the query is lowered to 39, you find that one document in the *blurbs* table mentions the word “raconteur” or “Paris”:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 39
and t1.index_any = "<accrue>(raconteur, Paris)"
score copy
```

-----  
40 A chef's chef and a raconteur's raconteur, Reginald  
Blotchet-Halls calls London his second home. "Th' palace  
. . .

### Using the *sort\_by* Column to Specify a Sort Order

---

The sort order specifies the collating sequence used to order the data in the result set. The default sort order is set by the *sort\_order* configuration parameter (for more information, see “Setting the Default Sort Order” on page 6-9).

Use the *sort\_by* pseudo column to return a result set with a sort order other than the default. With the Standard Full-Text Search engine, you can specify a single sort specification in the *sort\_by* pseudo

column. With the Enhanced Full-Text Search engine, you can specify up to 16 sort specifications in the *sort\_by* pseudo column.

Table 5-2 lists the values for the *sort\_by* pseudo column.

Table 5-2: Values for the *sort\_by* pseudo column

Value	Description
<i>fts_score desc</i>	Returns a result set sorted by score in descending order.
<i>fts_score asc</i>	Returns a result set sorted by score in ascending order.
<i>fts_timestamp desc</i>	Returns a result set sorted by a timestamp in descending order.
<i>fts_timestamp asc</i>	Returns a result set sorted by a timestamp in ascending order.
<i>column_name desc</i>	Returns a result set sorted according to the descending order of a column. <i>column_name</i> is the name of the source table's column.
<i>column_name asc</i>	Returns a result set sorted according to the ascending order of a column. <i>column_name</i> is the name of the source table's column.
<i>fts_cluster asc</i>	Returns a clustered result set. Only available with the Enhanced Full-Text Search engine. (See "Using Pseudo Columns to Request Clustered Result Sets" on page 5-6 for more information.)

► **Note**

Before you can sort by specific columns, you must modify the *style.vgw* and *style.ufl* files (see "Setting Up a Column to Use As a Sort Specification" on page 4-4).

For example, the following query sorts the documents by timestamp in ascending order:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 90
and t1.index_any = "<accrue>(raconteur, Paris)"
and t1.sort_by = "fts_timestamp asc"
```

### Using the *summary* Column to Summarize Documents

Use the *summary* pseudo column to have queries return only summaries of the documents that meet the search criteria, rather than returning entire documents. The *summary* column is not available by default; you must edit the *style.prm* file prior to creating the text index to enable summarization. See “Enabling Query-By-Example, Summarization, and Clustering” on page 4-1 for information about enabling the *summary* column.

For example, the following query returns only summaries of documents that include the words “Iranian” and “book” (in this example, the *style.prm* file is configured to display 255 characters):

```
select t1.score, t1.summary
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 70
and t1.index_any = "(Iranian <and> book)"
```

```
score summary
```

```
-----
78  They asked me to write about myself and my book, so here
    goes: I started a restaurant called "de Gustibus" with two
    of my fri
```

```
(1 row affected)
```

The Full-Text Search engine supports summaries of up to 255 bytes.

For additional examples of queries using summarization, see the sample script *sample\_text\_queries.sql* in the *SSYBASE/sds/text/sample/scripts* directory.

### Using Pseudo Columns to Request Clustered Result Sets

The clustering function analyzes a result set and groups rows into clusters so that the rows in each cluster are semantically more similar to each other, on average, than they are to other rows in other clusters. Ordering rows by subtopics can help you get a sense of the major subject areas covered in the result set. Clustering is only available with the Enhanced Full-Text Search Specialty Data Store.

Returning a clustered result set can be significantly slower than returning a nonclustered result set. If the response time of a query is critical, use a nonclustered result set.

### Preparing to Use Clustering

---

Before you request a clustered result set, you must build the text index with the clustering function enabled (see “Enabling Query-By-Example, Summarization, and Clustering” on page 4-1).

The Verity clustering algorithm attempts to group similar rows together, based on the values of the following configuration parameters:

- `cluster_style`
- `cluster_max`
- `cluster_effort`
- `cluster_order`

Use the `sp_text_cluster` system procedure to have a query use values that are different from the default values of these configuration parameters. (For values and how to set them for a query, see “`sp_text_cluster`” on page A-18.)

### Writing Queries Requesting a Clustered Result Set

---

To obtain a clustered result set, specify “`fts_cluster asc`” as the sort specification in the `sort_by` pseudo column of the query. For example:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "<many> <word> software"
and t1.max_docs = 10
and t1.sort_by = "fts_cluster asc"
```

Include any of the following pseudo columns in your query to return additional clustering information:

- `cluster_number` – contains the number of the cluster the row belongs to. Clusters are numbered starting with 1.
- `cluster_keywords` – contains the most common words found in the cluster. The `cluster_keywords` column contains a null value for each row that does not fit into any cluster.
- `max_docs` – limits the number of rows returned for each cluster. (In a nonclustered query, the `max_docs` column limits the total number of rows that are returned in a result set.)
- `score` – contains a value of 0 to 10000. The higher the score, the closer the row is to the cluster center. A score of 0 indicates the

row does not fit into any cluster. (In a nonclustered query, the *score* column can contain a value of 0 to 100.)

See the sample script named *sample\_text\_queries.sql* in the *\$\$YBASE/sds/text/sample/scripts* directory for examples of SQL statements using clustering.

## Full-Text Search Operators

The special search operators that you use to perform full-text searches are part of the Verity Search '97 search engine. Table 5-3 describes the Verity search operators provided by the Full-Text Search engine.

Table 5-3: Verity search operators

Operator Name	Description
<b>accrue</b>	Selects documents that contain at least one of the search elements specified in a query. The more search elements there are, the higher the score will be.
<b>and</b>	Selects documents that contain all the search elements specified in a query.
<b>complement</b>	Returns the complement of the <i>score</i> value (the <i>score</i> value subtracted from 100).
<b>in</b>	Selects documents that contain the search criteria in the document zone specified.
<b>like</b>	Selects documents that are similar to the sample documents or passages specified in a query.
<b>near</b>	Selects documents containing the specified search elements, where the closer the search terms are to each other in a document, the higher the document's score.
<b>near/n</b>	Selects documents containing two or more search terms within <i>n</i> number of words of each other, where <i>n</i> is an integer up to 1000. The closer the search terms are to each other in a document, the higher the document's score.
<b>or</b>	Selects documents that contain at least one of the search elements specified in a query.
<b>paragraph</b>	Selects documents that include all the search elements you specify within the same paragraph.
<b>phrase</b>	Selects documents that include a particular phrase. A phrase is a grouping of two or more words that occur in a specific order.



Table 5-3: Verity search operators (continued)

Operator Name	Description
product	Multiplies the score values for each of the items of the search criteria.
sentence	Selects documents that include all the specified words in the same sentence.
stem	Expands the search to include the specified word and its variations.
sum	Adds the score values for all items in the search criteria.
thesaurus	Expands the search to include the specified word and its synonyms.
topic	Specifies that the search term you enter is a topic.
wildcard	Matches wildcard characters included in search strings. Certain characters indicate a wildcard specification automatically.
word	Performs a basic word search, selecting documents that include one or more instances of the specified word.
yesno	Converts all nonzero score values to 100.

### Considerations When Using Verity Operators

Consider the following when you write full-text search queries:

- You **must** enclose the operators in angle brackets (<>) in the query. If they are not enclosed in angle brackets, the Full-Text Search engine issues error messages similar to the following:

```
Msg 20200, Level 15, State 0:
Server 'KRAZYKAT', Line 1:
Error E1-0111 (Query Builder): Syntax error in query string near
character 5
Msg 20200, Level 15, State 0:
Server 'KRAZYKAT', Line 1:
Error E1-0114 (Query Builder): Error parsing query: word(tasmanian)
Msg 20101, Level 15, State 0:
Server 'KRAZYKAT', Line 1:
VdkSearchNew failed with vdk error (-40).
Msg 20101, Level 15, State 0:
Server 'KRAZYKAT', Line 1:
VdkSearchGetInfo failed with vdk error (-11).
score copy
-----
(0 rows affected) score
```

- You must enclose the Verity language query in single quotes (') or double quotes ("). The Full-Text Search engine strips off the outermost quotes before it sends the query to Verity. For example, when you enter the query:

```
...where index_any = "'?own'"
```

the Full-Text Search engine sends the following query to Verity:

```
'?own'
```

- Search terms entered in mixed case automatically become case sensitive. Search terms entered in all uppercase or all lowercase are not automatically case sensitive. For example, a query on "Server" finds only the string "Server". A query on "server" or "SERVER" finds the strings "Server", "server", and "SERVER".
- You can use alternative syntax for the query expressions shown in Table 5-4.

Table 5-4: Alternative Verity syntax

Standard Query Expression	Alternative Syntax
<MANY><WORD> <i>string</i>	" <i>string</i> "
<MANY><STEM> <i>string</i>	' <i>string</i> '

When using the alternative syntax, remember that the Full-Text Search engine strips off the outermost quotes before it sends the query to Verity. For example, when you enter the query:

```
...where index_any = "'play'"
```

the Full-Text Search engine sends the following query to Verity:

```
'play'
```

This is the same as:

```
<MANY><STEM>play
```

### Using the Verity Operators

The following sections describe how to use the Verity operators shown in Table 5-3 on page 5-8. For complete information on the syntax for Verity operators, see the Verity Web site at:

<http://www.verity.com>

### *accrue*

---

The *accrue* operator selects documents that contain at least one of the search items specified in the query. There must be two or more search elements. Each result is relevance-ranked. For example, the following query searches for the word “restaurant” or “deli” or both in the *copy* column of the *blurbs* table:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 35
and t1.index_any = "<accrue>(restaurant, deli)"
```

### *and, or*

---

The *and* and *or* operators select documents that contain the specified search elements. Each result is relevance-ranked. The *and* operator selects documents that contain all the elements specified in the query. For example, the following query selects documents that contain both “Iranian” and “business”:

```
select t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "(Iranian <and> business)"
```

The *or* operator selects the documents that contain any of the search elements. For example, if the preceding query is rewritten to use the *or* operator, the query selects documents that contain the word “Iranian” or “business”:

```
select t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "(Iranian <or> business)"
```

### *complement*

---

The *complement* operator returns the complement of the *score* value for a document; that is, it subtracts the value of *score* from 100 and returns the result as the *score* value for the document.

### *in*

---

The *in* operator selects documents that contain the specified search element in one or more document zones. Document zones are created for a text index in the following two scenarios:

- When you create an index on two or more columns using `sp_create_text_index`, a document zone is created for each column in the text index (for more information, refer to “Specifying Multiple Columns When Creating a Text Index” on page 3-8). A document zone is not created when you create a text index on a single column. For example, if you specify the `au_id` and `copy` columns of the `blurbs` table when you create the text index, you can issue the following query:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 35
and t1.index_any = "gorilla <in> copy"
```

This returns rows that contain the word “gorilla” in the `copy` column. However, if you specify only the `copy` column of the `blurbs` table when you create the text index, this query does not return any rows.

- When you create an index that uses a filter, a document zone is created for each tag in the document (for more information, see “Using Filters on Text That Contains Tags” on page 4-6). You can limit your search to a particular tag by specifying the tag name after the `in` operator. For example, to search for the word “automotive” in a “title” tag in an HTML document, specify:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 35
and t1.index_any = "automotive <in> title"
```

### *like*

The `like` operator selects documents that are similar to the document(s) or passages you provide. The search engine analyzes the text to find the most important terms to use. If you specify multiple samples, the search engine selects important terms that are common across the samples. Each result is relevance-ranked.

The `like` operator accepts a single operand, called the query-by-example (QBE) specification. The QBE specification can be either literal text or document IDs. The document IDs are from the `IDENTITY` column in the source table. For example, to select documents that are similar to the document in the `copy` column in the row with an `IDENTITY` of “2”, enter:

```

select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 35
and t1.index_any = '<like> ( "{2}" )'

```

Before using literal text in the QBE specification, you must uncomment the following line in the *style.prm* file:

```
$define DOC-FEATURES "TF"
```

For more information, see “Enabling Query-By-Example, Summarization, and Clustering” on page 4-1.

See the sample script named *sample\_text\_queries.sql* in the *SSYBASE/sds/text/sample/scripts* directory for examples of SQL statements using QBE.

### *near, near/n*

---

The *near* operator selects documents that contain the items specified in the query and are near each other (“near” being a relative term). The documents in which the search words appear closest to each other receive the highest relevance-ranking.

The *near/n* operator specifies how far apart the items can be (*n* has a maximum value of 1000). The following example selects documents in which the words “raconteur” and “home” appear within 10 words of each other:

```

select t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "<near/10>(raconteur, home)"

```

### *or*

---

See “and, or” on page 5-11.

