

---

# **MarkLogic Server**

---

## **Administrator's Guide**

Release 3.2  
June, 2008

## Copyright

© Copyright 2002-2008 by Mark Logic Corporation. All rights reserved worldwide.

This Material is confidential and is protected under your license agreement.

Excel and PowerPoint are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. This document is an independent publication of Mark Logic Corporation and is not affiliated with, nor has it been authorized, sponsored or otherwise approved by Microsoft Corporation.

Contains LinguistX, from Inight Software, Inc. Copyright © 1996-2006. All rights reserved. [www.inight.com](http://www.inight.com).

Antenna House OfficeHTML Copyright © 2000-2006 Antenna House, Inc. All rights reserved.

Argus Copyright ©1999-2005 Icen Technology Ltd. All rights reserved.

---

---

**Table of Contents**

---

---

**Administrator's Guide**

Copyright .....	2
<b>1.0 Introduction .....</b>	<b>10</b>
1.1 Objectives .....	10
1.2 Audience .....	10
1.3 Scope and Requirements .....	10
<b>2.0 Administrative Interface .....</b>	<b>11</b>
2.1 Overview of the Admin Interface .....	11
2.2 Accessing the Admin Interface .....	12
2.3 Logging Off the Admin Interface .....	12
2.4 Creating and Managing Administrators .....	12
<b>3.0 Starting and Stopping MarkLogic Server .....</b>	<b>13</b>
3.1 Starting the Server .....	13
3.2 Stopping the Server .....	13
3.2.1 Using System Command to Stop MarkLogic Server .....	14
3.2.2 Using the Admin Interface to Stop MarkLogic Server .....	14
3.3 Restarting the Server .....	14
<b>4.0 Groups .....</b>	<b>16</b>
4.1 Overview of Groups .....	16
4.2 Example .....	17
4.3 Procedures for Configuring and Managing Groups .....	18
4.3.1 Creating a New Group .....	18
4.3.2 Viewing Group Settings .....	19
4.3.3 Deleting a Group .....	19
4.3.4 Configuring an SMTP Server .....	20
4.3.5 Restarting All Hosts in a Group .....	21
<b>5.0 HTTP Servers .....</b>	<b>22</b>
5.1 HTTP Server Overview .....	22
5.2 Procedures for Creating and Managing HTTP Servers .....	22
5.2.1 Creating a New HTTP Server .....	23
5.2.2 Viewing HTTP Server Settings .....	26
5.2.3 Deleting an HTTP Server .....	26

5.2.4	Canceling a Request .....	27
6.0	XDBC Servers .....	29
6.1	XDBC Server Overview .....	29
6.2	Procedures for Creating and Managing XDBC Servers .....	30
6.2.1	Creating a New XDBC Server .....	30
6.2.2	Viewing XDBC Server Settings .....	32
6.2.3	Deleting an XDBC Server .....	32
7.0	WebDAV Servers .....	33
7.1	WebDAV Server Overview .....	33
7.1.1	Accesses a Database for Read and Write, Not XQuery Execution .....	34
7.1.2	WebDAV Server Security .....	34
7.1.3	Directories .....	35
7.1.3.1	Automatic Directory Creation in a Database Settings .....	35
7.1.3.2	Properties and URIs of Directories .....	36
7.1.4	Server Root Directory .....	36
7.1.5	Documents in a WebDAV Server .....	37
7.2	Procedures for Creating and Managing WebDAV Servers .....	37
7.2.1	Creating a New WebDAV Server .....	37
7.2.2	Viewing WebDAV Server Settings .....	39
7.2.3	Deleting a WebDAV Server .....	39
7.3	WebDAV Clients .....	40
7.3.1	Tested WebDAV Clients .....	40
7.3.2	General Steps to Connect to a Server .....	41
7.3.3	Steps to Connect to a Web Folder in Windows Explorer .....	41
7.4	Example: Setting Up a WebDAV Server to Add/Modify Documents Used By Another Server .....	42
8.0	Databases .....	44
8.1	Understanding Databases .....	44
8.1.1	Schemas and Security Databases .....	44
8.1.2	Modules Database .....	45
8.1.3	Triggers Database .....	45
8.1.4	Database Settings .....	46
8.1.4.1	Basic Administrative Settings .....	46
8.1.4.2	Index Settings that Affect Documents .....	46
8.1.4.3	Reindexing Settings .....	49
8.1.4.4	Document and Directory Settings .....	49
8.1.4.5	Memory and Journal Settings .....	51
8.1.4.6	Other Settings .....	53
8.1.4.7	Merge Control Settings .....	54
8.1.5	Example of Databases in MarkLogic Server .....	54
8.2	Creating a New Database .....	55
8.3	Attaching Forests to the Database .....	55

8.4	Viewing Database Settings .....	56
8.5	Loading Documents into a Database .....	56
8.6	Detaching a Forest from a Database .....	57
8.7	Deleting a Database .....	58
9.0	Word Query Database Settings .....	60
9.1	Understanding the Word Query Configuration .....	60
9.1.1	Overview of Configuration Options .....	60
9.1.2	Understanding Which Elements are Included and Excluded .....	61
9.1.3	Adding a Weight to Boost or Lower the Relevance of an Included Element 63	
9.1.4	Specifying An Attribute Value for an Included Element .....	63
9.1.5	Understanding the Index Option Configuration .....	64
9.2	Configuring Customized Word Query Settings .....	64
10.0	Fields Database Settings .....	68
10.1	Overview of Fields .....	68
10.2	Understanding Field Configurations .....	69
10.2.1	Overview of Field Configuration Options .....	69
10.2.2	Understanding Which Elements are Included and Excluded .....	70
10.2.3	Adding a Weight to Boost or Lower the Relevance of an Included Element 71	
10.2.4	Specifying An Attribute Value for an Included Element .....	72
10.2.5	Understanding the Index Option Configuration .....	72
10.3	Field Word Lexicons .....	73
10.4	Configuring Fields .....	73
10.4.1	Configuring a New Field .....	73
10.4.2	Modifying an Existing Field .....	79
11.0	Understanding and Controlling Database Merges .....	80
11.1	Overview of Merges: Merges are Good .....	80
11.1.1	Dynamic and Self-Tuning .....	80
11.1.2	What Happens During a Merge .....	81
11.1.3	Dangers of Disabling Merges .....	81
11.1.4	Merges Will Change Scores .....	82
11.2	Setting Merge Policy .....	82
11.2.1	Overview of the Merge Policy Controls .....	82
11.2.2	Description on Merge Parameters .....	83
11.3	Blackout Periods for Merges .....	84
11.3.1	Understanding Merge Blackouts .....	85
11.3.2	Configuring Merge Blackout Periods .....	85
11.3.3	Deleting Merge Blackout Periods .....	86
11.4	Merges and Point-In-Time Queries .....	86
11.5	Monitoring a Merge .....	86
11.5.1	Messages in the ErrorLog.txt File .....	86

11.5.2	Database Status Page .....	87
11.6	Explicit Merge Commands .....	87
11.6.1	Manually Initiating a Merge .....	87
11.6.2	Cancelling a Merge .....	88
11.7	Configuring Merge Policy Rules .....	89
11.7.1	Determine the Baseline for Your Merges .....	89
11.7.2	If You Want to Avoid ‘Large’ Merges .....	89
11.7.3	Other Solutions .....	90
12.0	Backing Up and Restoring a Database .....	91
12.1	Backup and Restore Overview .....	91
12.1.1	Consistent, Database-Level Backup .....	92
12.1.2	Admin Interface .....	92
12.1.3	Backup and Restore Transactions .....	92
12.1.4	Backup Directory Structure .....	93
12.1.5	Phases of Backup or Restore Operation .....	94
12.1.5.1	Validation Phase .....	94
12.1.5.2	Copy Phase .....	95
12.1.5.3	Synchronization Phase .....	95
12.1.6	Notes about Backup and Restore Operations .....	95
12.2	Backing Up a Database .....	96
12.3	Restoring a Database .....	99
13.0	Hosts .....	101
13.1	Adding a Host to a Cluster .....	101
13.2	Changing the Group of the Host .....	102
13.3	Shutting Down or Restarting a Host .....	103
13.4	Clearing a Forest on a Host .....	103
13.5	Deleting a Forest on a Host .....	104
13.6	Leaving the Cluster .....	104
14.0	Forests .....	106
14.1	Creating a Forest .....	106
14.2	Making Backups of a Forest .....	107
14.3	Restoring a Forest .....	109
14.4	Clearing Data in a Forest .....	109
14.5	Deleting a Forest from a Host .....	110
15.0	Security Administration .....	112
15.1	Security Entities .....	112
15.2	Users .....	115
15.2.1	Creating a User .....	115
15.2.2	Viewing a User’s Configuration .....	117
15.2.3	Deleting a User .....	118

15.3	Roles .....	118
15.3.1	Creating a Role .....	120
15.3.2	Viewing a Role .....	121
15.3.3	Deleting a Role .....	122
15.4	Execute Privileges .....	122
15.4.1	Creating an Execute Privilege .....	123
15.4.2	Viewing an Execute Privilege .....	124
15.4.3	Deleting an Execute Privilege .....	124
15.5	URI Privileges .....	125
15.5.1	Creating a URI Privilege .....	126
15.5.2	Viewing a URI Privilege .....	126
15.5.3	Deleting a URI Privilege .....	127
15.6	Amps .....	127
15.6.1	Creating an Amp .....	128
15.6.2	Viewing an Amp .....	129
15.6.3	Deleting an Amp .....	130
15.7	Collections .....	130
15.7.1	Creating a Collection .....	131
15.7.2	Viewing a Collection .....	132
15.7.3	Removing a Permission from a Collection .....	132
15.7.4	Deleting a Collection .....	133
15.8	Realm .....	134
15.8.1	Setting the Realm .....	135
15.8.2	Changing the Realm .....	135
16.0	Text Indexing .....	137
16.1	Text Indexes .....	137
16.1.1	Understanding the Text Index Settings .....	138
16.1.2	Viewing Text Index Configuration .....	143
16.1.3	Configuring Text Indexes .....	145
16.2	Phrasing and Element-Word-Query Boundary Control .....	146
16.2.1	Phrasing Control .....	146
16.2.2	Element Word Query Throughs .....	148
16.2.3	Procedures .....	148
16.2.3.1	Viewing Phrasing and Element-Word-Query Settings .....	148
16.2.3.2	Configuring Phrasing and Element-Word-Query Settings .....	149
16.2.3.3	Deleting a Phrasing or Element-Word-Query Setting .....	151
16.3	Query Behavior with Reindex Settings Enabled and Disabled .....	152
16.3.1	Understanding the Reindexer Enable Settings .....	152
16.3.2	Query Evaluation According to the Lowest Common Denominator .....	153
16.3.3	Reindexing Does Not Apply to Point-In-Time Versions of Fragments .....	153
16.3.4	Example Scenario .....	154
17.0	Element and Attribute Range Indexes and Lexicons .....	155
17.1	Understanding Element and Attribute Range Indexes .....	155