### *phrase*

---

The *phrase* operator selects documents that contain a particular phrase (a group of two or more items that occur in a specific order). Each result is relevance-ranked. The following example selects the documents that contain the phrase “gorilla’s head”:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 50
and t1.index_any = "<phrase>(gorilla's head)"
```

### *paragraph*

---

The **paragraph** operator selects documents in which the specified search elements appear in the same paragraph. The closer the words are to each other in a paragraph, the higher the score the document receives in relevance-ranking. The following example searches for documents in which the words “text” and “search” occur within the same paragraph:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 50
and t1.index_any = "<many><paragraph>(text, search)"
```

### *product*

---

The **product** operator multiplies the *score* value for the documents for each of the search elements. To arrive at a document’s score, the Full-Text Search engine calculates a score for each search element and multiplies the *scores*. For example:

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 50
and t1.index_any = "<product>(cat, created)"
```

The *score* value for each search element is 78; however, because the *score* values for the items are multiplied, the document has a *score* value of 61 ( $.78 \times .78 = .61(100) = 61$ ).

### *sentence*

---

The **sentence** operator selects documents in which the specified search elements appear in the same sentence. The closer the words are to each other in a sentence, the higher the score the document receives in relevance-ranking. The following example searches for documents in which the words “tax” and “service” occur within the same sentence:

```

select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 50
and t1.index_any = "<many><sentence>(tax, service)"

```

### *stem*

---

The stem operator searches for documents containing the specified word and its variations. For example, if you specify the word “cook,” the Full-Text Search engine produces documents that contain “cooked,” “cooking,” “cooks,” and so on. To relevance-rank the result set, include the *many* modifier in the query (see “Operator Modifiers” on page 5-19).

The following query uses the stem operator to find documents that contain variations of the word “create,” that is, words that contain the word “create” as a stem. Notice that even though the first document contains a word in which “create” is not a perfect stem (“creative”), the document is still selected:

```

select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 10
and t1.index_any = "<many><stem>create"

```

```
score copy
```

```

-----
78  Anne Ringer ran away from the circus as a child.  A
    university creative writing professor and her family
    . . .
78  If Chastity Locksley didn't exist, this troubled world
    would have created her!  Not only did she master the mystic

```

### *sum*

---

The *sum* operator totals the *score* values for each search element, up to a maximum of 100. To arrive at a document’s score, the Full-Text Search engine calculates a score for each search element and totals those scores.

### *thesaurus*

---

The *thesaurus* operator searches for documents containing a synonym for a search element. For example, you might perform a search using the word “dog,” looking for documents that use any of its synonyms

("canine," "pooch," "pup," "watchdog," and so on). Each result is relevance-ranked.

The Full-Text Search engine supplies a default thesaurus. With the Enhanced Full-Text Search engine, you can create a custom thesaurus. For more information, see "Creating a Custom Thesaurus (Enhanced Version Only)" on page 4-7.

The following example uses the `thesaurus` operator to find a result set that contains synonyms for the word "crave." The first document is selected because it contains the word "want"; the second, because it contains the word "hunger":

```
select t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "<thesaurus>(crave)"
```

```
score copy
```

```
-----
78  They asked me to write about myself and my book, so here
    goes:  I started a restaurant called "de Gustibus" with two
    . . .
    of restaurant over another, when what they really want is a
    . . .
78  A chef's chef and a raconteur's raconteur, Reginald
    Blotchet-Halls calls London his second home. "Th' palace
    . . .
    his equal skill in satisfying our perpetual hunger for
    . . .
```

### ***topic*** (Enhanced Version Only)

The `topic` operator selects documents that meet the search criteria defined by the specified topic. For more information, see "Creating Topics (Enhanced Version Only)" on page 4-11. For example, use the following syntax to find documents that meet the criteria defined by the topic "engineering":

```
select t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = "<topic>(engineering)"
```



**wildcard**

The **wildcard** operator allows you to substitute wildcard characters for part of the item for which you are searching. Table 5-5 describes the wildcard characters and their attributes.

Table 5-5: Full-Text Search engine wildcard characters

Character	Function	Syntax	Locates
?	Specifies one alphanumeric character. You do not need to include the <b>wildcard</b> operator when you include the question mark in your query. The question mark is ignored in a set ([]) or in an alternative pattern ({}).	'?an'	"ran," "pan," "can," and "ban"
*	Specifies zero or more of any alphanumeric character. You do not need to include the <b>wildcard</b> operator when you include the asterisk in your query; you should not use the asterisk to specify the first character of a wildcard-character string. The asterisk is ignored in a set ([]) or in an alternative pattern ({}).	'corp*'	"corporate," "corporation," "corporal," and "corpulent"
[ ]	Specifies any single character in a set. If a word includes a set, you must enclose the word in backquotes ( ` ` ). Also, there can be no spaces in a set.	<wildcard> 'c[auo]t'	"cat," "cut," and "cot"
{ }	Specifies one of each pattern separated by a comma. If a word includes a pattern, you must enclose the word in backquotes ( ` ` ). Also, there can be no spaces in a set.	<wildcard> 'bank{s,er,ing}'	"banks," "banker," and "banking"
^	Specifies one of any character not included in a set. The caret (^) must be the first character after the left bracket ( [ ) that introduces a set.	<wildcard> 'st[^oa]ck'	Excludes "stock" and "stack," but locates "stick" and "stuck"
-	Specifies a range of characters in a set.	<wildcard> 'c[a-r]t'	Includes every three-letter word from "cat" to "crt"

To relevance-rank the result set, include the **many** modifier in the query (see "Operator Modifiers" on page 5-19).

For example, the following query searches for documents that include variations of the word "slingshot":

```
select t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id
and t1.index_any = '"slingshot*"'
```

```
score copy
```

```
-----
100  Albert Ringer was born in a trunk to circus parents, but
      another kind of circus trunk played a more important role
      . . .
      gorilla.  "Slingshotting" himself from the ring ropes,
      . . .
```

### *word*

---

The *word* operator searches for documents containing the specified word. To relevance-rank the result set, include the *many* operator in the query. The following example searches the *blurbs* table for documents containing the word "palates":

```
select t1.score, t2.copy
from i_blurbs t1, blurbs t2
where t1.id=t2.id and t1.score > 50
and t1.index_any = "<many><word>(palates)"
```

### *yesno*

---

The *yesno* operator converts all nonzero *score* values to 100. For example, if the score values for five documents are 86, 45, 89, 89, and 100, each of those documents is returned with a *score* value of 100. *score* values of 0 are not changed. The *yesno* operator is helpful for ensuring that all documents containing the search criteria are returned in the result set, regardless of the sort order.

## Operator Modifiers

The Verity query language includes modifiers that you can use with the operators to refine a search. The modifiers are described in Table 5-6.

Table 5-6: Verity operator modifiers

Modifier Name	Description	Works with These Operators	Example
<b>case</b>	Performs case-sensitive searches. If you enter search terms in mixed case, the search is automatically case sensitive.	wildcard word	<code>&lt;case&gt;&lt;word&gt;(Net)</code>
<b>many</b>	Counts the number of times that a word, stemmed word, or phrase occurs in a document. Relevance-ranks the document according to its density.	paragraph phrase sentence stem word wildcard	<code>&lt;many&gt;&lt;stem&gt;(write)</code>
<b>not</b>	Excludes documents that contain the items for which the query is searching.	and or	<code>cat&lt;and&gt;&lt;not&gt;elephant</code>
<b>order</b>	Specifies that the items in the documents occur in the same order in which they appear in the query.  Always place the <b>order</b> modifier just before the operator	near/n paragraph sentence	Simple syntax: <code>tidbits&lt;order&gt;&lt;paragraph&gt;king</code>  Explicit syntax: <code>&lt;order&gt;&lt;paragraph&gt;(tidbits,king)</code>



# 6

## System Administration

This chapter describes system administration issues for both the Standard and Enhanced versions of the Full-Text Search engine. Topics include:

- Starting the Full-Text Search Engine on UNIX 6-1
- Starting the Full-Text Search Engine on Windows NT 6-2
- Shutting Down the Full-Text Search Engine 6-4
- Modifying the Configuration Parameters 6-4
- Backup and Recovery for the Standard Full-Text Search Engine 6-13
- Backup and Recovery for the Enhanced Full-Text Search Engine 6-16

### Starting the Full-Text Search Engine on UNIX

---

Use the `startserver` utility to start the Full-Text Search engine on UNIX. The `startserver` utility is included in the `bin` directory of Adaptive Server. For example, to start a Full-Text Search engine named KRAZYKAT, enter:

```
startserver -f $SYBASE/install/RUN_KRAZYKAT
```

where the `-f` flag specifies the relative path to the `runserver` file. After you issue the command, the Full-Text Search engine issues a series of messages describing the settings of the configuration parameters.

### Creating the Runserver File

---

The `runserver` file contains start-up commands for the Full-Text Search engine. The `runserver` file can include the flags shown in Table 6-1.

Table 6-1: Definition of flags in the `runserver` file

Flag	Definition
<code>-Sserver_name</code>	Specifies the name of the Full-Text Search engine and is used to locate the configuration file and the network connection information in the <code>interfaces</code> file.

**Table 6-1: Definition of flags in the runserver file**

Flag	Definition
-t	Causes the Full-Text Search engine to write start-up messages to standard error.
- <i>errorlog_path</i>	Specifies the path to the error log file.
- <i>interfaces_file_path</i>	Specifies the path to the interfaces file.

A sample runserver file is copied to the `$$SYBASE/install` directory during installation. Make a copy of this file, renaming it `RUN_server_name`, where `server_name` is the name of the Full-Text Search engine. You must include the `LD_LIBRARY_PATH` environment variable in the runserver file. For example, the runserver file for a Full-Text Search engine named `KRAZYKAT` would be `RUN_KRAZYKAT` and would be similar to:

```
#!/bin/sh
#
SYBASE=$SYBASE/sds/text
export SYBASE

LD_LIBRARY_PATH="$SYBASE/lib:$LD_LIBRARY_PATH"
export LD_LIBRARY_PATH

$SYBASE/bin/txtsvr -SKRAZYKAT
```

The start-up command in the runserver file must consist of a single line and cannot include a return. If you have to carry the contents of the file over to a second or third line, include a backslash (\) for a line break.

## Starting the Full-Text Search Engine on Windows NT

You can start the Full-Text Search engine from Sybase Central™, as a service, or from the command line:

- From Sybase Central – see your Sybase Central documentation for information about starting servers.
- As a service – see “Starting the Full-Text Search Engine As a Service” on page 6-3.
- From the command line – use the following syntax:

```
%SYBASE%\sds\text\bin\txtsvr.exe -Sserver_name
[-t] [-i%SYBASE%path_to_sql.ini_file]
[-l%SYBASE%path_to_errorlog]
```

where:

- *-S* is the name of the Full-Text Search engine you are starting
- *-t* directs start-up messages to standard error
- *-i* is the path to the *sql.ini* file
- *-l* is the path to the error log

For example, to start a Full-Text Search engine named KRAZYKAT using the default *sql.ini* and error log files, and using *-t* to trace start-up messages, enter:

```
%SYBASE%\sds\text\bin\txtsvr.exe -SKRAZYKAT -t
```

The Full-Text Search engine is up and running when you see the start-up complete message.

### Starting the Full-Text Search Engine As a Service

Use the *instsvr* utility in Sybase Central to add the Full-Text Search engine to the list of items you can start and stop with the Services utility. *instsvr* is located in the *%SYBASE%\sds\text\bin* directory.

The *instsvr* utility uses the following syntax:

```
instsvr.exe service_name %SYBASE%\sds\text\bin\txtsvr.exe
"startup_parameters"
```

where:

- *service\_name* is the name of the Full-Text Search engine you are adding as a service. With Sybase Central, Sybase recommends you use a server name with the extension “\_TS” (for example, KRAZYKAT\_TS).
- *startup\_parameters* are any parameters you want used at start-up.

For example, to install a Full-Text Search engine named KRAZYKAT\_TS as a service, enter:

```
instsvr.exe KRAZYKAT_TS %SYBASE%\sds\text\bin\txtsvr.exe
"-SKRAZYKAT_TS -t"
```

#### ► Note

If you need to include more than one parameter (for example, *-i*), you must include all the parameters in one set of double quotes.

To configure Sybase Central to start and stop your Full-Text Search engine, you must provide a service name that begins with

“SYBTEXT\_*server\_name*”, where *server\_name* is the name of the Full-Text Search engine listed in the interfaces file. For example, if the name in the interfaces file is KRAZYKAT\_TS, run the following `instsvr` command to create a service that can be managed by Sybase Central:

```
instsvr SYBTEXT_KRAZYKAT_TS %SYBASE%\sds\text\bin\txtsvr.exe
"-SKRAZYKAT_TS -t"
```

## Shutting Down the Full-Text Search Engine

Use the following command to shut down the Full-Text Search engine from Adaptive Server:

```
server_name...sp_shutdown
```

where *server\_name* is the name of the Full-Text Search engine you are shutting down.

For example, to shutdown a Full-Text Search engine named KRAZYKAT, enter:

```
KRAZYKAT...sp_shutdown
```

## Modifying the Configuration Parameters

Each Full-Text Search engine has configuration parameters with default values, as shown in Table 6-2.

Table 6-2: Configuration parameters

Parameter	Description	Default Value
<code>batch_size</code>	Determines the size of the batches sent to the Full-Text Search engine.	500
<code>max_indexes</code>	The maximum number of text indexes that will be created in the Full-Text Search engine.	126
<code>max_stacksize</code>	Size (in kilobytes) of the stack allocated for client threads.	34,816
<code>max_threads</code>	Maximum number of threads available for the Full-Text Search engine.	50
<code>max_packetsize</code>	Packet size sent between the Full-Text Search engine and the Adaptive Server.	2048



Table 6-2: Configuration parameters (continued)

Parameter	Description	Default Value
<b>max_sessions</b>	Maximum number of sessions for the Full-Text Search engine.	100
<b>min_sessions</b>	Minimum number of sessions for the Full-Text Search engine.	10
<b>language</b>	Language used by the Full-Text Search engine.	us_english
<b>charset</b>	Character set used by the Full-Text Search engine.	iso_1
<b>vdkCharset</b>	Character set used by Verity Search '97.	850
<b>vdkLanguage</b>	Language used by Verity Search '97.	english0
<b>vdkHome</b>	Verity directory.	UNIX: <i>SSYBASE/sds/text/verity</i> Windows NT: <i>%SYBASE%\sds\text\verity</i>
<b>collDir</b>	Storage location of the Full-Text Search engine's collection.	UNIX: <i>SSYBASE/sds/text/collections</i> Windows NT: <i>%SYBASE%\sds\text\collections</i>
<b>default_Db</b>	Name of the Full-Text Search engine database that stores text index metadata.	<i>text_db</i>
<b>interfaces</b>	Full path to the directory in which the interfaces file used by the Full-Text Search engine is located.	UNIX: <i>SSYBASE/interfaces</i> Windows NT: <i>%SYBASE%\ini\sql.ini</i>
<b>sort_order</b>	Default sort order.	0
<b>errorLog</b>	Full path name to the error log file.	The directory in which you start Full-Text Search engine
<b>traceflags</b>	String containing numeric identifiers used to generate diagnostic information.	0
<b>srv_traceflags</b>	String containing numeric flag identifiers used to generate Open Server diagnostic information.	0

The Enhanced Full-Text Search engine has additional configuration parameters as shown in Table 6-3:

**Table 6-3: Configuration parameters for Enhanced version only**

Parameter	Description	Default Value
<code>cluster_style</code>	Clustering style to use.	Fixed
<code>cluster_max</code>	Maximum number of clusters to generate when <code>cluster_style</code> is set to Fixed.	0
<code>cluster_effort</code>	Amount of effort the Full-Text Search engine should expend on finding a good cluster.	Default
<code>cluster_order</code>	The order to return clusters and rows within a cluster.	0
<code>auto_online</code>	Specifies whether to bring indexes online automatically when the Full-Text Search engine is started. 0 indicates online is not automatic; 1 indicates automatic.	0
<code>backDir</code>	The default location for the placement of text index backup files.	UNIX: <code>SSYBASE/sds/text/backup</code> Windows NT: <code>%SYBASE%\sds\text\backup</code>
<code>knowledge_base</code>	The location of a knowledge base map for implementing the Verity topics feature.	null
<code>nocase</code>	Sets the case-sensitivity of the Full-Text Search engine. If you are using a case-sensitive sort order in Adaptive Server, set to 0. If you are using a case-insensitive sort order, set to 1.	0

A sample configuration file that includes all of these parameters is copied to your installation directory during installation. The sample configuration file is named `textsvr.cfg`. The entire sample configuration file is listed in Appendix B, "Sample Files."

### Modifying Values in the Standard Version

With Standard Full-Text Search Specialty Data Store, you use a configuration file to change the default values. The configuration file

is named *server\_name.cfg* and is in the *\$\$SYBASE* directory. *server\_name* is the name of the Full-Text Search engine.

- For UNIX, the *srvbuild* utility creates the configuration file when it builds the Full-Text Search engine.
- For Windows NT, you manually create the configuration file by copying a sample configuration file with default values.

To modify the default values, use a text editor to edit the configuration file. Uncomment the line that contains the configuration parameter you are modifying. You must restart the Full-Text Search engine for the new values to take effect.

### Modifying Values in the Enhanced Version

---

With Enhanced Full-Text Search Specialty Data Store, you can use the *sp\_text\_configure* system procedure to change the value of a configuration parameter. The syntax is:

```
sp_text_configure server_name, config_name,  
                config_value
```

where:

- *server\_name* is the name of the Full-Text Search engine
- *config\_name* is the name of the configuration parameter
- *config\_value* is the value you assign to the configuration parameter

For more information, see “*sp\_text\_configure*” on page A-21.

► **Note**

---

You can also modify the value of a configuration parameter by editing a configuration file as described in above.

---

### Setting the Default Language

---

The default language for Verity is set with the *vdkLanguage* configuration parameter. By default, *vdkLanguage* is set to “*english0*”.

You can configure Verity to use a different default language. Table 6-4 lists the locales supported by Sybase.

**Table 6-4: vdkLanguage configuration parameters**

Language	Default Locale Name
English	english0
German	german0
French	french0

Additional language adapters are available in the `$$SYBASE/sds/text/verity/common` directory; however, the Full-Text Search engine displays messages only in the languages shown in Table 6-4.

The `language` parameter is the language the Full-Text Search engine displays its error messages and Open Server and Open Client error messages. Set the `language` parameter to the Adaptive Server language.

For example, with the Standard Full-Text Search engine, to change the Verity language to Spanish in a server named KRAZYKAT, include the following line in the configuration file:

```
vdkLanguage = spanish0
```

With the Enhanced Full-Text Search engine, run the following:

```
sp_text_configure KRAZYKAT, 'vdkLanguage', 'spanish0'
```

For more information about the Verity languages, see the Verity Web site:

<http://www.verity.com>

### Setting the Default Character Set

The default character set for Verity is set with the `vdkCharset` parameter in the configuration file. The files used for the Verity

character sets are in `$$SYBASE/sds/text/verity/common`. Table 6-5 describes the character sets you can use with Verity.

**Table 6-5: Verity character sets**

Character Set	Description
850	Default
437	IBM PC character set
1252	Windows code page for Western European languages
mac1	Macintosh roman

The default character set for the Full-Text Search engine is set with the `charset` parameter. Set the `charset` parameter to the Adaptive Server character set.

For example, with the Standard Full-Text Search engine, to change the Verity character set to IBM PC in a server named KRAZYKAT, include the following line in the configuration file:

```
vdkCharset = 437
```

With the Enhanced Full-Text Search engine, run the following:

```
sp_text_configure KRAZYKAT, 'vdkCharset', '437'
```

For more information about the Verity character sets, see the Verity Web site:

<http://www.verity.com>

### Setting the Default Sort Order

By default, the Full-Text Search engine sorts the result set by the `score` pseudo column in descending order (the higher scores appear first). To change the default sort order, set the `sort_order` configuration parameter to one of the values in Table 6-6.

**Table 6-6: Sort order values for the configuration file**

Value	Description
0	Returns result sets sorted by the <code>score</code> pseudo column in descending order. The default value.
1	Returns result sets sorted by the <code>score</code> pseudo column in ascending order.

**Table 6-6: Sort order values for the configuration file (continued)**

Value	Description
2	Returns result sets sorted by a timestamp in descending order.
3	Returns result sets sorted by a timestamp in ascending order.

For example, with the Standard Full-Text Search engine, to change the default sort order to sort by descending timestamp in a server named KRAZYKAT, include the following line in the configuration file:

```
sort_order = 2
```

With the Enhanced Full-Text Search engine, enter:

```
sp_text_configure KRAZYKAT, 'sort_order', '2'
```

When you sort a result set by descending timestamp (value 2 in Table 6-6), the Full-Text Search engine returns the newest documents first. The newest documents are those that were inserted or updated most recently. When results are sorted by ascending timestamp (value 3 in Table 6-6), the Full-Text Search engine returns the oldest documents first.

Setting the default sort order is especially important if your query uses the *max\_docs* pseudo column. The *max\_docs* pseudo column limits the number of rows of the result set to the first *n* rows, ordered by the sort order. If you set *max\_docs* to a number smaller than the size of the result set, the sort order you select could exclude the rows that contain the information for which you are searching.

For example, if you sort by ascending timestamp, the latest document added to the table appears last in the result set. If the entire result set consists of 11 documents, and you set *max\_docs* to 10, the latest document does not appear in the result set. However, if you sort by descending timestamp, the latest document appears first in the result set.

## Setting Trace Flags

The *traceflags* parameter enable the logging of certain events when they occur within the Full-Text Search engine. Each trace flag is

uniquely identified by a number. Trace flags are described in Table 6-7.

**Table 6-7: Full-Text Search engine trace flags**

Trace Flag	Description
1	Traces connects, disconnects, and attention events from Adaptive Server.
2	Traces language events. Traces the SQL statement that Adaptive Server sent to the Full-Text Search engine.
3	Traces RPC events.
4	Traces cursor events. Traces the SQL statement sent to the Full-Text Search engine by Adaptive Server.
5	Writes the errors that display to the log.
6	Traces information about text indexes. Writes the search string being passed to Verity to the log, and writes the number of records that the search returns to the log.
7	Traces done packets.
8	Traces calls to the interface between the Full-Text Search engine and the Verity API.
9	Traces SQL parsing.
10	Traces Verity processing.
11	Disables Verity collection optimization.
12	Disables <code>sp_statistics</code> from returning information.
13	Traces backup operations. Available only with Enhanced Full-Text Search Specialty Data Store.
14	Logs Verity status and timing information.
15	Generates ngram index information for collections. ngrams increase the speed of wildcard searches. This trace flag is required for wildcard searches against data in unicode format.

You can enable and disable trace flags interactively, using the remote procedure calls (RPCs) `sp_traceon` and `sp_traceoff` in the Full-Text Search engine. For more information on these RPCs, see the *Adaptive Server Reference Manual*.

## Setting Open Server Trace Flags

---

Use the `srv_traceflags` parameter to turn on trace flags to log Open Server diagnostic information. Open Server trace flags are described in Table 6-8.

Table 6-8: Open Server trace flags

Trace Flag	Description
1	Traces TDS headers.
2	Traces TDS data.
3	Traces attention events.
4	Traces message queues.
5	Traces TDS tokens.
6	Traces Open Server events.
7	Traces deferred event queues.
8	Traces network requests.

For example, with the Standard Full-Text Search engine, to trace attention events on the server named KRAZYKAT, include the following line in the configuration file:

```
srv_traceflags = 3
```

With the Enhanced Full-Text Search engine, run the following:

```
sp_text_configure KRAZYKAT, 'srv_traceflags', '3'
```

## Setting Case Sensitivity

---

By default, the Full-Text Search engine is case sensitive. This means you must enter identifiers in the same case or they are not recognized. For example, if you have a table named *blurbs* (lowercase), you cannot issue an `sp_create_text_index` command that specifies the table name *BLURBS*. You must issue a command that uses the same case for the table name argument:

```
sp_create_text_index "KRAZYKAT", "i_blurbs", "blurbs", "", "copy"
```

With Enhanced Full-Text Search engine, use the `nocase` parameter to set the case sensitivity of the Full-Text Search engine. 0 indicates case



sensitive; 1 indicates case insensitive. Set the `nocase` parameter to the sort order case sensitivity in Adaptive Server.

For example:

```
sp_text_configure KRAZYKAT, 'nocase', '1'
```

changes the KRAZYKAT server to case insensitive.

► **Note**

---

The `nocase` parameter does not affect the case sensitivity of the Verity query. For information on Verity case sensitivity, see “Considerations When Using Verity Operators” on page 5-9.

---

## Backup and Recovery for the Standard Full-Text Search Engine

---

The Adaptive Server user database and the Verity collections are physically separate. Backing up your user database does **not** back up the Verity collections, and restoring your database from a backup does **not** restore your Verity collections. The backup and recovery procedures described in Chapter 21, “Backing Up and Restoring User Databases,” of the *System Administration Guide* apply only to the user database and `text_db` database in Adaptive Server.

Make sure you follow the recommended schedule for backing up your databases that is described in Chapter 20, “Developing a Backup and Recovery Plan,” of the *System Administration Guide*. Sybase recommends that when you back up a user database with text indexes, you also back up:

- The `text_db` database
- The text indexes

A regular backup schedule ensures the integrity of the text indexes, the Adaptive Server data, and the `text_events` table, all of which are integral to recovering your text indexes without having to drop and re-create them.

---

**► Note**

You do not have to back up the user database and text indexes at the same time to recover the text indexes. However, you must restore the user database before you restore the text index. Doing so restores the *text\_events* table, which the *sp\_redo\_text\_events* system procedure uses to bring the text indexes in sync with the user database.

---

If you have Enhanced Full-Text Search Specialty Data Store, use the automated process described in “Backup and Recovery for the Enhanced Full-Text Search Engine” on page 6-16.

---

**Backing Up Verity Collections**

---

Follow these steps to back up your Verity collections:

1. Shut down the Full-Text Search engine:

```
server_name...sp_shutdown
```

2. Back up the files. By default, the collections are located in:

```
$SYBASE/sds/text/collections
```

Each collection name consists of the database name, owner name, and index name in the format *db.owner.index*. For example, if you create a text index called *i\_blurbs* on the *pubs2* database, the full path to those files would be similar to:

```
$SYBASE/sds/text/collections/pubs2.dbo.i_blurbs
```

- In UNIX, back up the files by using the *tar* or *cpio* utility
  - In Windows NT, use a compression utility such as *PKZIP* to back up the files
3. For future reference, make a note of the time of the backup in a permanent location.
  4. Back up the user database and the *text\_db* database, using the *dump database* command. For more information on the *dump database* command, see the *Adaptive Server Reference Manual*.
  5. Restart the Full-Text Search engine. For instructions, see “Starting the Full-Text Search Engine on UNIX” on page 6-1 or “Starting the Full-Text Search Engine on Windows NT” on page 6-2.

## Restoring Verity Collections and Text Indexes from Backup

As Database Administrator, follow these steps to restore your Verity collections:

1. Restore the Adaptive Server user database and *text\_db* database. This returns the source tables, metadata, and *text\_events* table to a consistent and predictable state. See Chapter 21, "Backing Up and Restoring User Databases," in the *System Administration Guide* for more information.
2. Shut down the Full-Text Search engine:  

```
server_name...sp_shutdown
```
3. Restore your collections from the backup files created in step 2 in "Backing Up Verity Collections" on page 6-14.
4. Restart the Full-Text Search engine. For instructions, see "Starting the Full-Text Search Engine on UNIX" on page 6-1 or "Starting the Full-Text Search Engine on Windows NT" on page 6-2.
5. Log in to Adaptive Server, and run the *sp\_redo\_text\_events* system procedure in the restored database. For example, if you are restoring the *pubs2* database, you have to be in that database to run the system procedure, *sp\_redo\_text\_events*, as follows:  

```
sp_redo_text_events "from_date"
```

where *from\_date* is the date and time associated with the backup used to recover the collections.  
For example:  

```
sp_redo_text_events "10/31/97:16:45"
```

restores the collections up to October 31, 1997 at 4:45 PM. For more information, see "sp\_redo\_text\_events" on page A-12.
6. Run the *sp\_text\_notify* system procedure to notify the Full-Text Search engine that changes need to be propagated to the Verity collections. The Full-Text Search engine connects to Adaptive Server, reads all the unprocessed entries in the *text\_events* table and applies them to the text index. For more information, see "sp\_text\_notify" on page A-30.

Your text indexes and collections are now fully restored.

---

## Backup and Recovery for the Enhanced Full-Text Search Engine

---

Backup and recovery for the Enhanced Full-Text Search Specialty Data Store is automated with the `sp_text_dump_database` and `sp_text_load_index` system procedures. These system procedures provide a seamless interface for maintaining data and text index integrity.

The Adaptive Server user database and the Verity collections are physically separate. Backing up your user database does **not** back up the Verity collections, and restoring your database from a backup does **not** restore your Verity collections. The backup and recovery procedures described in Chapter 21, “Backing Up and Restoring User Databases,” of the *System Administration Guide* apply only to the user database and the `text_db` database in Adaptive Server.

Follow the recommended schedule for backing up your databases, as described in Chapter 20, “Developing a Backup and Recovery Plan,” of the *System Administration Guide*. Sybase recommends that when you back up a user database with text indexes, you also back up:

- The `text_db` database
- The text indexes

► **Note**

---

You do not have to back up the user database and text indexes at the same time to recover the text indexes. However, you must restore the user database before you restore the text index. This restores the `text_events` table, which the `sp_text_load_index` system procedure uses to bring the text indexes in sync with the user database.

---

A regular backup schedule ensures the integrity of the text indexes, the Adaptive Server data, and the `text_events` table, all of which are integral to recovering your text indexes without having to drop and re-create them.

If you have the Standard Full-Text Search Specialty Data Store, use the automated process described in “Backup and Recovery for the Standard Full-Text Search Engine” on page 6-13.

---

### Backing Up Verity Collections

---

The `sp_text_dump_database` system procedure backs up collections and (optionally) the user and `text_db` databases. `sp_text_dump_database` also

maintains the *text\_events* table by deleting entries that are no longer needed for recovery. It is available only with the Enhanced Full-Text Search engine.

During a backup, the Full-Text Search engine processes queries, but defers any update requests until the backup is complete. This eliminates the need to shut down and restart the Full-Text Search engine.

Run `sp_text_dump_database` from the database containing the text indexes you are backing up. The backup of the text indexes is placed in the directory specified in the `backDir` configuration parameter. The output of the `dump database` command is written to the Full-Text Search error log. Sybase recommends dumping the current database and the *text\_db* database at the time the text indexes are backed up. However, this is optional.

For example, to back up the text indexes, the *sample\_colors\_db* database to the `/work2/sybase/colorsbackup` directory, and the *text\_db* database to the `/work2/sybase/textdbbackup` directory, enter:

```
sp_text_dump_database @backupdbs =  
INDEXES_AND_DATABASES, @current_to = "to  
'/work2/sybase/colorsbackup'", @textdb_to="to  
'/work2/sybase/textdbbackup' "
```

► **Note**

---

It is important to back up the *text\_db* database whenever text indexes are backed up, since that database contains the metadata for all text indexes.

---

For more information, see “`sp_text_dump_database`” on page A-23.

## Restoring Collections and Text Indexes from Backup

---

The `sp_text_load_index` system procedure restores text indexes that have been backed up with the `sp_text_dump_database` system procedure.

As Database Administrator, perform the following procedures to restore your Verity collections:

1. Restore your Adaptive Server user database and *text\_db* database. This returns the source tables, metadata, and *text\_events* table to a consistent and predictable state. Follow the procedures described in Chapter 21, “Backing Up and Restoring

User Databases,” in the *System Administration Guide*, to restore user and *text\_db* databases.

2. Run `sp_text_load_index` to restore the Verity collection from the most recent index dump. The procedure resets the status of all *text\_events* table entries made since the last index dump to “unprocessed” and notifies the Full-Text Search engine to process those events.

**Example:**

To restore the *sample\_colors\_db* database and all of its text indexes:

1. Restore the *text\_db* database:

```
1> use master
2> go

1> load database text_db from '/work2/sybase/textdbbackup'
2> go
```

2. Restore the *sample\_colors\_db* database:

```
1> load database sample_colors_db from
'/work2/sybase/colorsbackup'
2> go
```

3. Bring the *text\_db* and *sample\_colors\_db* databases online:

```
1> online database text_db
2> online database sample_colors_db
3> go
```

4. Restore the text index:

```
1> use sample_colors_db
2> go

1> sp_text_load_index
2> go
```

For more information, see “`sp_text_load_index`” on page A-28.

# 7

## Performance and Tuning

The Full-Text Search engine is shipped with a default configuration. You can optimize the performance of the Full-Text Search engine by altering the default configuration so that it better reflects the needs of your site. This chapter describes ways in which you can enhance performance. Topics include:

- Updating Existing Indexes 7-1
- Increasing Query Performance 7-2
- Reconfiguring Adaptive Server 7-3
- Reconfiguring the Full-Text Search Engine 7-4
- Using `sp_text_notify` 7-5
- Configuring Multiple Full-Text Search Engines 7-5

### Updating Existing Indexes

---

The amount of time it takes to update records in a text index can be reduced by enabling (turning on) trace flag 11 or trace flag 12, or both:

- Enabling trace flag 11 disables Verity collection optimization. This means that Verity does not optimize the text index after you issue `sp_text_notify`, which is a performance gain. If trace flag 11 is turned off (the default), the Full-Text Search engine calls Verity to optimize the text index at the end of `sp_text_notify` processing, which can delay the completion of `sp_text_notify`.  
With Enhanced Full-Text Search Specialty Data Store, you can use the `sp_optimize_text_index` system procedure to optimize a text index at a later time if trace flag 11 is enabled. (For more information, see “`sp_optimize_text_index`” on page A-10.)
- Enabling trace flag 12 disables the Full-Text Search engine from returning `sp_statistics` information. If trace flag 12 is turned off (the default), an `update statistics` command is issued to the Full-Text Search engine, which can delay the completion of `sp_text_notify`.

If updates to the text index occur as often as every few seconds, you may improve performance by disabling the `update statistics` processing and the Verity optimization, or both, for most of the updates.

Trace flags 11 and 12 can be enabled and disabled interactively using the remote procedure calls `sp_traceon` and `sp_traceoff` in the Full-Text Search engine. (See the *Adaptive Server Reference Manual* for information on `sp_traceon` and `sp_traceoff`.)

## Increasing Query Performance

---

Two issues can significantly improve query performance:

- Limiting the number of rows returned by the Full-Text Search engine
- Ensuring the correct join order for queries

### Limiting the Number of Rows

---

Use the `max_docs` pseudo column to limit the number of rows returned by the Full-Text Search engine. The fewer the number of rows returned by the Full-Text Search engine, the faster Adaptive Server can process the join between the source table and the index table.

### Ensuring the Correct Join Order for Queries

---

The more tables and text indexes that are listed in a join, the greater the chance that the query will run slowly because of incorrect join order. Queries run fastest when the text index is queried first during a join between the text index and one or more tables.

To ensure correct join order:

- Make sure that a unique clustered or nonclustered index is created on the `IDENTITY` column of the table being indexed
- Limit joins to one base table and one text index

If a query is running slowly, use `showplan` or enable trace flag 11205, and examine the join order. Trace flag 11205 dumps remote queries to the Adaptive Server error log file. The fastest queries contain an `index_any` search condition in the `where` clause and query the text index first.

The slowest queries contain the `id` column in the text index `where` clause and query the indexed table first. In this case, rewrite the query or use `forceplan` to force the join order that is listed in your query. For more information about `forceplan`, see Chapter 10,



“Advanced Optimizing Techniques,” in the *Performance and Tuning Guide*.

## Reconfiguring Adaptive Server

---

You can improve the performance of the Full-Text Search engine by resetting the following Adaptive Server configuration parameters. (For information about setting configuration parameters with `sp_configure`, see Chapter 11, “Setting Configuration Parameters,” in the *System Administration Guide*.)

### *cis cursor rows*

---

The `cis cursor rows` parameter specifies the number of rows received by Adaptive Server during a single fetch operation. The default number for `cis cursor rows` is 50. Increasing this number increases the number of rows received by Adaptive Server from the Full-Text Search engine during a fetch operation. However, keep in mind that the larger the number you set for `cis cursor rows`, the more memory Adaptive Server allocates to that parameter.

### *cis packet size*

---

The `cis packet size` parameter determines the number of bytes contained in a single network packet. The default for `cis packet size` is 512. You must specify values for this parameter in multiples of 512. Increasing this parameter improves the performance of the Full-Text Search engine because, with a larger packet size, it returns fewer packets for each query. However, keep in mind that the larger the number you set for `cis packet size`, the more memory Adaptive Server allocates to that parameter.

The `cis packet size` parameter is dynamic; you do not need to reboot Adaptive Server for this parameter to take effect.

► **Note**

---

If you change the `cis packet size`, you must also change the `max_packet_size` parameter in the Full-Text Search engine configuration file to the same value.

---

You need to reboot the Full-Text Search engine for the `max_packetsize` parameter to take effect.

## Reconfiguring the Full-Text Search Engine

---

You can improve the performance of the Full-Text Search engine by reconfiguring the following Full-Text Search engine configuration parameters (see “Modifying the Configuration Parameters” on page 6-4):

### *batch\_size*

---

The `batch_size` configuration parameter determines the number of rows per batch the Full-Text Search engine indexes. `batch_size` has a default of 500 (that is, 500 rows of data indexed per batch). Performance improves if you increase the size of the batches that are indexed. However, the larger the batch size, the more memory the Full-Text Search engine allocates to this parameter.

When considering how large to set `batch_size`, consider the size of the data on which you are creating a text index. When creating the text index, the Full-Text Search engine allocates memory equal to (in bytes):

*(amount of space needed for data) x (batch\_size) = memory used*

For example, if the data you are indexing is 10,000 bytes per row, and `batch_size` is set to 500, then the Full-Text Search engine will need to allocate almost 5MB of memory when creating the text index.

Base the batch size you choose on the typical size of your data and the amount of memory available on your machine.

### *min\_sessions and max\_sessions*

---

`min_sessions` and `max_sessions` determine the minimum and maximum number of user connections allowed for the Full-Text Search engine. Each user connection requires about 5MB of memory. Do not set `max_sessions` to an amount that exceeds your available memory. Also, because the memory for `min_sessions` is allocated at start-up, if you set the number for `min_sessions` extremely high (to allow for a large number of user connections), a large percentage of your memory will be dedicated to user connections for the Full-Text Search engine.

You may improve the performance of the Full-Text Search engine by setting `min_sessions` equal to the average number of user sessions that will be used. Doing so prevents the Full-Text Search engine from having to allocate memory at the start of the user session.

### Using `sp_text_notify`

---

Review the needs of your site before you decide how often to issue `sp_text_notify`.

Using the `sp_text_notify` system procedure produces a load on the Full-Text Search engine as the system procedure reads the data and updates the text collections. Depending on the size of this load, the performance hit for issuing `sp_text_notify` can be substantial. Because of the performance implications, you must determine how up to date the indexes need to be. If they need to be current (close to real-time), then you will have to issue `sp_text_notify` frequently (as often as every 5 seconds). However, if your indexes do not need to be that current, it may be prudent to wait until the system is not active before you issue `sp_text_notify`.

► *Note*

---

You cannot issue `sp_text_notify` from within a transaction.

---

### Configuring Multiple Full-Text Search Engines

---

For tables that are used frequently, you can improve performance by placing the text indexes for these tables on separate Full-Text Search engines. Performance improves because users can spread their queries over a number of Full-Text Search engines, instead of sending all queries to a single engine. Each Adaptive Server can connect to multiple Full-Text Search engines, but each Full-Text Search engine can connect to only one Adaptive Server.

#### Creating Multiple Full-Text Search Engines at Start-Up

---

If you are initially creating multiple Full-Text Search engines, you can edit the `installtextserver` script so that it includes all of those Full-Text Search engines. For more information, see “Editing the `installtextserver` Script” on page 3-2.

## Adding Full-Text Search Engines

---

You can add Full-Text Search engines at a later date by issuing the `sp_addserver` command from `isql`. The `sp_addserver` command has the following syntax:

```
sp_addserver server_name [, server_class [, physical_name]]
```

where:

- `server_name` is the name used to address the server on your system (in this case, the Full-Text Search engine).
- `server_class` identifies the category of server being added. For the Full-Text Search engine, the value is “sds”.
- `physical_name` is the name in the interfaces file used by the server `server_name`.

For more information, see `sp_addserver` in the *Adaptive Server Reference Manual*.

Follow the steps described in “Configuring the Full-Text Search Engine” in the *Installation and Release Bulletin* for your platform, to configure additional Full-Text Search engines. Each Full-Text Search engine requires its own:

- Interfaces file entry
- Configuration file

All Full-Text Search engines use the same database (named `text_db` by default) for storing text index metadata and the same `vesaux` and `vesauxcol` tables.

For example, to add a Full-Text Search engine named BLUE, enter:

```
sp_addserver BLUE, sds, BLUE
```

After you configure and start the Full-Text Search engine, you can use the following syntax to see if you can connect to the Full-Text Search engine:

```
server_name...sp_show
```

For example, to connect to a server named BLUE, enter:

```
BLUE...sp_show
```

# A

## System Procedures

This appendix describes the Sybase-supplied system procedures used for updating and getting reports from system tables. Table A-1 lists the system procedures included with the Full-Text Search engine.

Table A-1: System procedures

Procedure	Description
<code>sp_clean_text_events</code>	Removes entries from the <i>text_events</i> table.
<code>sp_clean_text_indexes</code>	Cleans up stray indexes.
<code>sp_create_text_index</code>	Creates an external text index.
<code>sp_drop_text_index</code>	Drops text indexes.
<code>sp_clean_text_events</code>	Removes processed entries from the <i>text_events</i> table.
<code>sp_help_text_index</code>	Enhanced version only. Displays text indexes.
<code>sp_optimize_text_index</code>	Enhanced version only. Runs the Verity optimization routines.
<code>sp_redo_text_events</code>	Changes the status of entries in the <i>text_events</i> table and forces re-indexing of the modified table.
<code>sp_refresh_text_index</code>	Adds an entry to the <i>text_events</i> table for the update to a source table.
<code>sp_show_text_online</code>	Displays information about databases or indexes that are currently online.
<code>sp_text_cluster</code>	Enhanced version only. Displays or modifies clustering options.
<code>sp_text_configure</code>	Enhanced version only. Displays or modifies Full-Text Search engine configuration parameters.
<code>sp_text_dump_database</code>	Enhanced version only. Makes a backup copy of the text indexes in a database and optionally dumps the <i>text_db</i> and current databases.
<code>sp_text_kill</code>	Enhanced version only. Terminates all connections to a text index.
<code>sp_text_load_index</code>	Enhanced version only. Restores text indexes from a backup.
<code>sp_text_notify</code>	Notifies the Full-Text Search engine that the <i>text_events</i> table has been modified.
<code>sp_text_online</code>	Makes a database available to Adaptive Server.

## sp\_clean\_text\_events

### Function

Removes processed entries from the *text\_events* table.

### Syntax

```
sp_clean_text_events [up_to_date]
```

### Parameters

*up\_to\_date* – the date and time through which all processed entries will be deleted.

### Examples

```
1. sp_clean_text_events "01/15/98:17:00"
```

Removes data entered on or before January 15, 1998 at 5:00 p.m.

### Comments

- If the *up\_to\_date* parameter is specified, all entries having a date less than or equal to *up\_to\_date* and whose status is set to processed is deleted.
- If *up\_to\_date* is omitted, all entries whose status is set to processed is deleted.
- Remove entries from the *text\_events* table only after you have backed up the collection associated with the text index.
- With the Enhanced Full-Text Search engine, the *sp\_text\_dump\_database* system procedure automatically runs this.

### Messages

None

### Permissions

Any user can execute *sp\_clean\_text\_events*.

### See Also

*sp\_text\_dump\_database*

## sp\_clean\_text\_indexes

### Function

Removes indexes from the *vesaux* table that are not associated with a table.

### Syntax

```
sp_clean_text_indexes
```

### Parameters

None.

### Examples

1. `sp_clean_text_indexes`

### Comments

- This procedure reads entries from the *vesaux* and *vesauxcol* tables, verifying that both the source table and the corresponding index table exist. If either is missing, the index is dropped.

### Messages

- Fetch resulted in an error
- Unable to drop object definition for *index\_name!*

### Permissions

Any user can execute `sp_clean_text_indexes`.

## sp\_create\_text\_index

### Function

Creates an external text index.

### Syntax