17.2	Using Range Indexes for Element and Attribute Value Lexicons .....	157
17.3	Understanding Element and Attribute Word Lexicons .....	158
17.4	Viewing Element Range Index or Element Word Lexicon Settings .....	158
17.5	Defining Element Range Indexes or Element Word Lexicons .....	158
17.6	Viewing Attribute Range Index and Attribute Word Lexicon Settings .....	160
17.7	Defining Attribute Range Indexes or Attribute Word Lexicons .....	160
17.8	Defining Element or Attribute Value Lexicons .....	162
17.9	Deleting Range Indexes or Lexicons .....	162
<b>18.0</b>	<b>Fragments .....</b>	<b>164</b>
18.1	Choosing a Fragmentation Strategy .....	165
18.1.1	Fragment Roots .....	166
18.1.2	Fragment Parents .....	166
18.2	Defining Fragment Roots .....	167
18.3	Defining Fragment Parents .....	168
18.4	Viewing Fragment Rules .....	169
18.5	Deleting Fragment Rules .....	170
<b>19.0</b>	<b>Namespaces .....</b>	<b>171</b>
19.1	Defining Namespaces for a Group .....	171
19.2	Defining Namespaces for an HTTP or XDBC Server .....	172
19.3	Viewing Namespace Settings for a Group .....	173
19.4	Viewing Namespace Settings for an HTTP or XDBC Server .....	174
19.5	Deleting Namespaces for a Group .....	175
19.6	Deleting Namespaces for an HTTP or XDBC Server .....	175
<b>20.0</b>	<b>Understanding and Defining Schemas .....</b>	<b>177</b>
20.1	Understanding Schemas .....	177
20.2	Procedures For Defining Schemas .....	178
20.2.1	Adding a Schema Definition for a Group .....	178
20.2.2	Adding a Schema Definition for an HTTP or XDBC Server .....	179
20.2.3	Viewing Schema Definitions for a Group .....	180
20.2.4	Viewing Schema Definitions for an HTTP or XDBC Server .....	181
20.2.5	Deleting a Schema Definition for a Group .....	182
20.2.6	Deleting a Schema Definition for an HTTP or XDBC Server .....	182
<b>21.0</b>	<b>Log Files .....</b>	<b>184</b>
21.1	Understanding the Log Levels .....	184
21.2	Configuring Log Files .....	185
21.3	Viewing the System Log .....	186
21.4	Viewing the File Log .....	187
<b>22.0</b>	<b>Appendix A: ‘Hot’ versus ‘Cold’ Admin Tasks .....</b>	<b>188</b>
22.1	Groups .....	189

22.2	HTTP, XDBC, and WebDAV Servers .....	190
22.3	Databases .....	190
22.4	Hosts .....	190
22.5	Forests .....	191
22.6	Mimetypes .....	191
22.7	Security .....	191
23.0	Appendix B: Pre-defined Execute Privileges .....	192
24.0	Appendix C: Pre-defined Roles .....	204
24.1	admin .....	204
24.2	admin-builtins .....	204
24.3	domain-management .....	206
24.4	filesystem-access .....	206
24.5	merge .....	206
24.6	pipeline-management .....	207
24.7	security .....	207
24.8	trigger-management .....	209
	Technical Support .....	210

## 1.0 Introduction

MarkLogic Server is a powerful software solution for harnessing your digital content base. MarkLogic Server enables you to build complex applications that interact with large volumes of XML, SGML, HTML and other popular content formats. MarkLogic Server's unique architecture ensures that your applications are both scalable and high-performance, delivering query results at search-engine speeds while providing transactional integrity over the underlying content repository.

### 1.1 Objectives

This document describes administrative tasks required to manage the operation of MarkLogic Server on your system.

### 1.2 Audience

This document is intended for a technical audience, specifically the system administrator in charge of MarkLogic Server.

### 1.3 Scope and Requirements

This guide explains administrative tasks for MarkLogic Server running on the following platforms:

- Microsoft Windows XP SP2\*, Microsoft Windows 2003 Server (x86)
- Windows 2003 Server 64-bit Edition (x64)
- Sun Solaris 8, 9, and 10 (64-bit SPARC)
- Sun Solaris 10 (x64)
- Red Hat Enterprise Linux 3.0, and 4.0 (x86)
- Red Hat Enterprise Linux 3.0 and 4.0 (x64)

\* Microsoft Windows XP is supported for development only.

This document only explains the administrative tasks for the software. To learn how to get started using the software, or how to install the software, refer to the appropriate documents:

- *Getting Started With MarkLogic Server*
- *MarkLogic Server Installation Guide*

This document assumes that you have successfully completed all the tasks in *Getting Started with MarkLogic Server*. If not, be sure to complete these basic tasks before doing any administrative work for MarkLogic Server.

## 2.0 Administrative Interface

The MarkLogic Server administrative interface (or Admin Interface) is used to configure the MarkLogic Server software on your system. This chapter provides a general overview of the Admin Interface and includes the following sections:

- [Overview of the Admin Interface](#)
- [Accessing the Admin Interface](#)
- [Logging Off the Admin Interface](#)
- [Creating and Managing Administrators](#)

### 2.1 Overview of the Admin Interface

With the Admin Interface, you can complete any of the following tasks:

- Manage basic software configuration
- Create and configure groups
- Create and manage databases
- Create and manage new forests
- Back up and restore forest content
- Create and manage new web server and Java-language access paths
- Create and manage security configurations
- Tune system performance
- Configure namespaces and schemas
- Check the status of resources on your systems

The Admin Interface is implemented as a MarkLogic Server web application. By default, it runs on port 8001 of your hosts. If you have completed the basic tasks in the *Getting Started with MarkLogic Server* manual, then accessing the Admin Interface requires that you enter a user name and password. After you have been authenticated, you should not need to re-enter your user name and password to complete any of the other tasks outlined in this guide during the current session.

Some configurations changes require the server to restart to reflect the changes. Configuration changes that do not require the server to restart to reflect the changes are defined as “hot”. In an Enterprise Edition clustered deployment, “cold” tasks will require all of the hosts in the cluster to restart their instance of MarkLogic Server in order to reflect the changes. In an Enterprise Edition single-server deployment and in all Standard Edition deployments, “cold” tasks will cause MarkLogic Server to restart in order to reflect the changes. For a list of which tasks are “hot” and which are “cold,” see “Appendix A: ‘Hot’ versus ‘Cold’ Admin Tasks” on page 188.

## 2.2 Accessing the Admin Interface

To access the Admin Interface, complete the following procedure:

1. Open the following URL in a browser:

<http://localhost:8001/>

**Note:** If you are not accessing the Admin Interface from the same system on which MarkLogic Server is running, you will have to use the IP address or domain name of the server instead of `localhost`.

2. Log in with your admin user name and password.

The summary screen for the Admin Interface displays.

**Note:** If you have already logged on as an admin user during this session, you do not have to log in again.

From the summary screen, you can see and click on many of the items configured in MarkLogic Server. The summary screen displays all of the Databases, App Servers, Groups, Forests, Security objects, and Hosts configured for your system. If you click on any object or category, the Admin Interface takes you to a more detailed page for the object or category.

## 2.3 Logging Off the Admin Interface

To log off the Admin Interface, close the browser window used to access the Admin Interface. This action is sufficient to end the current session and force the user to authenticate again starting another session.

## 2.4 Creating and Managing Administrators

MarkLogic Server administrators are managed by defining which user has the Admin role. For the procedures for creating, managing and removing administrators, see “Security Administration” on page 112.

## 3.0 Starting and Stopping MarkLogic Server

Use the following procedures to start and stop MarkLogic Server:

- [Starting the Server](#)
- [Stopping the Server](#)
- [Restarting the Server](#)

### 3.1 Starting the Server

To start MarkLogic Server, use the appropriate system command for your platform:

Platform	Command
Microsoft Windows	Select Start > Programs > MarkLogic Server > Start MarkLogic Server  <b>Note:</b> When you start MarkLogic Server from the Start menu, the Windows service configuration for MarkLogic Server is set to start automatically. Also, if you are using Windows Vista, to start the service you must right-click the Start MarkLogic Server link in the Start menu and choose Run as Administrator, then choose to allow the action.
Red Hat Linux	<code>/etc/init.d/MarkLogic start</code>
Sun Solaris	<code>/etc/init.d/MarkLogic start</code>

### 3.2 Stopping the Server

There are two ways to perform a clean shutdown of MarkLogic Server:

- [Using System Command to Stop MarkLogic Server](#)
- [Using the Admin Interface to Stop MarkLogic Server](#)

### 3.2.1 Using System Command to Stop MarkLogic Server

You can stop MarkLogic Server with the appropriate system command for your platform:

Platform	Command
Microsoft Windows	Select Start > Programs > MarkLogic Server > Stop MarkLogic Server  <b>Note:</b> If you are using Windows Vista, to stop the service you must right-click the Stop MarkLogic Server link in the Start menu and choose Run as Administrator, then choose to allow the action.
Red Hat Linux	<code>/etc/init.d/MarkLogic stop</code>
Sun Solaris	<code>/etc/init.d/MarkLogic stop</code>

### 3.2.2 Using the Admin Interface to Stop MarkLogic Server

To stop the server from the Admin Interface, complete the following procedure:

1. Click the Hosts icon on the left tree menu.
2. Click on the name of the host you want to shut down.
3. Click the Status tab on the top right.
4. Click Shutdown.
5. A confirmation message displays while shutting down. Click OK to shut down the server.

**Note:** MarkLogic Server must be running in order for you to use the Admin Interface. Once you have stopped the server, you will no longer be able to access the Admin Interface until you start MarkLogic Server again; to restart the server, run the system command for your platform as described in “Starting the Server” on page 13.

### 3.3 Restarting the Server

To restart the server from the Admin Interface, complete the following procedure:

1. Click the Hosts icon on the left tree menu.
2. Click the Status tab on the top right.

3. Click Restart.
4. A confirmation message displays while restarting. Click OK to restart MarkLogic Server.

You may also manually stop and start the server as described above.

**Note:** The restart operation normally completes within a few seconds. It is possible, however, for it to take longer under some conditions (for example, if the Security database needs to run recovery or if the connectivity between hosts in a cluster is slow). If it takes longer than a few seconds for MarkLogic Server to restart, than the Admin Interface might return a `503: Service Unavailable` message. If you encounter this situation, wait several seconds and then reload the Admin Interface.

## 4.0 Groups

This chapter describes Groups in MarkLogic Server, and includes the following sections:

- [Overview of Groups](#)
- [Example](#)
- [Procedures for Configuring and Managing Groups](#)

### 4.1 Overview of Groups

The following are the basic definitions for Group, Host, and cluster:

- A *Group* is a set of similarly configured Hosts within a cluster.
- A *Host* is an instance of MarkLogic Server running on a single machine.
- A *cluster* is a set of Hosts that work together.

For Standard Edition configurations, you can only use one group at a time (because there is only one host). For Enterprise Edition configurations with multiple hosts, you can have as many group configurations as makes sense in your environment.

Groups allow you to have several configurations, each of which applies to a distinct set of Hosts. Different configurations are often needed when different hosts perform different tasks, or when the hosts have different system capabilities (disk space, memory, and so on). In Enterprise Edition clusters, a common configuration is to have one group defined for the *evaluator* nodes (hosts that service query requests) and another group defined for the *data* nodes (hosts to which forests are attached).

HTTP, XDBC, and WebDAV servers are defined at the Group level and apply to all hosts within the group. Schemas and namespaces can also be defined at the group level to apply group-wide.

The Group configuration page allows you to define configuration information for memory settings, SMTP server settings, and other configuration settings. The values for the settings are set at installation time based on your system memory configuration at the time of the installation. For a description of each configuration option, see the help tab of the group configuration page in the Admin Interface.

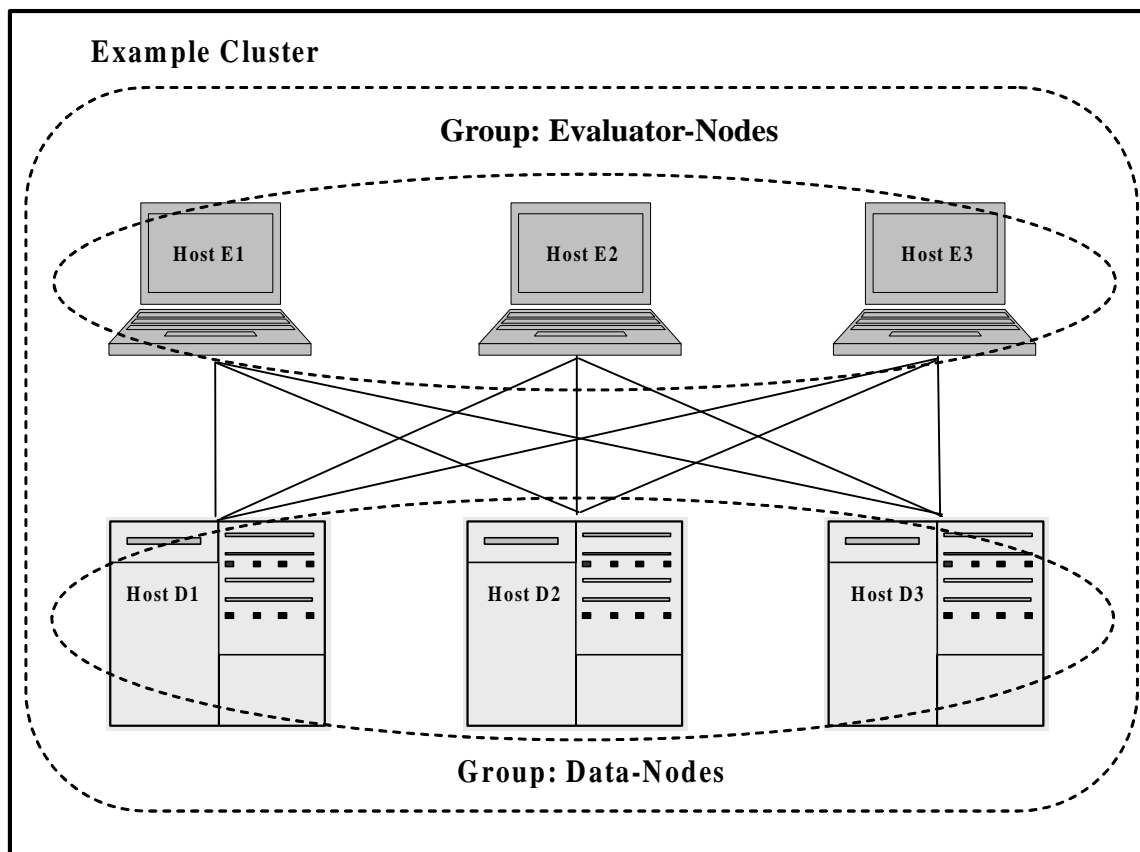
## 4.2 Example

The relationships between a cluster, a Group and a Host in MarkLogic Server may be best illustrated with an example.

In this example, each machine is set up as a **Host** within the Example **Cluster**.

**Hosts** E1, E2 and E3 belong to a **Group** called Evaluator-Nodes. They are configured with HTTP servers and XDBC servers to run user applications. All **Hosts** in Evaluator-Nodes have the same MarkLogic Server configuration.

**Hosts** D1, D2 and D3 belong to a **Group** called Data-Nodes. **Hosts** in Data-Nodes are configured with data forests and interact with Evaluator-Nodes to service data requests. See the sections on Databases, Forests and Hosts for details on configuring data forests.



**Note:** If you are administering a single-host, Standard Edition MarkLogic Server environment, the host is automatically added to a Default group during the installation process. You will only have one host in the group and will not be able to add other hosts to the group. To set up a multiple-host cluster, Enterprise Edition is required.

## 4.3 Procedures for Configuring and Managing Groups

The following procedures describe how to create and manage groups in MarkLogic Server:

- [Creating a New Group](#)
- [Viewing Group Settings](#)
- [Deleting a Group](#)
- [Configuring an SMTP Server](#)
- [Restarting All Hosts in a Group](#)

### 4.3.1 Creating a New Group

To create a new group, perform the following steps:

1. Log into the Admin Interface.
2. Click the Groups icon on the left tree menu.
3. Click the Create tab on the Group Summary page.

The Add Group page will display.

**Create Group**

Summary Create Help

ok cancel

**group** -- *The server specification.*

**group name**   
The name of a group of servers.  
**Required. You must supply a value for group-name.**

**list cache size\***   
The size of the list cache, in megabytes.

**list cache partitions\***   
The number of list cache partitions.

4. Go to the group name field and enter a short hand name for the group.

MarkLogic Server will use this name to refer to the group.

5. You can change the value of list cache size, compressed tree cache size and expanded tree cache size or leave the defaults. They specify the amount of memory dedicated to caching term list, tree data in compressed form and tree data in expanded form.
6. Go to the system log level. System log level specifies the minimum log level messages sent to the operating system. Log levels are listed in decreasing level of log details. You may change the system log level or leave it at the default level.
7. Go to the file log level. File log level specifies the minimum log level messages sent to the log file. Log levels are listed in decreasing level of log details. You may change the file log level or leave it at the default level.
8. The rotate log files field specifies how often to start a new log file. You may change this field or use the default value provided.
9. The keep log files field specifies how many log files are kept. You may change this field or use the default value provided.
10. Click OK.

Adding a group is a hot administrative task; the server does not need to restart to reflect your changes.

### **4.3.2 Viewing Group Settings**

To view the settings for a particular group, perform the following steps:

1. Log into the Admin Interface.
2. Click the Groups icon on the left tree menu.
3. Click the Configure tab at the top right.
4. Locate the group for which you want to view settings.
5. Click the icon for this group.
6. View the settings.

### **4.3.3 Deleting a Group**

You must drop all hosts assigned to a group before you can delete a group. To delete a group, perform the following steps:

1. Log into the Admin Interface.
2. Click the Groups icon on the left tree menu.

3. Click the Configure tab at the top right.
4. Locate the Group to be deleted.
5. Click on Hosts to check that there is no host assigned to the group. All hosts assigned to a group must be dropped before the group can be deleted. Dropping a host from a group does not drop the host from the cluster.
6. Click the icon for this group again.
7. Click Delete. Deleting a group deletes it from the system.
8. A confirmation message displays. Click OK to permanently delete the group.

Deleting a group is a hot operation; the server does not need to restart to reflect your changes.

#### 4.3.4 Configuring an SMTP Server

The installation process configures an SMTP server based on the environment at installation time. A single SMTP server is configured for all of the hosts in a group. The SMTP configuration is used when applications use the `xcmp:email` function.

To change the SMTP server or the SMTP timeout for the system (the time after which SMTP requests fail with an error), perform the following steps:

1. Log into the Admin Interface.
2. Click the Groups icon on the left tree menu.
3. Click the Configure tab at the top right.
4. In the SMTP Relay field, enter the hostname for your SMTP server.
5. In the SMTP Timeout field, enter the time (in seconds) after which requests will time out.
6. Click OK.

Changing any SMTP settings is a hot operation; the server does not need to restart to reflect your changes.

### 4.3.5 Restarting All Hosts in a Group

Perform the following steps to restart all the hosts in a group from the Admin Interface:

1. Click the Groups icon on the left tree menu.
2. Click the name of the group you want to restart, either from the menu tree or from the Group Summary page.
3. Click the Status tab on the top right.
4. Click Restart.
5. A confirmation message displays while restarting. Click OK to restart all of the hosts in the MarkLogic Server group.

**Note:** The restart operation normally completes within a few seconds. It is possible, however, for it to take longer under some conditions (for example, if the Security database needs to run recovery or if the connectivity between hosts in a cluster is slow). If it takes longer than a few seconds for MarkLogic Server to restart, than the Admin Interface might return a `503: Service Unavailable` message. If you encounter this situation, wait several seconds and then reload the Admin Interface.