```
sp_create_text_index server_name, index_table_name,  
                    table_name, option_string, column_name  
                    [, column_name ... ]
```

### Parameters

*server\_name* – is the name of the Full-Text Search engine.

*index\_table\_name* – is the name of the index table. *index\_table\_name* has the form [*dbname*.*owner*.]*table*, where:

- *dbname* is the name of the database containing the index table.
- *owner* is the name of the owner of the index table.
- *table* is the name of the index table.

*table\_name* – is the name of the source table containing the text being indexed. *table\_name* has the form [*dbname*.*owner*.]*table*.

*option\_string* – is a placeholder for index creation options.

*column\_name* – is the name of the column indexed by the text index.

### Examples

```
1. sp_create_text_index "blue", "i_blurbs", "blurbs",  
   " ", "copy"
```

Creates a text index and an index table named *i\_blurbs* on the *copy* column of the *blurbs* table.

### Comments

- Up to 16 columns can be indexed in a single text index.
- Columns of the following datatypes can be indexed:
  - Standard version: *char*, *varchar*, *nchar*, *nvarchar*, *text*, *image*, *datetime*, and *smalldatetime*
  - Enhanced version: all datatypes in the Standard version, plus *int*, *smallint*, and *tinyint*



- The content of *option\_string* is not case sensitive.
- *option\_string* uses a null string (" ") to specify "No Options".
- Assign the value "empty" to *option\_string* to create a text index that you will immediately drop. This creates the Verity collection directory and the style files, but does not populate the collections. For example, when you configure an individual table for clustering, you create the text index and immediately drop it. After you edit the *style.prm* file, you re-create the text index. For more information, see "Editing Individual style.prm Files" on page 4-3.
- `sp_create_text_index` writes entries to the *vesaux* table and tells the Full-Text Search engine to create the text index.
- Execution of `sp_create_text_index` is synchronous. The Adaptive Server process executing this system procedure remains blocked until the index is created. The time required to index large amounts of data may be take as long as several hours to complete.
- When you create a text index on two or more columns, each column in the text index is placed into its own document zone. The name of the zone is the name of the column. The zones can be used to limit your search to a particular column. For more information, see "in" on page 5-11.

### Messages

- Can't run `sp_create_text_index` from within a transaction
- '*column\_name*' cannot be NULL.
- Column '*column\_name*' does not exist in table '*table\_name*'
- Index table mapping failed - Text Index creation aborted
- Invalid text index name - '*index\_name*' already exists
- '*parameter*' is not in the current database
- Server name '*server\_name*' does not exist in `sys.servers`.
- '*table\_name*' does not exist
- '*table\_name*' is not a valid object name
- Table '*table\_name*' does not have an identity column - text index creation aborted

- Text index creation failed
- User '*user\_name*' is not a valid user in the database

**Permissions**

Any user can execute `sp_create_text_index`.

## sp\_drop\_text\_index

### Function

Drops the index table and text indexes.

### Syntax

```
sp_drop_text_index "table_name.index_table_name"  
[, "table_name.index_table_name" ...]
```

### Parameters

*table\_name* – is the name of the table associated with the text indexes you are dropping. *table\_name* has the form [*dbname*.*owner*.]*table*, where:

- *dbname* is the name of the database containing the table.
- *owner* is the name of the owner of the table.
- *table* is the name of the table.

*index\_table\_name* – is the name of the index table and text index you are dropping. *index\_table\_name* has the form [*dbname*.*owner*.]*index*.

### Examples

1. `sp_drop_text_index "blurbs.i_blurbs"`

Drops the index table and text index associated with the *blurbs* table.

### Comments

- First, the `sp_drop_text_index` system procedure issues a remote procedure call (RPC) to the Full-Text Search engine to delete the Verity collection. Then, it removes the associated entries from the *vesaux* and *vesauxcol* tables, drops the index table, and removes the index table object definition.
- Up to 255 indexes can be specified in a single `sp_drop_text_index` request.
- If *database* and *owner* are not specified, the current owner and database are used.

**Messages**

- Can't run `sp_drop_text_index` from within a transaction.
- Index '`index_name`' is not a Text Index
- '`parameter_name`' is not a valid name
- Server name '`server_name`' does not exist in `sys.servers`
- Unable to drop index table '`table_name`'. This table must be dropped manually
- User '`user_name`' is not a valid user in the '`database_name`' database
- `vs_drop_index` failed with code '`code_name`'.

**Permissions**

Any user can execute `sp_drop_text_index`.

## sp\_help\_text\_index

(Enhanced version only)

### Function

Displays a list of text indexes for the current database.

### Syntax

```
sp_help_text_index [index_table_name]
```

### Parameters

*index\_table\_name* – is the name of the text index you want to display.

### Examples

1. `sp_help_text_index`  
Displays all indexes.
2. `sp_help_text_index "i_blurbs"`  
Displays information about the text index *i\_blurbs*.

### Comments

- `sp_help_text_index` is available only with Enhanced Full-Text Search Specialty Data Store.
- If the *index\_table\_name* parameter is specified, information about that text index is displayed. This information includes the name of the text index, the name of the Verity collection for the index, the name of the source table, the name of the IDENTITY column, and the name of the Full-Text Search engine that created the index.
- If *index\_table\_name* is omitted, a list of all text indexes in the current database is displayed

### Messages

- No text indexes found in database '*database\_name*'
- Text index '*index\_name*' does not exist in database '*database\_name*'
- Object must be in the current database

### Permissions

Any user can execute `sp_help_text_index`.

## sp\_optimize\_text\_index

(Enhanced version only)

### Function

Performs optimization on a text index.

### Syntax

```
sp_optimize_text_index index_table_name
```

### Parameters

*index\_table\_name* – is the name of the text index you want to optimize.

*index\_table\_name* has the form [*dbname*.*owner*.]*table*, where:

- *dbname* is the name of the database containing the index table. If present, the *owner* or a placeholder is required.
- *owner* is the name of the owner of the index table.
- *table* is the name of the index table.

### Examples

1. `sp_optimize_text_index "i_blurbs"`

Optimizes the text index *i\_blurbs* to improve query performance.

### Comments

- `sp_optimize_text_index` is available only with Enhanced Full-Text Search Specialty Data Store.
- This system procedure causes the Full-Text Search engine to run the specified text index through the Verity optimization routines.
- `sp_optimize_text_index` is useful for optimizing a text index that has been updated with Verity optimization disabled (trace flag 11 turned on).

### Messages

- '*index\_table\_name*' is not in the current database
- '*index\_table\_name*' does not exist
- Index '*index\_table\_name*' is not a Text Index
- This procedure is not supported against remote server '*server\_name*'

**Permissions**

Any user can execute `sp_optimize_text_index`.

**See Also**

“Updating Existing Indexes” on page 7-1

## sp\_redo\_text\_events

### Function

Changes the status of entries in the *text\_events* table and forces the re-indexing of the modified columns.

### Syntax

```
sp_redo_text_events [from_date [, to_date]]
```

### Parameters

*from\_date* – is the starting date and time in a date range of entries to be modified.

*to\_date* – is the ending date and time in the specified date range of the entries to be modified.

### Examples

```
1. sp_redo_text_events "01/05/98:17:00",  
   "02/12/98:08:30"
```

Re-indexes columns that were modified between January 5, 1998 at 5:00 p.m. and February 12, 1998 at 8:30 a.m.

### Comments

- Resets the status to “unprocessed” for all entries in the *text\_events* table that currently have a status of “processed.” The Full-Text Search engine is notified that a re-index operation is required.
- Useful for synchronizing a text index after a recovery of the Verity collection from a backup. When you use the Enhanced Full-Text Search engine, this procedure is run automatically during *sp\_text\_load\_index*.
- If *to\_date* is omitted, all entries between *from\_date* and the current date with a status of “processed” are reset to “unprocessed.”
- If both *from\_date* and *to\_date* are omitted, all entries in the *text\_events* table with a status of “processed” are reset to “unprocessed.”

### Messages

- *to\_date* cannot be specified without *from\_date*
- You have not specified the full range.



**Permissions**

Any user can execute `sp_redo_text_events`.

## sp\_refresh\_text\_index

### Function

Records modifications in the *text\_events* table when you change source data.

### Syntax

```
sp_refresh_text_index table_name, column_name, rowid,  
mod_type
```

### Parameters

*table\_name* – is the name of the source table being updated. *table\_name* has the form [*dbname*.*owner*.]*table*, where:

- *dbname* is the name of the database containing the table.
- *owner* is the name of the owner of the table.
- *table* is the name of the table.

*column\_name* – is the name of the column being updated.

*rowid* – is the IDENTITY column value of the changed row.

*mod\_type* – specifies the type of the change. Must be *insert*, *update*, or *delete*.

### Examples

```
1. sp_refresh_text_index "blurbs", "copy", 2.000000,  
"update"
```

Records in the *text\_events* table that you have updated the *copy* column of the *blurbs* table. The row you have updated has an *id* of 2.000000.

### Comments

- The user maintains the consistency of the text index. You must run *sp\_refresh\_text\_index* anytime you update source data that has been indexed so that the *text\_events* table reflects the change. This keeps the collections in sync with your source data. The collections are not updated until you run *sp\_text\_notify*.
- You can create triggers that issue *sp\_refresh\_text\_index* for non-*text* and non-*image* columns. For more information on creating

triggers, see “Propagating Changes to the Text Index” on page 3-9.

### Messages

- Column '*column\_name*' does not exist in table '*table\_name*'
- Invalid *mod\_type* specified ('*mod\_type*'). Correct values: INSERT, UPDATE, DELETE
- Owner '*owner\_name*' does not exist
- Table '*table\_name*' does not exist
- '*table\_name*' is not a valid name.
- Text event table not found

### Permissions

Any user can execute `sp_refresh_text_index`.

### See Also

`sp_text_notify`

## sp\_show\_text\_online

### Function

Displays information about databases or indexes that are currently online.

### Syntax

```
sp_show_text_online server_name [, {INDEXES |  
    DATABASES} ]
```

### Parameters

*server\_name* – is the name of the Full-Text Search engine to which the request is sent.

INDEXES | DATABASES – specifies whether the request should contain data about online indexes or online databases. The default is INDEXES.

### Examples

1. **exec sp\_show\_text\_online KRAZYKAT**

Displays all indexes that are currently online in the KRAZYKAT Full-Text Search engine.

2. **exec sp\_show\_text\_online KRAZYKAT, DATABASES**

Displays all databases that are currently online in the KRAZYKAT Full-Text Search engine.

### Comments

- `sp_show_text_online` issues a remote procedure call (RPC) to the Full-Text Search engine to retrieve information about the indexes or the databases that are currently online.
- If the results of this procedure do not list a database, use `sp_text_online` to bring the desired database online.

### Messages

- `sp_show_text_online` failed for server *server\_name*.
- The parameter value '*value*' is invalid
- The RPC sent to the server returned a failure return code
- The second parameter must be INDEXES or DATABASES

**Permissions**

Any user can execute `sp_show_text_indexes`.

**See Also**

`sp_text_online`

## sp\_text\_cluster

(Enhanced version only)

### Function

Displays or changes clustering parameters for the active thread.

### Syntax

```
sp_text_cluster server_name, cluster_parameter [,
cluster_value]
```

### Parameters

*server\_name* – is the name of the Full-Text Search engine.

*cluster\_parameter* – is the name of the clustering parameter.  
Values are shown in Table A-2.

*cluster\_value* – is the value you assign to the clustering parameter for the active thread. Values are shown in Table A-2.

Table A-2: Clustering configuration parameters

Values for <i>cluster_parameter</i>	Values for <i>cluster_value</i>
<i>cluster_style</i>	<p>Specifies the type of clustering to use. Valid values are:</p> <ul style="list-style-type: none"> <li><b>fixed</b> – generates a fixed number of clusters. The number is set by the <b>cluster_max</b> parameter.</li> <li><b>coarse</b> – automatically determines the number of clusters to generate, based on fewer, coarse grained clusters.</li> <li><b>medium</b> – automatically determines the number of clusters to generate, based on medium sized clusters.</li> <li><b>fine</b> – automatically determines the number of clusters to generate, based on smaller, finer grained clusters.</li> </ul>
<i>cluster_max</i>	<p>Specifies the maximum number of clusters to generate when <b>cluster_style</b> is set to <b>fixed</b>. A value of 0 means that the search engine determines the number of clusters to generate.</p>

Table A-2: Clustering configuration parameters (continued)

Values for <i>cluster_parameter</i>	Values for <i>cluster_value</i>
<i>cluster_effort</i>	<p>Specifies the amount of effort (time) that the search engine should expend on finding a good clustering. Valid values are:</p> <ul style="list-style-type: none"> <li>• <i>effort_default</i> – the search engine spends the default amount of time. You can also use the Verity term “default” if you enclose it in double quotes (“”).</li> <li>• <i>high</i> – the search engine spends the longest time.</li> <li>• <i>medium</i> – the search engine spends less time.</li> <li>• <i>low</i> – the search engine spends the least amount of time.</li> </ul>
<i>cluster_order</i>	<p>Specifies the order in which to return the rows within the clusters. Valid values are:</p> <ul style="list-style-type: none"> <li>• “0” – indicates rows are returned in order of similarity to the cluster center. This means the first row returned for a cluster is the one that is most prototypical of the rows in the cluster.</li> <li>• “1” – indicates that rows are returned in the same relative order in which they were submitted for clustering. For example, if cluster 1 contains the first, third and seventh rows found for the query, they will be returned in that relative order within the cluster.</li> </ul>

### Examples

1. `sp_text_cluster KRAZYKAT, cluster_order, "1"`  
Changes the *cluster\_order* parameter to 1 for the active thread.
2. `sp_text_cluster KRAZYKAT, cluster_style`  
Displays the current value of the *cluster\_style* parameter.

### Comments

- The Verity clustering algorithm attempts to group similar rows together, based on the values of the clustering parameters.
- If the *cluster\_parameter* parameter is specified, but the *cluster\_value* parameter is omitted, `sp_text_cluster` displays the value of the clustering parameter that is specified.
- `sp_text_cluster` does not modify the value of the clustering configuration parameter. The *cluster\_value* is valid only for the thread that is currently executing. To modify the default values, use the `sp_text_configure` system procedure.

- For information on how to request a clustered result set, see “Using Pseudo Columns to Request Clustered Result Sets” on page 5-6.

**Messages**

- This procedure is not supported against remote server '*server\_name*'
- The parameter value '*value*' is invalid
- sp\_text\_cluster failed (status = *status*)

**Permissions**

Any user can execute sp\_text\_cluster.

**See Also**

sp\_text\_configure



## sp\_text\_configure

(Enhanced version only)

### Function

Displays or changes Full-Text Search engine configuration parameters.

### Syntax