## 5.0 HTTP Servers

This chapter describes HTTP servers and provides procedures for configuring them. The following sections are included:

- [HTTP Server Overview](#)
- [Procedures for Creating and Managing HTTP Servers](#)

### 5.1 HTTP Server Overview

MarkLogic Server enables you to write XQuery-based web applications by connecting sets of XML content to HTTP servers that can access stored XQuery programs. These applications can return XHTML or XML content to a browser or other HTTP-enabled client application.

HTTP servers are defined at the group level and are accessible by all hosts within the group. Each HTTP server provides access to a set of XQuery programs that reside within a specified directory structure. Each host in the group must have access to the directory structure or mirror the directory structure along with the program files. An HTTP server executes the XQuery programs against the database to which it is connected.

HTTP servers follow the MarkLogic Server security model, as do WebDAV and XDBC servers. The server authenticates access to those programs using user IDs and passwords stored in the security database for that HTTP server. (Each HTTP server is connected to a database, and each database is in turn connected to a security database in which security objects such as users are stored.)

HTTP servers can execute XQuery code, either from a specified location on the file system or from a Modules database.

Granular access control to the system and to the data is achieved through the use of privileges and permissions. For details on configuring security objects in MarkLogic Server, see “Security Administration” on page 112. For conceptual information on the MarkLogic Server security model, see *Understanding and Using Security*.

### 5.2 Procedures for Creating and Managing HTTP Servers

Use the following procedures to create and manage HTTP servers:

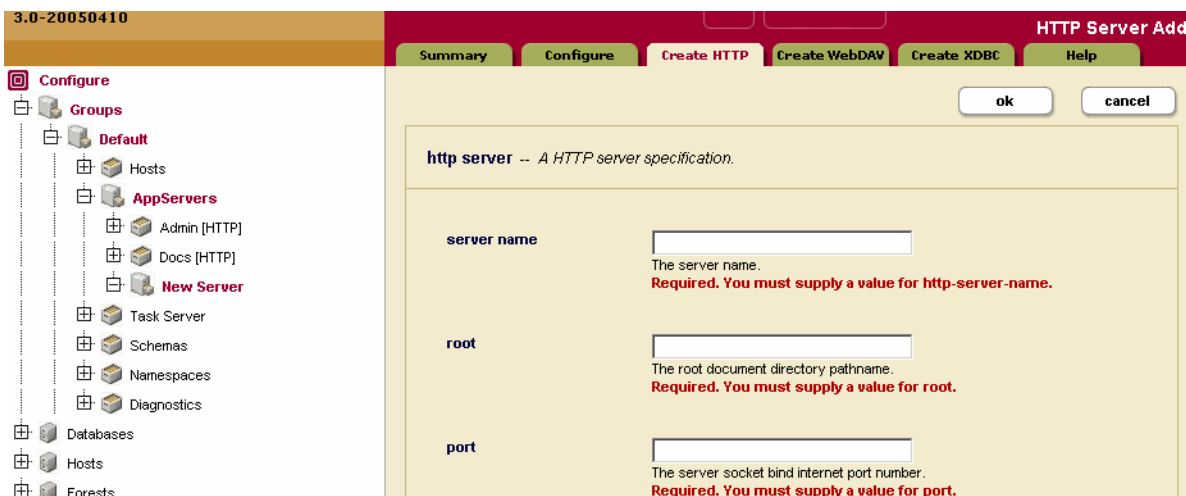
- [Creating a New HTTP Server](#)
- [Viewing HTTP Server Settings](#)
- [Deleting an HTTP Server](#)
- [Canceling a Request](#)

## 5.2.1 Creating a New HTTP Server

To create a new server, complete the following steps:

1. Click the Groups icon in the left frame.
2. Click the group in which you want to define the HTTP server (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Click the Create HTTP tab at the top right.

The HTTP Server Create page will display:



5. In the Server Name field, enter a shorthand name for this HTTP server.

MarkLogic Server will use this name to refer to this server on display screens and in user interface controls.

- In the Root directory field, enter the name of the directory in which you will store your XQuery programs.

**Note:** Unless you specify a shared drive, all hosts in the group will need to have a copy of the XQuery programs in the directory specified above.

The root directory is either a fully-qualified pathname or is relative to the directory in which MarkLogic Server is installed. The following table shows the default installation directory for each platform:

Platform	Program Directory
Microsoft Windows	C:\Program Files\MarkLogic
Red Hat Linux	/opt/MarkLogic
Sun Solaris	/opt/MARKlogic

**Warning** Do not create HTTP server root directories named Docs, Data or Admin. These directories are reserved by MarkLogic Server for other purposes. Creating HTTP server root directories with these names can result in unpredictable behavior of the server and may also complicate the software upgrade process.

- In the Port field, enter the port number through which you want to make this HTTP server available.

The port number must not be assigned to any other HTTP, XDBC or WebDAV server.

- Go to the Database field and select the database to be accessed by this HTTP server.

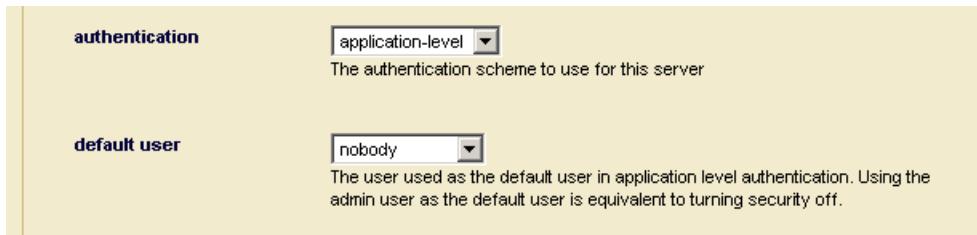
Multiple HTTP, XDBC, and WebDAV servers can access the same database.

- Scroll to the Authentication field. Select an authentication scheme: basic, digest, digestbasic, or application-level.

The screenshot shows a configuration panel with a light yellow background. It contains two sections:

- time limit:** A text input field containing the value '600'. Below it is the text: 'The maximum number of seconds allowed to service a request.'
- authentication:** A dropdown menu with 'basic' selected. Below it is the text: 'The authentication scheme to use for this server'

If you select application-level authentication, you will also need to fill in a Default User. Any one accessing the HTTP server is automatically logged in as the Default User until the user logs in explicitly. If you want anyone to have full access to the application, use an admin user (admin) as the Default User.



The screenshot shows a configuration panel with a light yellow background. It contains two sections: 'authentication' and 'default user'. The 'authentication' section has a dropdown menu set to 'application-level' with the text 'The authentication scheme to use for this server' below it. The 'default user' section has a dropdown menu set to 'nobody' with the text 'The user used as the default user in application level authentication. Using the admin user as the default user is equivalent to turning security off.' below it.

10. Scroll to the Privilege field near the bottom of the screen. This field represents the privilege needed to access (login to) the server. You may leave this field blank.

A user accessing the HTTP server must have the execute privilege selected in order to access the HTTP server. If you chose application-level authentication above, you should ensure that the default user has the selected privilege.



The screenshot shows a configuration panel with a light yellow background. It contains one section: 'privilege'. It has a dropdown menu that is currently blank with the text 'The privilege restricting access to the server.' below it.

11. Scroll to the top or bottom and click OK.

The HTTP server is now added. Adding an HTTP server is a “hot” admin task; the changes take effect immediately.

### 5.2.2 Viewing HTTP Server Settings

To view the settings for a particular HTTP server, complete the following steps:

1. Click the Groups icon in the left frame.
2. Click the group which contains the HTTP server you want to view (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Locate the HTTP server for which you want to view settings, either in the tree menu or on the summary page.
5. Click the icon for the HTTP server.
6. View the settings.

### 5.2.3 Deleting an HTTP Server

To delete the settings for an HTTP server, complete the following steps:

1. Click the Groups icon in the left frame.
2. Click the group which contains the HTTP server you want to delete (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Locate the HTTP server you want to delete, either in the tree menu or on the summary page.
5. Click the icon for the HTTP server.
6. Click Delete.
7. A confirmation message displays. Confirm the delete and click OK.

Deleting an HTTP server is a “cold” admin task; the server restarts to reflect your changes.

## 5.2.4 Canceling a Request

You can cancel a request in the App Server Status page of the Admin Interface (Groups > *group\_name* > App Servers > *app\_server\_name* > Status tab).

**App Server Status**

Summary | Configure | **Status** | Create HTTP | Create WebDAV | Create XDBC | Help

**App Server: myAppServer[HTTP]** show less

**appserver status** -- A detailed view of this appserver's activity.

**App Server** myAppServer [HTTP]  
**Database** apidoc  
**Hosts** raymond.marklogic.com

Host	Threads	Requests	Updates	Average Time	Request Rate	Oldest Request	Expanded Tree Cache		
							Hits	Misses	Ratio
raymond.marklogic.com	2	1	0	2.8 s	0.1	2.8 s	460224	34389	93%
	<b>2</b>	<b>1</b>	<b>0</b>	<b>2.8 s</b>	<b>0.1</b>	n/a	<b>460224</b>	<b>34389</b>	<b>93%</b>

Query	#	Average Time	Oldest Time	Expanded Tree Cache		
				Hits	Misses	Ratio
/cq-eval.xqy	1	2.8 s	2.8 s	0	0	n/a
<b>Total</b>	<b>1</b>	<b>2.8 s</b>	<b>2.8 s</b>	<b>0</b>	<b>0</b>	<b>n/a</b>

Host	Query	User	Client IP	Time	Expanded Tree Cache			
					Hits	Misses	Ratio	
raymond.marklogic.com	/cq-eval.xqy	admin	182.16.1.131	2.8 s	0	0	n/a	[cancel]
	<b>Total</b>				<b>0</b>	<b>0</b>	<b>n/a</b>	

To cancel a long-running request (for example, a long-running query statement or update statement), perform the following steps:

1. Click the Group menu item in the Admin Interface.
2. Navigate to the App Server in which the request was issued, either from the tree menu or from the summary page.
3. Click the Status tab.
4. Click the Show More button.
5. At the bottom right of the App Server Status page, click the cancel button on the row for the query you want to cancel.

6. Click OK on the Cancel Request confirmation page. If the request is already completed when the confirmation page occurs, the page will indicate that the request cannot be found.

The request is canceled and the App Server Status page appears again.

## 6.0 XDBC Servers

This chapter describes XDBC servers and provides procedures for configuring them. The following sections are included:

- [XDBC Server Overview](#)
- [Procedures for Creating and Managing XDBC Servers](#)

### 6.1 XDBC Server Overview

XDBC (XML Database Connector) servers are defined at the group level and are accessible by all hosts within the group. Each XDBC server provides access to a specific forest, and to a library of XQuery programs that reside within a specified directory structure. Applications execute by default against the database that is connected to the XDBC server.

XDBC Servers allow XML Contentbase Connector (XCC) applications to communicate with MarkLogic Server. XCC is an API used to communicate with MarkLogic Server from Java or .NET middleware applications. XDBC servers also allow old-style XDBC applications to communicate with MarkLogic Server, although XDBC applications cannot use certain 3.1 and newer features (such as point-in-time queries). Both XCC and XDBC applications use the same wire protocol.

XQuery requests submitted via XCC return results as specified by the XQuery code. These results can include XML and a variety of other datatypes. It is the XCC application's responsibility to parse, process and interpret these results in a manner appropriate to the variety of datatypes available. There are a number of publicly available libraries for assisting with this task, or you may write your own code. In order to accept connections from XCC-enabled applications, MarkLogic Server must be configured with an XDBC Server listening on the designated port. Each XDBC Server connects by default to a specific database within MarkLogic Server, but XCC provides the ability to communicate with any database in the MarkLogic Server cluster to which your application connects (and for which you have the necessary permissions and privileges).

XDBC servers follow the MarkLogic Server security model, as do HTTP and WebDAV servers. The server authenticates access to those programs using user IDs and passwords stored in the security database for that XDBC server. (Each XDBC server is connected to a database, and each database is in turn connected to a security database in which security objects such as users are stored.)

Granular access control to the system and to the data is achieved through the use of privileges and permissions. For details on configuring security objects in MarkLogic Server, see “Security Administration” on page 112. For conceptual information on the MarkLogic Server security model, see *Understanding and Using Security*.

## 6.2 Procedures for Creating and Managing XDBC Servers

Use the following procedures to create and manage XDBC servers:

- [Creating a New XDBC Server](#)
- [Viewing XDBC Server Settings](#)
- [Deleting an XDBC Server](#)

For the procedure to cancel a running request on an XDBC server, see “Canceling a Request” on page 27.

### 6.2.1 Creating a New XDBC Server

To create a new server, complete the following steps:

1. Click the Groups icon.
2. Click the group in which you want to define the XDBC server (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Click the Create XDBC tab at the top right.

The XDBC Server Create page displays.

The screenshot shows the 'XDBCServer Add' configuration page. The left-hand navigation tree is expanded to 'AppServers', showing sub-items like 'Admin [HTTP]', 'Docs [HTTP]', 'New XDBC Server', 'Task Server', 'Schemas', 'Namespaces', and 'Diagnostics'. The main configuration area has three fields: 'xdbc server name', 'library', and 'port'. Each field has a text input box, a description, and a red error message: 'Required. You must supply a value for xdbc-server-name.', 'Required. You must supply a value for library.', and 'Required. You must supply a value for port.' respectively. The page also has 'ok' and 'cancel' buttons at the top right.

5. In the XDBC Server Name field, enter a shorthand name for this XDBC server.

MarkLogic Server will use this name to refer to this server on display screens and in user interface controls.

- In the Library directory field, enter the name of the directory in which you will store any XQuery routines that can be referenced by submitted XQuery applications.

**Note:** Unless you specify a shared drive, all hosts in the group will need to have a copy of the XQuery programs in the directory specified above.

The root directory is either a fully-qualified pathname or is relative to the directory in which MarkLogic Server is installed. The following table shows the default installation directory for each platform:

Platform	Program Directory
Microsoft Windows	C:\Program Files\MarkLogic
Red Hat Linux	/opt/MarkLogic
Sun Solaris	/opt/MARKlogic

**Warning** Do not create XDBC server library directories named Docs, Data or Admin. These directories are reserved by MarkLogic Server for other purposes. Creating XDBC server library directories with these names can result in unpredictable behavior of the server and may also complicate the software upgrade process.

- In the Port field, enter the port number through which you want to make this XDBC server available.

The port number must not be assigned to any other XDBC or HTTP server.

- In the Database field, select the database to be accessed by this XDBC server.

Multiple HTTP, XDBC, and WebDAV servers can access the same database.

- Scroll to the Privilege field near the bottom of the screen. This field represents the privilege needed to access (login to) the server. You may leave this field blank.

A user accessing the XDBC server must have the execute privilege selected in order to access the XDBC server (or be a member of the Admin role).



10. Scroll to the top or bottom and click OK.

The new XDBC server is added. Adding a XDBC server is a “hot” admin task; the changes take effect immediately.

### 6.2.2 Viewing XDBC Server Settings

To view the settings for an XDBC server, complete the following steps:

1. Click the Groups icon.
2. Click the group which contains the XDBC server you want to view (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Locate the XDBC server for which you want to view settings, either in the tree menu or on the summary page.
5. Click the icon for the XDBC server.
6. View the settings.

### 6.2.3 Deleting an XDBC Server

To delete the settings for an XDBC server, complete the following steps:

1. Click on the Groups icon.
2. Click on the group which contains the XDBC server you want to delete (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Locate the XDBC server to be deleted, either in the tree menu or on the summary page.
5. Click the icon for this XDBC server.
6. Click Drop.
7. A confirmation message displays. Confirm the delete and click OK.

Deleting an XDBC server is a “cold” admin task; the server restarts to reflect your changes.

## 7.0 WebDAV Servers

A WebDAV server in MarkLogic Server is similar to an HTTP server, but has the following important differences:

- WebDAV servers cannot execute XQuery code.
- WebDAV servers support the WebDAV protocol to allow WebDAV clients to have read and write access (depending on the security configuration) to a database.
- A WebDAV server only accesses documents and directories in a database; it does not access the file system directly.

This chapter describes WebDAV servers in MarkLogic Server and includes the following sections:

- [WebDAV Server Overview](#)
- [Procedures for Creating and Managing WebDAV Servers](#)
- [WebDAV Clients](#)
- [Example: Setting Up a WebDAV Server to Add/Modify Documents Used By Another Server](#)

### 7.1 WebDAV Server Overview

WebDAV (Web-based Distributed Authoring and Versioning) is a protocol that extends the HTTP protocol to provide the ability to write documents through these HTTP extensions. You need a WebDAV client to write documents, but you can still read them through HTTP (through a web browser, for example). For information about WebDAV clients supported in MarkLogic Server, see “WebDAV Clients” on page 40. For general information about WebDAV and the WebDAV protocol, see the following web site:

<http://webdav.org>

This section provides an overview of WebDAV servers in MarkLogic Server, and includes the following topics:

- [Accesses a Database for Read and Write, Not XQuery Execution](#)
- [WebDAV Server Security](#)
- [Directories](#)
- [Server Root Directory](#)
- [Documents in a WebDAV Server](#)

### 7.1.1 Accesses a Database for Read and Write, Not XQuery Execution

In MarkLogic Server, WebDAV servers are defined at the group level and apply to all hosts within the group. Each WebDAV server provides access to a single database for reading and writing (dependent on the needed security permissions).

In the Admin console, you configure a WebDAV server to access a database. Documents stored in that database are accessible for reading via HTTP. The database is also accessible via WebDAV clients for reading, modifying, deleting, and adding documents. When you add a document via a WebDAV client (by dragging and dropping, for example), you are actually loading a document directly into the database.

When accessing a database via a WebDAV server, you cannot execute XQuery code. Unlike an HTTP server, there is no Modules database for a WebDAV server. You can, however, configure a database as the Modules database of an HTTP or XDBC server and you can configure the same database for access from a WebDAV server. Then, you can edit code from the WebDAV server that executes from an HTTP or XDBC server. For an example of this configuration, see “Example: Setting Up a WebDAV Server to Add/Modify Documents Used By Another Server” on page 42.

### 7.1.2 WebDAV Server Security

WebDAV servers follow the MarkLogic Server security model, as do HTTP and XDBC servers. The server authenticates users with user IDs and passwords stored in the security database for that WebDAV server, and the server controls access to objects in the database with privileges and roles. (Each WebDAV server is connected to a database, and each database is in turn connected to a security database in which security objects such as users are stored.)

You can configure application-level security if you want everyone who accesses the WebDAV server to effectively log in as the same user with no password. For example, if you want everyone to log in as *guest*, where *guest* has both read and write privileges and has a predefined set of default privileges, set the authentication scheme to application-level and set the default user to *guest*.

**Note:** Because users who have write permissions to the database on a WebDAV server can load documents into the database via a WebDAV client, be sure to configure appropriate default permissions on those users so that documents they load (for example, by dragging and dropping files into a WebDAV folder) have the needed permissions for other users to read and write, according to your security policy. You can achieve such granular access control to the system and to the data through the use of privileges and permissions. For information on using security features in MarkLogic Server, see “Security Administration” on page 112 and the chapters related to security in the *Developer’s Guide*.

### 7.1.3 Directories

A WebDAV directory is analogous to a file system directory. A directory must exist in order to view (via a WebDAV client) any documents in that directory (just like in a filesystem, where you must navigate to a directory in order to access any files in that directory). Each document in a directory has a URI that includes the directory URI as a prefix. Also, each directory visible from a WebDAV server must have the WebDAV root as its prefix, and there must exist a directory with the WebDAV root in the database.

For example, if you have a WebDAV root of `http://marklogic.com/`, then the URI of all documents and all directories must begin with that root in order to be visible from a WebDAV client. Also, the directory with a URI `http://marklogic.com/` must exist in the database. Therefore, a document with a URI of `http://marklogic.com/file.xml` is visible from this WebDAV server, and a directory with a URI of `http://marklogic.com/dir/` is also visible. A directory with a URI of `/dir/` and a document with a URI of `/dir/file.xml` is not visible from this server, however, because its URI does not begin with the WebDAV root.

The following sections describe further details about directories:

- [Automatic Directory Creation in a Database Settings](#)
- [Properties and URIs of Directories](#)

For more details on directories and properties, see the “Property Documents and Directories” chapter of the *Developer’s Guide*.

#### 7.1.3.1 Automatic Directory Creation in a Database Settings

In the configuration for a database in the Admin console, there is a directory creation setting. The directory creation setting specifies whether directories are created automatically when you create a document.