```
sp_text_configure server_name [, config_name [,  
config_value]]
```

### Parameters

*server\_name* – is the name of the Full-Text Search engine.

*config\_name* – is the name of the configuration parameter to be displayed or modified.

*config\_value* – is the value you assign to the configuration parameter.

### Examples

1. `sp_text_configure KRAZYCAT, backdir, "/data/backup"`  
Changes the backup destination directory to `/data/backup`.
2. `sp_text_configure KRAZYCAT, backdir`  
Displays the backup destination directory.

### Comments

- When you execute `sp_text_configure` to modify a dynamic parameter:
  - The configuration and run values are updated
  - The configuration file is updated
  - The change takes effect immediately
- When you execute `sp_text_configure` to modify a static parameter:
  - The configuration value is updated
  - The configuration file is updated
  - The change takes effect only when you restart the Full-Text Search engine

- When issued with no parameters, `sp_text_configure` displays a report of all Full-Text Search engine configuration parameters and their current values.
- If the `config_name` parameter is specified, but the `config_value` parameter is omitted, `sp_text_configure` displays the report for the configuration parameter specified.
- For information on the individual configuration parameters, see “Modifying the Configuration Parameters” on page 6-4.

#### Messages

- Configuration value cannot be specified without a configuration option
- This procedure is not supported against remote server 'server\_name'
- `sp_text_configure` failed - possible invalid configuration option ('`config_name`')

#### Permissions

Any user can execute `sp_text_configure`.

## sp\_text\_dump\_database

(Enhanced version only)

### Function

Makes a backup copy of a text index.

### Syntax

```
sp_text_dump_database backupdbs [, current_to] [,
current_with] [, current_stripe01 [, ... [,
current_stripe31]]] [, textdb_to] [, textdb_with]
[, textdb_stripe01 [, ... [, textdb_stripe31]]]
```

### Parameters

*backupdbs* – specifies whether the current database and the *text\_db* database are backed up before the text index is backed up. Valid values are shown in Table A-3.

Table A-3: Values for backupdbs

Value	Description
CURRENT_DB_AND_INDEXES	Indicates that the current database is backed up before the text indexes are backed up.
TEXT_DB_AND_INDEXES	Indicates that the <i>text_db</i> database is backed up before the text indexes are backed up.
INDEXES_AND_DATABASES	Indicates that the current and <i>text_db</i> databases are backed up before the text indexes are backed up.
ONLY_INDEXES	Indicates that only the text indexes are backed up.

*current\_to* – is the *to* clause of the *dump database* command for dumping the current database. Use this only if you specify **CURRENT\_DB\_AND\_INDEXES** or **INDEXES\_AND\_DATABASES** for the *backupdbs* parameter.

*current\_with* – is the *with* clause of the *dump database* command for dumping the current database. Use this only if you specify **CURRENT\_DB\_AND\_INDEXES** or **INDEXES\_AND\_DATABASES** for the *backupdbs* parameter.

*current\_stripe* – is the *stripe* clause of the *dump database* command for dumping the current database. Use this only if you specify

CURRENT\_DB\_AND\_INDEXES or INDEXES\_AND\_DATABASES for the *backupdbs* parameter.

*textdb\_to* – is the *to* clause of the *dump database* command for dumping the *text\_db* database. Use this only if you specify INDEXES\_AND\_DATABASES for the *backupdbs* parameter. Use this only if you specify TEXT\_DB\_AND\_INDEXES or INDEXES\_AND\_DATABASES for the *backupdbs* parameter.

*textdb\_with* – is the *with* clause of the *dump database* command for dumping the *text\_db* database. Use this only if you specify TEXT\_DB\_AND\_INDEXES or INDEXES\_AND\_DATABASES for the *backupdbs* parameter.

*textdb\_stripe* – is the *stripe* clause of the *dump database* command for dumping the *text\_db* database. Use this only if you specify TEXT\_DB\_AND\_INDEXES or INDEXES\_AND\_DATABASES for the *backupdbs* parameter.

### Examples

1. `sp_text_dump_database ONLY_INDEXES`

Only text indexes are backed up.

2. `sp_text_dump_database CURRENT_DB_AND_INDEXES, "to '/data/db1backup'"`

The current database is dumped to */data/db1backup* before the text indexes are backed up.

3. `sp_text_dump_database @backupdbs = "TEXT_DB_AND_INDEXES", @textdb_to = "to '/data/textdbbackup'"`

The *text\_db* database is dumped to */data/textdbbackup* before the text indexes are backed up.

4. `sp_text_dump_database @backupdbs = "INDEXES_AND_DATABASES", @current_to = "to '/data/db1backup'", @textdb_to = "to '/data/textdbbackup'"`

The current database is dumped to */data/db1backup* and the *text\_db* database is dumped to */data/textdbbackup* before the text indexes are backed up.

### Comments

- The Full-Text Search engine concatenates the values of *current\_to*, *current\_with*, and *current\_stripe01* to *current\_stripe31* to dump

- database *currentdbname*** and then executes the **dump database** command. The output from the execution of the **dump database** command is sent to the Full-Text Search error log.
- The Full-Text Search engine concatenates the values of *textdb\_to*, *textdb\_with*, and *textdb\_stripe01* to *textdb\_stripe31* to the string “dump database *currentdbname*” and then executes the **dump database** command. The output from the execution of the **dump database** command is sent to the Full-Text Search error log.
  - All entries in the *text\_events* table that have a “processed” status in the current database are deleted when all indexes have been backed up.
  - The backup files for the Verity collections are stored in the directory specified in the *backDir* configuration parameter.

#### Messages

- The parameter value '*value*' is invalid
- Server name '*server*' does not exist in syssservers
- Attempt to dump database '*database\_name*' failed - use the 'dump database' command
- Attempt to backup text indexes on server '*server\_name*' failed
- Attempt to clean *text\_events* in database '*database\_name*' failed (date = '*date*')
- Parameter '*parameter\_name*' is required when dumping database '*database\_name*'
- Dumping database '*database\_name*' - check Full Text Search SDS error log for status

#### Permissions

Any user can execute **sp\_text\_dump\_database**.

#### See Also

**dump\_database** in the *Adaptive Server Reference Manual*

## sp\_text\_kill

(Enhanced version only)

### Function

Terminates all connections to a text index.

### Syntax

```
sp_text_kill index_table_name
```

### Parameters

*index\_table\_name* – is the name of the text index from which all connections will be terminated. *index\_table\_name* has the form [*dbname*.*owner*.]*table*, where:

- *dbname* is the name of the database containing the index table. If present, the *owner* or a placeholder is required.
- *owner* is the name of the owner of the index table.
- *table* is the name of the index table.

### Examples

```
1. sp_text_kill "i_blurbs"
```

Terminates all existing connections to the text index *i\_blurbs*.

### Comments

- `sp_text_kill` is available only with Enhanced Full-Text Search Specialty Data Store.
- This system procedure causes the Full-Text Search engine to terminate all connections to the specified index, except for the connection that initiated the request.
- Attempts to drop a text index that is currently in use will fail. `sp_text_kill` can be used to terminate all existing connections so that the index can be successfully dropped.

### Messages

- Index '*index\_table\_name*' is not a text index
- This procedure is not supported against remote server '*server\_name*'
- '*index\_table\_name*' does not exist

- Only the System Administrator (SA) may execute this procedure

**Permissions**

Only user "sa" can execute sp\_text\_kill.

**See Also**

sp\_drop\_text\_index

## sp\_text\_load\_index

(Enhanced version only)

### Function

Restores a text index backup.

### Syntax

```
sp_text_load_index
```

### Parameters

None.

### Examples

1. `sp_text_load_index`

Restores all text indexes in the current database.

### Comments

- Run `sp_text_load_index` after the `text_db` database and the current database have been fully recovered.
- `sp_text_load_index` restores the Verity collections from the most recent backup. The Full-Text Search engine then runs `sp_redo_text_events` and `sp_text_notify` to reapply all entries in the `text_events` table since the date and time the index was backed up.

### Messages

- Server name '`server_name`' does not exist in `sys.servers`
- Unable to restore text indexes for server '`server_name`'
- This procedure is not supported against remote server '`server_name`'
- Update to `text_events` table in database `database_name` failed for server '`server_name`' - `text_events` not rolled forward

### Permissions

Any user can execute `sp_text_load_index`.



**See Also**

sp\_redo\_text\_events; sp\_text\_notify

## sp\_text\_notify

### Function

Notifies the Full-Text Search engine that the *text\_events* table has been modified.

### Syntax

```
sp_text_notify [{true | false}] [, server_name]
```

### Parameters

*true* – causes the procedure to run synchronously.

*false* – causes the procedure to run asynchronously.

*server\_name* – is the name of the Full-Text Search engine you are notifying.

### Examples

```
1. sp_text_notify true
```

### Comments

- You must run *sp\_text\_notify* after you issue *sp\_refresh\_text\_index* to inform the Full-Text Search engine that the source tables have been modified.
- If you do not specify *true* or *false*, *sp\_text\_notify* runs synchronously.
- If no server name is specified, all Full-Text Search engines are notified.

### Messages

- Can't run *sp\_text\_notify* from within a transaction
- Notification failed, server = '*server\_name*'
- Server name '*server\_name*' does not exist in *syssservers*
- The parameter value '*value*' is invalid

### Permissions

Any user can execute *sp\_text\_notify*.

### See Also

*sp\_refresh\_text\_index*

## sp\_text\_online

### Function

Makes a database available for full-text searches to Adaptive Server.

### Syntax

```
sp_text_online [server_name], [database_name]
```

### Parameters

*server\_name* – is the name of the Full-Text Search engine.

*database\_name* – is the name of the database that you are bringing online.

### Examples

```
1. sp_text_online @database_name = pubs2
```

Makes the *pubs2* database available for full-text searches using the Full-Text Search engine.

### Comments

- If a database is not specified, all databases are brought online for full-text searches.
- If a server name is not specified, all Full-Text Search engines listed in the *vesaux* table are notified.
- With the Enhanced Full-Text Search engine, databases are brought online automatically if the *auto\_online* configuration parameter is set to 1.

### Messages

- All Databases using text indexes are now online
- Databases containing text indexes on server '*database\_names*' are now online
- Server name '*server\_name*' is now online"
- Server name '*server\_name*' does not exist in *sys.servers*.
- The parameter value '*value*' is invalid
- The specified database does not exist
- *vs\_online* failed for server '*server\_name*'

**Permissions**

Any user can execute `sp_text_online`.

# B

## Sample Files

This appendix contains the following:

- The text of the default configuration file (*textsvr.cfg*)
- An overview of the *sample\_text\_main.sql* sample script
- A list of all the sample files provided by the Full-Text Search engine
- An overview of the *getsend* program

### Default *textsvr.cfg* Configuration File

---

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; @(#) File: textsvr.cfg 1.11 03/03/98
;
; Full Text Search Specialty Data Store -- 11.5
; Sample Configuration File
;
; The installation procedure places this file in the
; "SYBASE" directory.
;
; Lines with a semi-colon in column 1 are comment lines.
;
; Modification History:
; -----
; 11-21-97 Create file for Standard Full Text Search SDS
; 03-02-98 Add trace flags and config values for
; Enhanced Full Text Search SDS
;
;
; Copyright (c) 1997, 1998 Sybase, Inc.
; Emeryville, CA.
; All rights reserved.
;
;
; DIRECTIONS
;
; Modifying the textsvr.cfg file:
; -----
; An installation can run the Text Search SDS product
; as supplied, with no modifications to configuration
; parameters. Default values from the executable program
; are in effect.
;
; The "textsvr.cfg" file is supplied with all configuration
```

```

; parameters commented out.
;
; The hierarchy for setting configuration values is:
;
; default value internal to the executable program (lowest)
; configuration file value (overrides default value)
; command line argument (overrides default value and *.cfg file)
;
; Command line arguments are available to override
; settings for these options:
;
; -i<file specification for interfaces file>
; -t (no arg) directs text server to write start-up
; information to stderr (default is DO NOT write start-up
; information)
;
; To set configuration file parameters, follow these steps:
;
; (1) If changing the server name to other than "textsvr":
; (1A) Copy "textsvr.cfg" to "your_server_name.cfg"
; Example: text_server_115.cfg
; (1B) Modify the [textsvr] line to [your_server_name]
; Example: [text_server_115]
; The maximum length of "your_server_name" is 30 characters.
;
; (2) Set any configuration values in the CONFIG VALUES SECTION below.
; Remove the semi-colon from column 1.
;
; ::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
;
; DEFINITIONS OF TRACE FLAG AND SORT ORDER VALUES
;
; "traceflags" parameter, for text server
; Available "traceflags" values: 1,2,3,4,5,6,7,8,9,10,11,12, 13
;
; 1 trace connect/disconnect/attention events
; 2 trace language events
; 3 trace rpc events
; 4 trace cursor events
; 5 log error messages returned to the client
; 6 trace information about indexes
; 7 trace senddone packets
; 8 write text server/Verity api interface records to the log
; 9 trace sql parser
; 10 trace Verity processing
; 11 disable Verity collection optimization
; 12 disable returning of sp_statistics information
; 13 trace backup operations (Enhanced Full Text Search only)
;
; "srv_traceflags" parameter, for Open Server component of text server
; Available "srv_traceflags" values: 1,2,3,4,5,6,7,8
; 1 trace TDS headers

```

```

; 2 trace TDS data
; 3 trace attention events
; 4 trace message queues
; 5 trace TDS tokens
; 6 trace open server events
; 7 trace deferred event queue
; 8 trace network requests
;
; "sort_order" parameter
; Available "sort_order" values: 0,1,2,3
; 0 order by score, descending (default)
; 1 order by score, ascending
; 2 order by timestamp, descending
; 3 order by timestamp, ascending
;
; ::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
;
; CONFIG VALUES SECTION
;
; The "textsvr.cfg" file is supplied with the values commented out.
; To override value(s) in the executable program:
; - Set required value(s) below
; - Remove the semicolon from column 1
;
[textsvr]
;min_sessions = 10
;max_sessions = 100
;batch_size = 500
;sort_order = 0
;defaultDb = text_db
;errorLog = textsvr.log
;language = us_english
;charset = iso_1
;vdkLanguage = english0
;vdkCharset = 850
;traceflags = 0
;srv_traceflags = 0
;max_indexes = 126
;max_packet_size = 2048
;max_stacksize = 34816
;max_threads = 50
;collDir = <${SYBASE location on UNIX}>/sds/text/collections
;collDir = <%SYBASE% location on Win-NT>\sds\text\collections
;vdkHome = <${SYBASE location on UNIX}>/sds/text/verity
;vdkHome = <%SYBASE% location on Win-NT>\sds\text\verity
;interfaces = <${SYBASE location on UNIX}>/interfaces
;interfaces = <%SYBASE% location on Win-NT>\ini\sql.ini
; ::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::::
;
; The parameters in this section apply only to the Enhanced Full Text
; Search SDS.
; If defined to a Standard Full Text Search engine they will be ignored.