If you are using a WebDAV server to load documents into a database, we recommend you use the Admin Interface to set the directory creation setting for your database to `automatic`. If you create a WebDAV server that accesses a database with directory creation set to `automatic`, the root directory (required in order to access the database via a WebDAV client) is automatically created. Automatic directory creation also helps if you are loading documents manually (using the `xdmp:document-load` function, for example) whose URIs include directory hierarchies that do not exist in the database. Any directory implied by a URI is automatically created with directory creation set to `automatic`.

You can also manually create and delete directories in XQuery using the `xdmp:directory-create` and `xdmp:directory-delete` built-in functions.

For details on all of the directory creation settings, see “Basic Administrative Settings” on page 46.

### 7.1.3.2 Properties and URIs of Directories

A directory is stored as a properties document in a MarkLogic Server database. Like a document, a directory has a URI, but the URI must end in a forward slash (/). Use the `xdmp:document-properties("uri_name")` function to retrieve the properties document for a URI, or the `xdmp:document-properties()` function to retrieve all of the properties documents in the database.

Properties are in the `http://marklogic.com/xdmp/property` namespace. When you create a directory (either automatically or manually), the system creates a properties document in the database with a child element named `directory`. For example, if you have a directory in your database with a URI `/myCompany/marketing/`, the following query return the following results:

```
xdmp:document-properties("/myServer/Marketing/")
=>
<prop:properties xmlns:prop="http://marklogic.com/xdmp/property">
  <prop:directory/>
</prop:properties>
```

The properties document returned does not contain the URI of the directory, but just an empty element (`prop:directory`) indicating the existence of a directory.

The `xdmp:document-properties()` function returns the properties documents for all documents in the database. Whenever there is a directory element in the properties document, there is a directory in the database, and calling the XQuery `xdmp:node-uri` built-in function on that element returns the URI of the directory. For example, the following query returns the URIs for all of the directories in a database:

```
declare namespace prop="http://marklogic.com/xdmp/property"

for $x in xdmp:document-properties()/prop:properties/prop:directory
return <directory-uri>{xdmp:node-uri($x)}</directory-uri>
```

**Note:** It is possible to create a document with a URI that ends in a forward slash (/). To avoid confusion with directory URIs, the best practice is to avoid creating documents with URIs that end in a forward slash.

### 7.1.4 Server Root Directory

Each WebDAV server has a concept of a *root*. The root is the top-level directory accessible from the server; you can access any documents or directories in the database that are children of the root. The root therefore serves as a prefix for all document and directory URIs accessible through the WebDAV server. You enter the WebDAV root in the Admin console. The root can be any valid URI. The root should always end with a forward slash (/), and if it does not, the Admin console will append one to the string provided.

The root should be a unique string that can serve as the top of a directory structure. It is common practice to use a WebDAV root of the form `http://<company_domain>/`, but that is not required. The following are some examples of WebDAV roots:

```
http://myCompany/marketing/
```

```
/myCompany/marketing/
```

**Note:** Directories cannot end in two forward slashes (//). Therefore, you cannot create a directory with a URI `http://`. If you specify a root of `http://myCompany` for a WebDAV server and `directory creation` is set to `automatic` in the database, a directory with the URI `http://myCompany/` is automatically created in the database.

Whatever the root, any documents accessible through the WebDAV server must have URIs that begin with the root. Also, any documents created through a WebDAV client (for example, by dragging and dropping into a web folder) will be loaded with URIs beginning with the WebDAV root.

For example, a document with URI `/myCompany/marketing/strategy.doc` is accessible (given the necessary security permissions) via the WebDAV server with the root `/myCompany/marketing/`, and you can create that document by dragging a document named `strategy.doc` into a web folder configured to access the WebDAV server described above.

### 7.1.5 Documents in a WebDAV Server

The main purpose of a WebDAV server is to make it easy for people to store, retrieve, and modify documents in a database. The documents can be any type, whether they are text documents such as `.txt` files or source code, binary documents such as image files or Microsoft Word files, or XML documents. Because the documents are stored in a database, you can create applications that use the content in those documents for whatever purpose you need. You can also use the database backup and restore features to easily back up the content in the database.

## 7.2 Procedures for Creating and Managing WebDAV Servers

This section includes procedures to perform the following actions:

- [Creating a New WebDAV Server](#)
- [Viewing WebDAV Server Settings](#)
- [Deleting a WebDAV Server](#)

For the procedure to cancel a running request on a WebDAV server, see “Canceling a Request” on page 27.

### 7.2.1 Creating a New WebDAV Server

To create a new server, complete the following steps:

1. Click the Groups icon.
2. Click the group in which you want to define the WebDAV server (for example, Default).

3. Click the App Servers icon on the left tree menu.
4. Click the Create WebDAV tab at the top right.

The WebDAV Server Create page displays.

5. Go to the WebDAV Server Name field and enter a shorthand name for this WebDAV server.

MarkLogic Server will use this name to refer to this server on display screens and in user interface controls.

6. Go to the root field and enter the name of WebDAV root. This root is a string that represents the top-level of the WebDAV URI hierarchy. Any document accessible through this WebDAV server must have a URI that begins with this root string. For more details on the root, see “Server Root Directory” on page 36.

If the root directory does not contain a forward slash, the Admin Interface adds one for you.

7. Go to the Port field and enter the port number through which you want to make this WebDAV server available.

The port number must not be assigned to any other server.

8. Go to the Database field and select the database to be accessed by this WebDAV server.

Multiple HTTP, XDBC, and WebDAV servers can be connected to the same database.

**Note:** If you are using a database with a WebDAV server, the directory creation setting on the database should be set to `automatic`, which will automatically create the root directory and other directories for any documents added to the database (if the directory does not already exist). For more information on directories, see “Directories” on page 35.

9. Scroll to the Privilege field near the bottom of the screen. This field represents the privilege needed to access (login) the server. You may leave this field blank.
10. Scroll to the top or bottom and click OK.

The new WebDAV server is added. Adding a WebDAV server is a “hot” admin task.

## 7.2.2 Viewing WebDAV Server Settings

To view the settings for a WebDAV server, complete the following steps:

1. Click the Groups icon.
2. Click the group which contains the WebDAV server you want to view (for example, Default).
3. Click the App Servers icon on the left tree menu.
4. Locate the WebDAV server for which you want to view settings, either in the tree menu or on the summary page.
5. Click the icon for this WebDAV server.
6. View the settings.

## 7.2.3 Deleting a WebDAV Server

To delete the settings for a WebDAV server, complete the following steps:

1. Click the Groups icon.
2. Click the group which contains the WebDAV server you want to delete (for example, Default).
3. Click the WebDAVServers icon on the left tree menu.
4. Click the Configure tab at the top right.
5. Locate the WebDAV server to be deleted, either in the tree menu or on the summary page.
6. Click the icon for this WebDAV server.
7. Click Delete.
8. A confirmation message displays. Confirm the delete and click OK.

Deleting a WebDAV server is a “cold” admin task; the server restarts to reflect your changes.

## 7.3 WebDAV Clients

A WebDAV client allows you to log into a WebDAV server to read, modify, insert, add, or delete documents. This section lists the supported WebDAV clients for MarkLogic Server and provides some general and specific procedures. The following topics are included:

- [Tested WebDAV Clients](#)
- [General Steps to Connect to a Server](#)
- [Steps to Connect to a Web Folder in Windows Explorer](#)

### 7.3.1 Tested WebDAV Clients

The following table lists WebDAV clients that have been tested with MarkLogic Server:

WebDAV Client	How to Get It	Notes
Windows Explorer	Part of Windows 2000, Windows XP	Allows drag and drop from Windows. For instructions on setting up, see “Steps to Connect to a Web Folder in Windows Explorer” on page 41.
PerlDAV	<a href="http://www.webdav.org/perl原因/">http://www.webdav.org/perl原因/</a>	A command line, perl-based WebDAV client. Designed to be scriptable and to allow you to send individual WebDAV calls.
XML Spy	Altova Software ( <a href="http://www.altova.com/">http://www.altova.com/</a> )	Allows you to open, edit, and save XML files in XML Spy. Use the File > Open URL menu item in XML Spy.
jEdit DAV plug-in	Available on <a href="http://developer.marklogic.com">developer.marklogic.com</a>	Allows you to view and edit database documents in jEdit 4.2. This version is available from <a href="http://developer.marklogic.com">developer.marklogic.com</a> .

For detailed information on these clients, see the documentation accompanying these products.

**Note:** Directory and document names in WebDAV (and in MarkLogic Server databases) are case-sensitive, but some WebDAV clients (Windows Explorer, for example) are not case-sensitive. While Windows recognizes case, it treats the directory named `NewFolder` as the same directory as one named `newFolder`. Therefore, directory or document names that differ only in case might cause confusion when using Windows Explorer or other case-insensitive WebDAV clients. If possible, avoid assigning names to directories or documents that differ only by case ( for example, `NewFolder` VS `newFolder`).

### 7.3.2 General Steps to Connect to a Server

Each WebDAV client has its own way of connecting to a WebDAV server, but the general steps to connect to a WebDAV server are as follows:

1. Start the WebDAV client.
2. Enter the connection information for the WebDAV server. This includes the servername and port number of the WebDAV server. For example, if you have a WebDAV server running on port 9001 on a machine named `marklogic.myCompany.com`, enter the following URL in the appropriate place for your WebDAV client:

```
http://marklogic.myCompany.com:9001/
```

3. If prompted, enter a username and password for the WebDAV server. You will be prompted for a username or password unless you have configured application-level security.

**Note:** The user who logs into the WebDAV server must have the needed privileges (granted via roles) to access the documents and directories under the WebDAV root directory. Also, if you want the WebDAV user to create documents under the WebDAV root, then that user must have the needed URI privileges (granted via roles) to create documents under the root. The lack of any needed privileges and/or permissions can cause the WebDAV login or other WebDAV activities to fail. For details on URI privileges and document permissions, see *Understanding and Using Security*.

4. Use whatever browsing mechanism the client supports to add, remove, or modify documents and directories. For example, in Windows Explorer, double click on folders to expand them, drag and drop documents into folders, rename documents and directories, and so on.

### 7.3.3 Steps to Connect to a Web Folder in Windows Explorer

If you are running Windows, perform the following steps to use the Windows Explorer WebDAV client:

1. Double-click the My Network Places icon on your desktop.
2. In My Network Places, double-click the Add Network Places icon.
3. In the Add Network Place Wizard, enter your WebDAV server address and port number. For example, if you have a WebDAV server running on port 9001 on a machine named `marklogic.myCompany.com`, enter the following URL:

```
http://marklogic.myCompany.com:9001/
```

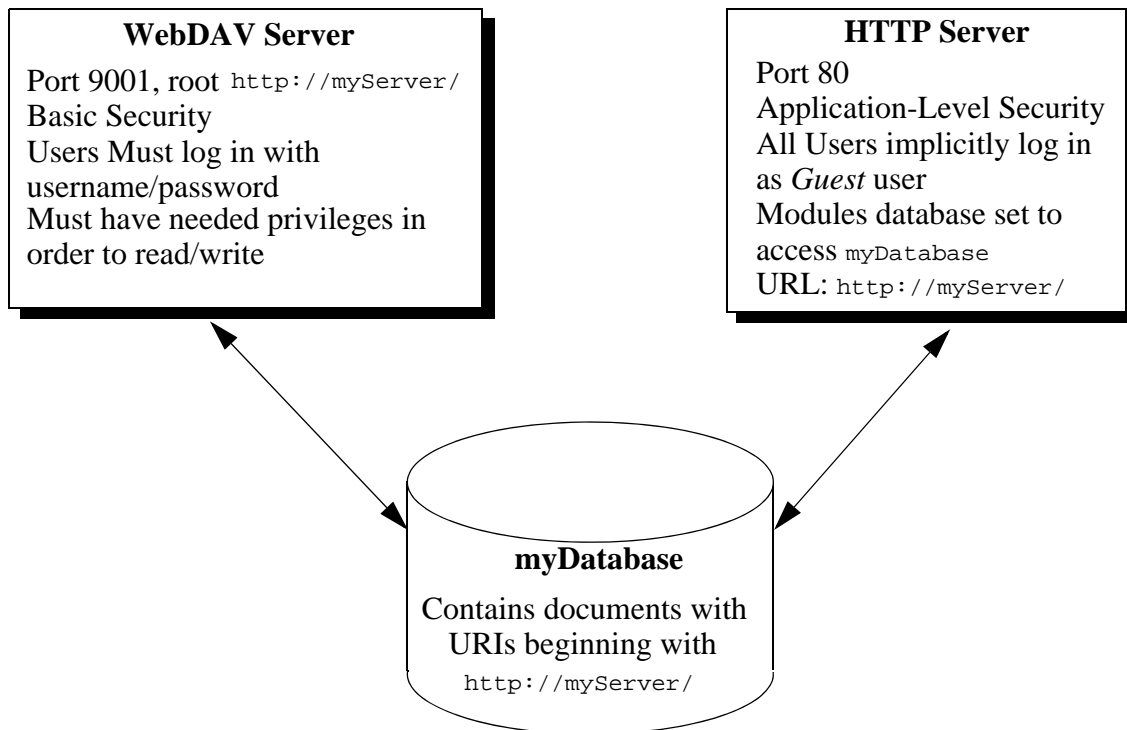
4. Click Next.

5. If prompted, enter your username and password for the WebDAV server.
6. Enter a name for the network place and click finish.

You can now use this folder like other Windows folders to drag and drop documents, rename documents, and so on. When you drag and drop a file into a WebDAV folder connected to a MarkLogic Server WebDAV server, you will actually load that document into the database.

#### 7.4 Example: Setting Up a WebDAV Server to Add/Modify Documents Used By Another Server

You can use a WebDAV server to provide privileged users write access to a database (via a WebDAV client). That database, in turn, might also be used as a Modules database in one or more other servers (HTTP, WebDAV, and/or XDBC) to provide read and execute access. Consider the scenario shown in the following figure:



In this scenario, all users can view the content by going to the URL `http://myServer/` in their web browsers. No password is needed to access this server because it is set up with application-level security, using a default user named *Guest*. The *Guest* user only has read permissions. If there is content that you do not want the *Guest* user to access, load that content with privileges that the *Guest* user does not have.

Meanwhile, users with the proper privileges can log in through a WebDAV client to access the WebDAV server at port 9001. Because the WebDAV server is configured with basic security, users are prompted for a username and password when they access the server through the WebDAV client (or through a web browser connected to port 9001). From the WebDAV client, they can add documents, edit documents, or read documents according to the database security policy.

For information about a Modules database, see “Modules Database” on page 45.

## 8.0 Databases

This section introduces basic forest management procedures. Later sections in this guide introduce some concepts for tuning the performance of your databases. For information on database backup and restore operations, see “Backing Up and Restoring a Database” on page 91. The following topics are included:

- [Understanding Databases](#)
  - [Schemas and Security Databases](#)
  - [Modules Database](#)
  - [Triggers Database](#)
  - [Database Settings](#)
  - [Example of Databases in MarkLogic Server](#)
- [Creating a New Database](#)
- [Attaching Forests to the Database](#)
- [Viewing Database Settings](#)
- [Loading Documents into a Database](#)
- [Detaching a Forest from a Database](#)
- [Deleting a Database](#)

### 8.1 Understanding Databases

A *database* in MarkLogic Server serves as a layer of abstraction between forests and HTTP, WebDAV, or XDBC servers. A database is made up of data *forests* that are configured on hosts within the same cluster but not necessarily in the same group. It enables a set of one or more forests to appear as a single contiguous set of content for query purposes. A forest is in turn made in-memory and on-disk structures called *stands*. Each stand is comprised of self-consistent XML, binary, and/or text fragments, stored in document order. When fragmentation rules are in place, XML documents may span multiple stands.

Multiple HTTP, XDBC, and WebDAV servers can be connected to the same database, allowing different applications to be deployed over a common content base. A database can also span forests that are configured on multiple hosts enabling data scalability through hardware expansion.

#### 8.1.1 Schemas and Security Databases

The installation process creates five databases by default - *Documents*, *Schemas*, *Security*, *Modules*, and *Triggers*. Every database points to a security database and a schema database. Security configuration information is stored in the security database and schemas are stored in the schemas database. A database can point back to itself for the security and schemas databases,

storing the security information and schemas in the same repository as the documents. However, security objects created through the Admin Interface are stored in the *Security* database by default. Mark Logic recommends leaving databases connected to *Security* as their security database.

### 8.1.2 Modules Database

The *modules* database is a database that is used to store executable XQuery code. During installation, a database named *Modules* is created, but any database can be used as a modules database, as long as the HTTP or XDBC server is configured to use it as a modules database. Also, it is possible to use the same database to store executable modules, to store queryable documents, and/or to store triggers.

If you use a modules database, each executable document in the database must have the root (specified in the HTTP or XDBC server) as a prefix to its URI. Also, the database should have `automatic` directory creation enabled, because all directories implied by a document URI must exist in order for the document to be executable. For information about directories and roots, see “Directories” on page 35 and “Server Root Directory” on page 36.

For example, if you are using a modules database and specify a root in an HTTP or XDBC server of `http://marklogic.com/`, the following documents are executable from that server:

```
http://marklogic.com/default.xqy
http://marklogic.com/myXQueryFiles/search_db.xqy
```

but the following files are not executable (because they do not have URIs that start with the root):

```
http://mycompany.com/default.xqy
/myXQueryFiles/search_db.xqy
```

In order to execute any documents in a modules database, the documents must be loaded with execute permissions. You can do this either by loading the documents as a user with default privileges that include execute permissions, or by setting those permissions on the document after it loads. For information on using permissions, privileges, and other security features in MarkLogic Server, see “Security Administration” on page 112 and the chapters related to security in the *Developer’s Guide*.

### 8.1.3 Triggers Database

The *triggers* database is a database that is used to store triggers. During installation, a database named *Triggers* is created, but any database can be used as a triggers database. Also, it is possible to use the same database to store executable modules, to store queryable documents, and/or to store triggers. A triggers database is required if you are using the Content Processing Framework. For details on the Content Processing Framework, see *Content Processing Framework*.

## 8.1.4 Database Settings

Each database has settings that control various aspects of a database such as memory allocation, indexing options, and so on. You configure these settings in the Admin Interface. You can configure the following basic types of settings for each database:

- [Basic Administrative Settings](#)
- [Index Settings that Affect Documents](#)
- [Reindexing Settings](#)
- [Document and Directory Settings](#)
- [Memory and Journal Settings](#)
- [Other Settings](#)
- [Merge Control Settings](#)

### 8.1.4.1 Basic Administrative Settings

The administrative settings configure properties such as the database name and which security and schema databases a database uses. These settings take effect immediately after any changes are made in the Admin Interface.

Database Setting	Description
database name	The name of the database.
security database	The name of the security database which this database accesses.
schema database	The name of the schemas database which this database accesses.
triggers database	The name of the triggers database which this database accesses.

### 8.1.4.2 Index Settings that Affect Documents

When you change any index settings for a database, the new settings take effect based on whether reindexing is enabled (`reindexer enabled`). For more details on text indexes, see “Text Indexing” on page 137.

In general, adding index options will have the effect of slowing document loading and increasing the size of database files.

Database Setting	Description
language	Specifies the default language for content in this database. Any content without an <code>xml:lang</code> attribute will be indexed in the language specified here. You should have a license key if you specify a non-English language; if you specify a non-english language and do not have a license for that language, the stemming and tokenization will be generic.
stemmed searches	Stemmed word searches enabled. Stemmed searches match not only the exact word in the search, but also words that come from the same stem and mean the same thing (for example, a search for <code>be</code> will also match the term <code>is</code> ). For more details on stemmed searches, see the chapter “Understanding and Using Stemmed Searches” in the <i>Developer’s Guide</i> .
word searches	Unstemmed word searches enabled. Enables searches for exact matches of words.
word positions	Index word positions for faster phrase and <code>cts:near-query</code> searches.
fast phrase searches	Speeds up phrase searches by eliminating some false positive results.
fast case sensitive searches	Speeds up case sensitive searches by eliminating some false positive results.
fast diacritic sensitive searches	Speeds up diacritic-sensitive searches by eliminating some false positive results.
fast element word searches	Speeds up element-word searches by eliminating some false positive results.
element word positions	Index element word positions for faster element-based phrase and <code>cts:near-query</code> searches.
fast element phrase searches	Speeds up element phrase searches by eliminating some false positive results.
element value positions	Index element word positions for faster element-based phrase and <code>cts:near-query</code> searches that use <code>cts:element-value-query</code> .
trailing wildcard searches	Faster wildcard searches with the wildcard at the end of the search pattern (for example, <code>abc*</code> ). For more details about wildcard searches, see the chapter “Understanding and Using Wildcard Searches” in the <i>Developer’s Guide</i> .