```

```

;
:auto_online = 0
;backDir = <$$SYBASE location on UNIX>/sds/text/backup
;backDir = <%SYBASE% location on Win-NT>\sds\text\backup
;knowledge_base =
;nocase = 0
;cluster_max = 0
;cluster_order = 0
;cluster_style = Fixed
;cluster_effort = Default

```

### The *sample\_text\_main.sql* Script

The installation of the Full-Text Search engine copies the *sample\_text\_main.sql* script to the *\$\$SYBASE/sds/text/sample/scripts* directory. This script illustrates the following operations:

- Setting up a text index.
- Modifying data and propagating changes to the collections. This includes inserts, updates, and deletes.
- Dropping a text index.

Execution of this script is not required for installation or configuration; Sybase supplies the script as a sample.

Before you run the *sample\_text\_main.sql* script:

- Your Adaptive Server and Full-Text Search engine must be configured and running.
- Use a text editor to edit the *sample\_text\_main.sql* script. Change “YOUR\_TEXT\_SERVER” to the name of your Full-Text Search engine in Step 4 in the *sample\_text\_main.sql* script.
- Verify that your *model* database contains a *text\_events* table. If your *model* database is **not** configured this way, you need to:
  - Modify the *sample\_text\_main.sql* script to exit after creating the database
  - Apply the *installevent* script to the new database (see “Running the *installevent* Script” on page 3-4)
  - Execute the remainder of the sample script

Direct the script as input to your Adaptive Server. For example, to run the *sample\_text\_main.sql* script on an Adaptive Server named MYSVR:

```

isql -Ulogin -Ppassword -SMYSVR
-i $$SYBASE/sds/text/sample/scripts/sample_text_main.sql -omain.out

```



When you finish with this sample environment, log in to your Adaptive Server and drop the sample database. For example:

```
1> use master
2> go
1> drop database sample_colors_db
2> go
```

The `sample_text_main.sql` script can be rerun.

## Sample Files Illustrating Full-Text Search Engine Features

---

The Full-Text Search engine supplies a set of sample files for illustrating text server operations. The files are located in the `$$YBASE/sds/text/sample/scripts` directory. Execution of the sample files is not required for installation, configuration, or operation of a Full-Text Search engine.

### Custom Thesaurus

---

The following files illustrate how to set up and use a custom thesaurus:

- `sample_text_thesaurus.ctl` – is a sample control file.
- `sample_text_thesaurus.sql` – provides sample queries using the custom thesaurus created by the sample control file.

You can create a custom thesaurus only with the Enhanced Full-Text Search engine. The scripts can be rerun.

### Topics

---

The following files illustrate how to set up and use topics:

- `sample_text_topics.otl` – is a sample outline file.
- `sample_text_topics.kbm` – is a sample knowledge base map.
- `sample_text_topics.sql` – provides sample queries using the defined topics.

Topics is available only with the Enhanced Full-Text Search engine. The scripts can be rerun.

### Clustering, Summarization, and Query-by-Example

---

The following files illustrate how to set up and use clustering, summarization and query-by example:

- *sample\_text\_setup.sql* – creates a sample environment.
- *sample\_text\_queries.sql* – issues queries against the environment and drops the environment.

You can use these scripts only with the Enhanced Full-Text Search engine. These scripts can be rerun as a pair.

### *getsend* Sample Program

---

The Enhanced Full-Text Search engine supplies a program named *getsend* to load *text* or *image* data from a file into a column defined in Adaptive Server.

The required source and header files, a makefile, and directions for building and running the program are included in the directory:

*SSYBASE/sds/text/sample/source*

Refer to the *README.TXT* file and *getsend.c* file for information on how to use the program.

# C

## Unicode Support

The Unicode standard, a subset of the International Standards Organization's ISO 10646 standard, is an international character set. Unicode is identical to the Basic Multilingual Plane (BMP) of ISO 10646, which supports all the major scripts and languages in the world. Therefore, it is a superset of all existing character sets.

The major advantages of Unicode are:

- Provides single-source development. This means you develop an application once and it can then be localized for multiple locales and in multiple languages. By using a single unified character set, you do not have to modify your applications to take into account differences between character sets, thus reducing development, testing, and support costs.
- Allows you to mix different languages in the same database. An all-Unicode system does not require that you design your database to keep track of the character set of your data.

The Enhanced Full-Text Search engine supports Unicode. To use this feature, you need to obtain and install the Unicode Developer's Kit (also known as UDK). This contains everything you need to set up a Unicode-enabled client/server database system.

To configure the Full-Text Search engine to store data in Unicode format, set the `charset` configuration value to `utf8` (see "Modifying the Configuration Parameters" on page 6-4).

► *Note*

---

If you issue wildcard searches against data in Unicode format, turn on trace flag 15. For more information, refer to "Setting Trace Flags" on page 6-10,

---



# Index

## Symbols

- , (comma)
  - in SQL statements xix
- { (curly braces)
  - in SQL statements xix
- ... (ellipsis) in SQL statements xxi
- () (parentheses)
  - in SQL statements xix
- [] (square brackets)
  - in SQL statements xix
- <>(angle brackets), enclosing Verity operators in 5-9

## A

- accrue operator 5-8, 5-11
- Adaptive Server
  - connecting to a Full-Text Search engine 1-1
  - processing a full-text query 2-7
- and operator 5-8, 5-11
  - with the not modifier 5-19
- Angle brackets, enclosing Verity operators in 5-9
- Attention events, tracing 6-11
  - Open Server 6-12
- auto\_online configuration parameter 3-9, 6-6, A-31

## B

- backDir configuration parameter 6-6, 6-17, A-25
- Backup and recovery
  - for the Enhanced version 6-16
  - for the Standard version 6-13
- Backup files
  - default location of 6-6
- Backup operations, tracing 6-11
- batch\_size configuration parameter 6-4

- and performance 7-4
- Brackets. *See* Square brackets [] and Angle brackets <>

## C

- case operator modifier 5-19
- Case sensitivity
  - in queries 5-10
  - setting for the Full-Text Search engine 6-12
  - in SQL xx
- Character sets
  - setting the default 6-8 to 6-9
- charset configuration parameter 6-5
  - setting the default 6-8
- cis cursor rows configuration
  - parameter 7-3
- cis packet size configuration
  - parameter 7-3
- cluster\_effort configuration
  - parameter 5-7, 6-6
  - values for A-19
- cluster\_keywords pseudo column 5-2, 5-7
- cluster\_max configuration parameter 5-7, 6-6
  - values for A-18
- cluster\_number pseudo column 5-2, 5-7
- cluster\_order configuration
  - parameter 5-7, 6-6
  - values for A-19
- cluster\_style configuration parameter 5-7, 6-6
  - values for A-18
- Clustering 5-6
  - configuring for all tables 4-2
  - configuring for individual tables 4-3
  - enabling 4-1
  - modifying values of parameters for A-18
  - setting up 5-7

- in a sort specification 5-5
  - writing queries for 5-7
- collDir configuration parameter 6-5
- Collections 2-2
  - See also* text indexes
  - backing up in the Enhanced
    - version 6-16, A-23
  - backing up in the Standard
    - version 6-14
  - backup and recovery in the Enhanced
    - version 6-16
  - backup and recovery in the Standard
    - version 6-13
  - creating A-4
  - default character set 6-8
  - default language 6-7
  - disabling optimization 6-11, 7-1
  - displaying the names of A-9
  - dropping A-7
  - location of 2-2
  - setting the location of 6-5
  - modifying data in 3-9
  - optimizing A-10
  - performance issues when
    - updating 7-5
  - populating with data 3-7
  - and reindexing A-12
  - restoring from backup in Enhanced
    - version 6-16
  - restoring from backup in Standard
    - version 6-15, 6-17
- Columns
  - valid datatypes to index 2-1
- Comma (.)
  - in SQL statements xix
- Commands in Verity. *See* Operators (commands)
- complement operator 5-8, 5-11
- Component Integration Services
  - connecting to a Full-Text Search
    - engine 1-1
- Configuration file
  - editing parameter values 6-6
  - sample B-1 to B-4
- Configuration parameters 6-4 to 6-6
  - See also* individual configuration parameters
  - auto\_online A-31
  - backDir 6-17, A-25
  - batch\_size parameter and
    - performance 7-4
  - charset 6-8
  - cluster\_effort 5-7, A-19
  - cluster\_max 5-7, A-18
  - cluster\_order 5-7, A-19
  - cluster\_style 5-7, A-18
  - displaying values in the Enhanced
    - version A-21
  - language 6-7
  - max\_sessions parameter and
    - performance 7-4 to 7-5
  - min\_sessions parameter and
    - performance 7-4 to 7-5
  - modifying values in the Enhanced
    - version 6-7, A-21
  - modifying values in the Standard
    - version 6-6
  - nocase 6-12
  - sort\_order 5-4, 6-9
  - srv\_traceflags 6-12
  - vdkCharset 6-8
  - vdkLanguage 6-7
- Configuration parameters, Adaptive Server
  - cis cursor rows 7-3
  - cis packet size 7-3
- Connecting to a Full-Text Search
  - engine 7-6
- Connections, number of user 7-4
- Conventions
  - See also* Syntax
  - directory paths xviii
  - used in manuals xix
- Curly braces ({} )
  - in SQL statements xix
- Cursor events, logging 6-11
- Custom thesaurus 4-7
  - and creating the control file 4-9

- and examining the default thesaurus 4-8
- and the `mksyd` utility 4-10
- and replacing the default thesaurus 4-10

## D

- Databases
  - bringing online for full-text searches 3-9
- Databases, bringing online automatically 6-6
- Datatypes
  - and indexing 3-7
  - of indexed columns 2-1, A-4
- `default_Db` configuration parameter 6-5
- Defining multiple Full-Text Search engines 3-2
- delete operations
  - creating triggers for 3-10
- Deletes
  - and updating the text indexes 2-4
  - from the `text_events` table A-2
  - from the `vesaux` table A-3
- Document filters 2-1
- Document zones
  - and multiple columns in a text index 3-8
  - using with the `in` operator 5-11
- `dump database` command
  - and the `sp_text_dump_database` system procedure 6-17, A-25
  - using in the Standard version 6-14

## E

- Ellipsis (...) in SQL statements xxi
- `errorLog` configuration parameter 6-5
- Error log file
  - setting the path name of 6-5
  - specifying in the `runserver` file 6-2
- Error logging 6-11
- Events, logging 6-10 to 6-12

## F

- Filters, document 2-1
  - creating 4-6
  - and document zones 5-12
- `forceplan`
  - and forcing join orders 7-2
- Full-Text Search engine
  - changing the name of 3-2
  - configuring multiple engines 3-2, 7-5 to 7-6
  - connecting to 7-6
  - document filters 2-1
  - how queries are processed 2-6 to 2-7
  - notifying of updates to the `text_events` table A-30
  - operators 5-8 to 5-18
  - relationship of components 2-6
  - shutting down 6-4
  - starting as a service 6-3
  - starting for UNIX platforms 6-1
  - starting for Windows NT 6-2 to 6-4
  - starting with Sybase Central 6-2
- Full-text search queries
  - bringing databases online for 3-9
  - and case sensitivity 5-10
  - components of 5-1
  - processing a 2-7
  - and requesting clustered result sets 5-7
  - sort order specifications 5-4 to 5-5
  - and using topics 4-14
  - using alternative syntax 5-10
- Full-Text Search Specialty Data Store
  - components of 2-1 to 2-6

## G

- `getsend` program B-6

## H

- `highlight` pseudo column 5-2

**I****IDENTITY columns**

- adding a unique index 3-7
- adding to existing source table 3-6
- displaying with the text index A-9
- example of adding 3-12
- joining with the index table 2-3, 2-7
- scale and precision required 3-6
- in the source table 2-1
- id* pseudo column 2-3, 5-2
  - mapping to the IDENTITY column in the source table 3-6
  - and query optimization 7-2
- index\_any* pseudo column 5-2
  - and query optimization 7-2
- Index table**
  - contents of 2-3
  - creating 3-7, A-4
  - dropping A-7
  - and the *id* column 3-6
  - in a query 2-6
  - joining with the source table 2-3
  - and pseudo columns 2-4, 5-2
- in operator 5-8, 5-11
- insert operations
  - creating triggers for 3-10
- Inserts
  - and updating the text indexes 2-4
- installevent installation script
  - editing 3-4
  - example of using 3-11
  - using 3-4
- installtextserver installation script
  - and creating multiple Full-Text Search engines 7-5
  - editing 3-2, 3-3
  - location of 3-2
- instsvr.exe utility 6-3
- Integrity, maintaining 2-2
- Interfaces
  - tracing calls between Full-Text Search engine and Verity 6-11
- interfaces configuration parameter 6-5
- Interfaces file

- setting the location of 6-5
- specifying in the runserver file 6-2

**J**

- Joining the source table with the text index 2-2, 2-3, 2-6, 3-6, 5-1
  - and increasing performance of 7-2
- Join order
  - ensuring correct 7-2

**K**

- keys* modifier 4-9
- knowledge\_base configuration parameter 4-14, 6-6
- Knowledge base map
  - creating 4-14
  - defining the location of 4-14

**L**

- Language
  - setting the default 6-7 to 6-8
- language configuration parameter 6-5
  - setting the default 6-7
- Language events, logging 6-11
- like operator 5-8, 5-12
  - enabling literal text in the QBE specification 4-1
- list: keyword 4-9
- Logging events using trace flags 6-10 to 6-12

**M**

- Maintaining integrity 2-2
- many operator modifier 5-19
- max\_docs* pseudo column 5-3
  - with clustered result sets 5-7
  - and increasing query performance 7-2
  - and sort orders 6-10



- max\_indexes configuration parameter 6-4
- max\_packet\_size configuration parameter 6-4
- max\_sessions configuration parameter 6-5
  - and performance 7-4 to 7-5
- max\_stack\_size configuration parameter 6-4
- max\_threads configuration parameter 6-4
- Metadata 2-2
- min\_sessions configuration parameter 6-5
  - and performance 7-4 to 7-5
- mksyd utility
  - and creating a custom thesaurus 4-10
  - and examining the default thesaurus 4-8
- mktopics utility 4-13

## N

- Naming the Full-Text Search engine 6-5
- near/n operator 5-8, 5-13
  - with the order modifier 5-19
- near operator 5-8, 5-12, 5-13
- Network requests, tracing 6-12
- nocase configuration parameter 6-6, 6-12
- not operator modifier 5-19

## O

- Online databases. *See* Databases, bringing online
- Open Server events, tracing 6-12
- Open Server trace flags 6-12
- Operator modifiers
  - case 5-19
  - many 5-19
  - not 5-19
  - order 5-19
- Operators (commands) 5-8 to 5-18
  - accrue 5-8, 5-11
  - and 5-8, 5-11
  - complement 5-8, 5-11
  - enclosing in angle brackets 5-9

- in 5-8, 5-11
- like 5-8, 5-12
- near 5-8, 5-12, 5-13
- near/n 5-8, 5-13
- or 5-8, 5-11
- paragraph 5-8, 5-14
- phrase 5-8, 5-13
- product 5-9, 5-14
  - and relevance-ranking 5-3 to 5-4
- sentence 5-9, 5-14
- stem 5-9, 5-15
- sum 5-9, 5-15
- thesaurus 5-9, 5-15
- topic 5-9, 5-16
- wildcard 5-9, 5-17
- word 5-9, 5-18
- yesno 5-9, 5-18

- Optimization, disabling 6-11, 7-1

- order operator modifier 5-19

- or operator 5-8, 5-11

- with the not modifier 5-19

- Outline file for topics 4-12

## P

- paragraph operator 5-8, 5-14
  - with the many modifier 5-19
  - with the order modifier 5-19

- Parameters

- of a search 2-4

- Parentheses ()

- in SQL statements xix

- Performance and tuning

- adding a unique index 3-7

- and using multiple Full-Text Search engines 7-5

- disabling text index optimization 7-1

- increasing query performance 7-2 to 7-3

- reconfiguring Adaptive Server 7-3 to 7-4

- reconfiguring the Full-Text Search engine 7-4 to 7-5

- and sp\_text\_notify 7-5

phrase operator 5-8, 5-13  
 with the *many* modifier 5-19  
 Procedures. *See* System procedures  
 Processed events  
   removing from the *text\_events*  
   table A-2  
 Processing full-text searches 2-6  
 product operator 5-9, 5-14  
 Propagating changes to the  
   collections 2-4  
 Proxy tables as a source table 2-2  
 Pseudo columns 2-4  
   *cluster\_keywords* 5-2, 5-7  
   *cluster\_number* 5-2, 5-7  
   *highlight* 5-2  
   *id* 5-2  
   in a query 2-6  
   *index\_any* 5-2  
   *max\_docs* 5-3, 5-7  
   *score* 5-3 to 5-4  
   *sort\_by* 5-3, 5-4 to 5-5, 5-7  
   *summary* 5-3, 5-6

## Q

QBE specification. *See*  
 Query-by-example  
 Queries  
   and pseudo columns 2-4  
 Queries, full-text search  
   bringing databases online for 3-9  
   and case sensitivity 5-10  
   components of 5-1  
   ensuring the correct join order 7-2  
   increasing performance of 7-2 to 7-3  
   processing of 2-6, 2-7  
   requesting clustered result sets 5-7  
   sort order specifications 5-4 to 5-5  
   and using topics 4-14  
   using alternative syntax 5-10  
 Query-by-example  
   configuring for all tables 4-2  
   configuring for individual tables 4-3  
   enabling 4-1

and the *like* operator 5-12

## R

Ranking documents. *See*  
 Relevance-ranking  
 Recovery  
   and synchronizing a text index with  
   the source table A-12  
   for the Enhanced version 6-16  
   for the Standard version 6-13  
 Relevance-ranking 5-3 to 5-4  
   *See also score* pseudo column  
 Remote procedure calls  
   *sp\_traceoff* 6-11, 7-2  
   *sp\_traceon* 6-11, 7-2  
 Remote tables as a source table 2-2  
 Replicating text indexes 3-10  
 RPC events, logging 6-11  
 RPCs. *See* Remote procedure calls  
 Runserver file 6-1  
   flags for 6-1

## S

Sample files  
   configuration file B-1 to B-4  
   illustrating clustering B-6  
   illustrating custom thesaurus 4-8, B-5  
   illustrating query-by-example B-6  
   illustrating summarization B-6  
   illustrating topics feature 4-12, B-5  
 Sample program *getsend* B-6  
 Sample scripts  
   *sample\_text\_main.sql* 3-6, 3-10, B-4  
   *score* pseudo column 2-4, 5-3 to 5-4  
   with clustered result sets 5-7  
   and default sort order 6-9  
   and the *many* modifier 5-19  
   sorting by 5-5  
   *score* values  
   how Sybase reports 5-4  
 Scripts, sample  
   *sample\_text\_main.sql* 3-6, 3-10, B-4

- Search parameters 2-4
- sentence operator 5-9, 5-14
  - with the *many* modifier 5-19
  - with the *order* modifier 5-19
- Sessions, number of user 7-4
- showplan
  - and examining join orders 7-2
- Shutting down the Full-Text Search engine 6-4
- sort\_by* pseudo column 5-3
  - and requesting a clustered result set 5-7
  - and specifying a sort order 5-4 to 5-5
  - and setting up a defined column as a sort specification 4-4
- sort\_order* configuration parameter 5-4, 6-5, 6-9
- Sort orders
  - and clustered result sets 5-5, 5-7
  - by column 4-4, 5-5
  - in a query 5-4 to 5-5
  - max\_docs* and sort order 6-10
  - by *score* 5-5
  - setting the default 6-9
  - by timestamp 5-5, 6-10
- Sort specifications
  - setting up a defined column to sort by 4-4
- Source tables
  - adding an *IDENTITY* column to 3-6
  - changes to data A-14, A-30
  - contents of 2-1
  - and displaying text indexes A-9
  - in a query 2-6
- sp\_addserver* system procedure 7-6
- sp\_clean\_text\_events* system procedure A-2
- sp\_clean\_text\_indexes* system procedure A-3
- sp\_create\_text\_index* system procedure 3-7, A-4 to A-6
  - creating indexes that use a filter 4-6
  - example of using 3-12
  - specifying multiple columns 3-8
- sp\_drop\_text\_index* system procedure A-7 to A-8
- sp\_help\_text\_index* system procedure A-9
- sp\_optimize\_text\_index* system procedure 7-1, A-10 to A-11
- sp\_redo\_text\_events* system procedure A-12 to A-13
  - and restoring text indexes in Standard version 6-15
- sp\_refresh\_text\_index* system procedure A-14 to A-15
  - modifying data in the collections 3-9
  - running automatically 3-10
- sp\_show\_text\_online* system procedure A-16 to A-17
- sp\_statistics* system procedure
  - disabling 6-11, 7-1
- sp\_text\_cluster* system procedure A-18 to A-20
- sp\_text\_configure* system procedure 6-7, A-21 to A-22
- sp\_text\_dump\_database* system procedure 6-16, A-23 to A-25
- sp\_text\_kill* system procedure A-26 to A-27
- sp\_text\_load\_index* system procedure 6-17, A-28 to A-29
- sp\_text\_notify* system procedure A-30
  - and modifying data in the collections 3-9
  - and performance issues 7-5
  - and restoring text indexes in Standard version 6-15
  - and turning off optimization 7-1
- sp\_text\_online* system procedure 3-9, A-31 to A-32
  - example 3-13
- sp\_traceoff* remote procedure call 6-11, 7-2
- sp\_traceon* remote procedure call 6-11, 7-2
- SQL parsing, tracing 6-11
- Square brackets [ ]
  - in SQL statements xix

**srv\_traceflags** configuration  
     parameter 6-5, 6-12  
**Starting the Full-Text Search engine**  
     from Sybase Central 6-2  
     on UNIX platforms 6-1  
     on Windows NT 6-2 to 6-4  
     as a service 6-3  
**startserver** utility 6-1  
**Start-up**  
     and setting the number of user  
         connections 7-4  
**Start-up commands**  
     and the runserver file 6-1  
     on Windows NT 6-3  
**stem** operator 5-9, 5-15  
     with the many modifier 5-19  
**style.dft** file 4-6  
**style.prm** file  
     editing an existing collection's A-5  
     editing for an existing collection 4-3  
     editing the master 4-2  
     and enabling Verity functionality 4-1  
     location of an existing collection 4-3  
     location of master 4-2  
**style.ufl** file 4-4, 4-6  
**style.vgw** file 4-4, 4-6  
**Summarization**  
     configuring for all tables 4-2  
     configuring for individual tables 4-3  
     enabling 4-1  
     writing queries requesting 5-6  
**summary** pseudo column 5-3  
     enabling before using 4-1  
     using 5-6  
**sum** operator 5-9, 5-15  
**Sybase Central**, starting from 6-2  
**Symbols in SQL statements** xix  
**Synonym list for a custom thesaurus** 4-9  
**synonyms**: statement 4-9  
**Syntax**, alternative Verity 5-10  
**Syntax conventions**, Transact-SQL xviii  
**syservers** table  
     adding Full-Text Search engines 7-6  
**System procedures**

*See also individual system procedures*  
 list of A-1

**sp\_clean\_text\_events** A-2  
**sp\_clean\_text\_indexes** A-3  
**sp\_create\_text\_index** A-4 to A-6  
**sp\_drop\_text\_index** A-7 to A-8  
**sp\_help\_text\_index** A-9  
**sp\_optimize\_text\_index** A-10 to A-11  
**sp\_redo\_text\_events** A-12 to A-13  
**sp\_refresh\_text\_index** A-14 to A-15  
**sp\_show\_text\_online** A-16 to A-17  
**sp\_text\_cluster** A-18 to A-20  
**sp\_text\_configure** A-21 to A-22  
**sp\_text\_dump\_database** A-23 to A-25  
**sp\_text\_kill** A-26 to A-27  
**sp\_text\_load\_index** A-28 to A-29  
**sp\_text\_notify** A-30  
**sp\_text\_online** A-31 to A-32  
**System tables**  
     updating A-1

## T

**TDS data**, tracing 6-12  
**TDS headers**, tracing 6-12  
**TDS tokens**, tracing 6-12  
**text\_db** database 2-2  
     backing up in the Enhanced  
         version 6-16, A-23  
     backing up in the Standard  
         version 6-13, 6-14  
     changing the name of 3-2, 3-5  
     restoring from backup in Enhanced  
         version 6-16, 6-17  
     restoring from backup in Standard  
         version 6-15  
     and the *vesauxcol* table 2-3  
     and the *vesaux* table 2-3  
**text\_events** table 2-4  
     backing up in the Enhanced  
         version 6-16  
     backing up in the Standard  
         version 6-13  
     changing the status of entries A-12

- columns in 2-4
- creating 3-4
- example of creating 3-11
- recording inserts, updates, and deletes A-14
- removing entries from A-2
- restoring from backup in Enhanced version 6-16, 6-17
- restoring from backup in Standard version 6-15
- and `sp_text_dump_database` 6-17, A-25
- and `sp_text_load_index` 6-18

Text documents, types of 2-1

Text indexes

- backing up in the Enhanced version 6-16, A-23
- backing up in the Standard version 6-13, 6-14
- bringing online A-31
- creating 3-7, A-4
- creating and batch sizes 7-4
- displaying a list of A-9
- displaying online A-16
- dropping A-7
- example of creating 3-11 to 3-13
- and the index table 2-3
- metadata 2-2
- that include multiple columns 3-8
- optimizing A-10
- performance issues when updating 7-5
- placing on multiple Full-Text Search engines 7-5
- and reindexing A-12
- replicating 3-10
- restoring from backup in Enhanced version 6-16
- restoring from backup in Standard version 6-15, 6-17
- setting location of backup files 6-6
- and tracing information 6-11
- update using `text_events` table 2-4
- updating 7-1
- using a document filter with 4-6

- `textsrvr.cfg` file
  - sample B-1 to B-4
- Thesaurus, custom 4-7
  - and creating the control file 4-9
  - and examining the default thesaurus 4-8
  - and the `mksyd` utility 4-10
  - and replacing the default thesaurus 4-10
- thesaurus operator 5-9, 5-15
  - using a custom thesaurus 4-7
- Timestamp
  - sorting by 6-10
- topicEDITOR 4-13
- topic operator 4-14, 5-9, 5-16
- Topics
  - creating a knowledge base map 4-14
  - creating an outline file 4-12
  - creating a topic set directory 4-13
  - creating complex relationships 4-13
  - description of 4-11
  - executing queries using 4-14
  - sample files 4-12
  - troubleshooting 4-15
- Topic set directories 4-13
  - mapping to 4-14
- Trace flags 6-10
  - enabling trace flags 11 and 12 7-1
- Open Server 6-12
  - setting to examine join orders 7-2
- traceflags configuration parameter 6-5
- Triggers for running
  - `sp_refresh_text_index` 3-10

## U

- Unicode
  - and wildcard searches 6-11
- Unique index
  - adding to an IDENTITY column 3-7
  - example of creating 3-12
- update operations
  - creating triggers for 3-10
- Updates

- and updating the text indexes 2-4
- update statistics
  - disabling 7-1
- Updating indexes 7-1
- User
  - connections 7-4
  - sessions 7-4
- User databases
  - backing up in the Enhanced version A-23
  - backing up in the Standard version 6-14, 6-16
  - bringing online automatically 6-6
  - bringing online for full-text searches 3-9, A-31
  - displaying a list of text indexes for A-9
  - displaying online A-16
  - restoring from backup in Enhanced version 6-16, 6-17
  - restoring from backup in Standard version 6-15
- User table. *See* Source table

## V

- vdkCharset configuration parameter 6-5
  - setting the default 6-8
- vdkHome configuration parameter 6-5
- vdkLanguage configuration parameter 6-5
  - setting the default 6-7
- Verity
  - setting the Verity directory 6-5
  - tracing Verity processing 6-11
- Verity collections. *See* Collections
- Verity query. *See* Full-text search queries
- Verity Search '97 1-1
- vesauxcol table
  - columns in 2-3
  - removing entries when dropping text indexes A-7
  - updating 3-7
- vesaux table
  - columns in 2-3

- creating entries A-5
- removing entries from A-3
- removing entries when dropping text indexes A-7
- updating 3-7

## W

- wildcard operator 5-9, 5-17
  - using with data in Unicode format 6-11
  - with the case modifier 5-19
  - with the many modifier 5-19
- Windows NT
  - directory paths xviii
- word operator 5-9, 5-18
  - with the case modifier 5-19
  - with the many modifier 5-19
- writetext command, using triggers
  - with 3-10

## Y

- yesno operator 5-9, 5-18

## Z

- Zones. *See* Document zones