Database Setting	Description
trailing wildcard word positions	Index word positions for trailing wildcard searches.
fast element trailing wildcard searches	Faster wildcard searches with the wildcard at the end of the search pattern within a specific element, but slower document loads and larger database files.
three character searches	Enables wildcard searches where the search pattern contains three or more consecutive non-wildcard characters (for example, abc*x, *abc, a?bcd). For more details about wildcard searches, see the chapter “Understanding and Using Wildcard Searches” in the <i>Developer’s Guide</i> .
three character word positions	Index word positions for three-character wildcard queries.
two character searches	Enables wildcard searches where the search pattern contains two or more consecutive non-wildcard characters (for example, ab*). For more details about wildcard searches, see the chapter “Understanding and Using Wildcard Searches” in the <i>Developer’s Guide</i> .
word lexicon	Maintains a lexicon of all of the words in a database, with uniqueness determined by a specified collation. If you specify the collation <a href="http://marklogic.com/collation/">http://marklogic.com/collation/</a> , that specifies the UCA root collation, which is useful for many locales. The lexicon is capitalization-sensitive and diacritic-sensitive (therefore, there is a different entry for Ford and ford).
uri lexicon	Maintains a lexicon of all of the URIs used in a database. The URI lexicon speeds up queries that constrain on URIs. It is like a range index of all of the URIs in the database. To access values from the URI lexicon, use the <code>cts:uris</code> or <code>cts:uri-match</code> APIs.
collection lexicon	Maintains a lexicon of all of the collection URIs used in a database. The collection lexicon speeds up queries that constrain on collections. It is like a range index of all of the collection URIs in the database. To access values from the collection lexicon, use the <code>cts:collections</code> or <code>cts:collection-match</code> APIs.

### 8.1.4.3 Reindexing Settings

The reindexing settings enable or disable reindexing and allow you to force reindexing of older fragments.

Database Setting	Description
<code>reindexer enable</code>	When set to <code>true</code> , index configuration changes automatically initiate a background reindexing operation on the entire database. When set to <code>false</code> , any new index settings take effect for future documents loaded into the database; existing documents retain the old settings until they are reloaded or until you set <code>reindexer enabled</code> to <code>true</code> . For information on how the reindexer effects queries, see “Query Behavior with Reindex Settings Enabled and Disabled” on page 152.
<code>reindexer throttle</code>	Sets the priority of system resources devoted to reindexing. Higher numbers give reindexing a higher priority.
<code>reindexer timestamp</code>	Specifies the timestamp of fragments to force a reindex/refragment operation. Click the <code>get current timestamp</code> button to enter the current system timestamp. When you set this parameter to a timestamp and <code>reindex enable</code> is set to <code>true</code> , it causes a reindex and refragment operation on all fragments in the database that have a timestamp equal to or less than the specified timestamp. Note that if you restore a database that has a timestamp set, if there are fragments in the restored content that are older than the specified content, they will start to reindex as soon as they are restored.

### 8.1.4.4 Document and Directory Settings

The document and directory settings affect the default settings for how documents and directories are created in the database.

Database Setting	Description
directory creation	<p>Specifies if directories should be automatically created when a document is created. If you are using the database to store documents accessible via a WebDAV server or as a Modules database, this setting should be set to <code>automatic</code>. The following are the settings:</p> <ul style="list-style-type: none"> <li>• <code>automatic</code>—directories are automatically created based on the URI of a document.</li> <li>• <code>manual-enforced</code>—requires that the directory hierarchy corresponding to the URI exists before creating a document. If you create a document where the corresponding directory hierarchy does not exist, an error is raised. For example, if you try to create a document with the URI:               <pre>http://marklogic.com/file.xml</pre>               then the directory with URI <code>http://marklogic.com/</code> must exist. Otherwise, an error is raised. This setting provides the same behavior as a file system.</li> <li>• <code>manual</code>—directories are not automatically created, but documents can still be created without corresponding directories.</li> </ul> <p>For more information about directories, see “Directories” on page 35. For more information about Modules databases, see “Modules Database” on page 45.</p>
maintain last modified	<p>Creates and updates the last-modified property each time a document is created or updated. The default is <code>true</code>.</p>
maintain directory last modified	<p>Creates and updates the last-modified property on a directory each time a directory is created or updated. If set to <code>true</code>, update operations on documents in a directory will also update the directory last-modified timestamp, which can cause some contention when multiple documents in the directory are being updated. If your application is experiencing contention during these type of updates (for example, if you see deadlock-detected messages in the error log), set this property to <code>false</code>. The default is <code>false</code>.</p>
inherit permissions	<p>When set to <code>true</code>, documents and directories automatically inherit permissions from their parent directory (if permissions are not set explicitly when creating the document or directory). If there are any default permissions on the user who is creating the document or directory, those permissions are combined with any inherited permissions.</p>

Database Setting	Description
<code>inherit collections</code>	When set to <code>true</code> , documents and directories automatically inherit collection settings from their parent directory (if collections are not set explicitly when creating the document or directory). If there are any default collections on the user who is creating the document or directory, those permissions are combined with any inherited collections.
<code>inherit quality</code>	When set to <code>true</code> , documents and directories automatically inherit any quality settings from their parent directory (if quality is not set explicitly when creating the document or directory).

### 8.1.4.5 Memory and Journal Settings

The memory and journal settings are automatically configured at installation time. The memory settings configure the memory limits for the system, and the journal settings control the transactional journal, used for recovery if a database transaction fails. The default settings should be sufficient for most systems. Depending on the system workload, setting the memory settings incorrectly can adversely affect performance; if you need to change the settings, contact Mark Logic Support.

Database Setting	Description
<code>in memory limit</code>	The maximum number of fragments in an in-memory stand. An in-memory stand contains the latest version of any new or changed fragments. Periodically, in-memory stands are written to disk as a new stand in the forest. Also, if a stand accumulates a number of fragments beyond this limit, it is automatically saved to disk by a background thread.
<code>in memory list size</code>	The size, in megabytes, of the in-memory list storage.
<code>in memory tree size</code>	The size, in megabytes, of the in-memory tree storage. The <code>in memory tree size</code> should be at least 1 or 2 megabytes larger than the largest binary or text document you plan on loading into the database.
<code>in memory range index size</code>	The size, in megabytes, of the in-memory range index storage.

Database Setting	Description
journal size	<p>The size, in megabytes, of each journal file. The system uses journal files for recovery operations if a transaction fails to complete successfully. The default value should be sufficient for most systems; it is calculated at database configuration time based on the size of your system. If you change the other memory settings, however, the journal size should equal the sum of the <code>in memory list size</code> and the <code>in memory tree size</code>. Additionally, you should add space to the journal size if you use range indexes (particularly if you use a lot of range indexes or have extremely large range indexes), as range index data can take up journal space. Also, if your transactions span multiple forests, you may also need to add journal size, as each journal must keep the lock information for all of the documents in the transaction, not just for the documents that reside in the forest in which the journal exists.</p> <p>When you change the journal size, the next time the system creates a new journal, it will use the new size limit; existing journals will continue to use the old size limit until they are replaced with new ones (for example, when a journal fills up, when a forest is cleared, or when the system is cleanly shutdown and restarted).</p>
preallocate journals	<p>Set to <code>true</code> to preallocate journal file disk space, set to <code>false</code> to only allocate space as needed. Preallocating the disk space can help reduce fragmentation of the journal files (as long as the filesystem is not fragmented when you set this property). There are two journal files per forest.</p>
preload mapped data	<p>Specifies whether memory mapped data (for example, range indexes and word lexicons) is loaded into memory when a forest is mounted to the database. Preloading the memory mapped data improves query performance, but uses more memory, especially if you have a lot of range indexes and/or lexicons. Also, it will cause a lot of disk I/O at database startup time, slowing the system performance during the time the mapped data is read into memory. If you do not preload the mapped data, it will be paged into memory dynamically when a query requests data that needs it, slowing the query response time.</p>

### 8.1.4.6 Other Settings

The following are the remaining database configuration options.

Database Setting	Description
<code>position list max size</code>	<p>The maximum size, in megabytes, of the position list portion of the index for a given term. If the position list size for a given term grows larger than the limit specified, then the position information for that term is discarded. The default value is 128, the minimum value is 1, and the maximum value is 512. For example, position queries (<code>cts:near-query</code>) for frequently occurring words that have reached this limit (words like <i>a</i>, <i>an</i>, <i>the</i>, and so on) are resolved without using the indexes. Even though those types of words are resolved without using the indexes, this limit helps improve performance by making the indexes smaller and more efficient in relation to the content actually loaded in the database.</p>
<code>format compatibility</code>	<p>Specifies the version compatibility that MarkLogic Server applies to the indexes for this database during request evaluation. Setting this to a value other than <code>automatic</code> specifies that all forest data has the specified on-disk format, and it disables the automatic checking for index compatibility information. The automatic detection occurs during database startup and after any database configuration changes, and can take some time and system resources for very large forests and for very large clusters. The default value of <code>automatic</code> is recommended for most installations.</p>
<code>index detection</code>	<p>Specifies whether to auto-detect index compatibility between the content and the current database settings. This detection occurs during database startup and after any database configuration changes, and can take some time and system resources for very large forests and for very large clusters. Setting this to <code>none</code> also causes queries to use the current database index settings, even if some settings have not completed reindexing. The default value of <code>automatic</code> is recommended for most installations.</p>
<code>expunge locks</code>	<p>Specifies if MarkLogic Server will automatically expunge any lock fragments created using <code>xdmp:lock-acquire</code> with specified timeouts. If you set this to <code>none</code>, the lock fragments will remain in the database after the locks expire (although they will no longer be locking any documents) until they are explicitly removed with <code>xdmp:lock-release</code>. Setting this to <code>none</code> is only recommended to speed cluster startup time for extremely large clusters. The default setting of <code>automatic</code>, which cleans up the locks as they expire, is recommended for most installations.</p>

### 8.1.4.7 Merge Control Settings

The merge control settings allow you to control when merges occur, set merge parameters, and set up blackout periods where you do not want merges to occur. You can access the merge control settings by clicking the Admin Interface menu item for Database > *db\_name* > Merge Controls. Use caution when adjusting the merge parameters or disabling merges, as merges are necessary for optimal database performance. For explanations of the merge control settings and more details on controlling merges, see “Understanding and Controlling Database Merges” on page 80.

### 8.1.5 Example of Databases in MarkLogic Server

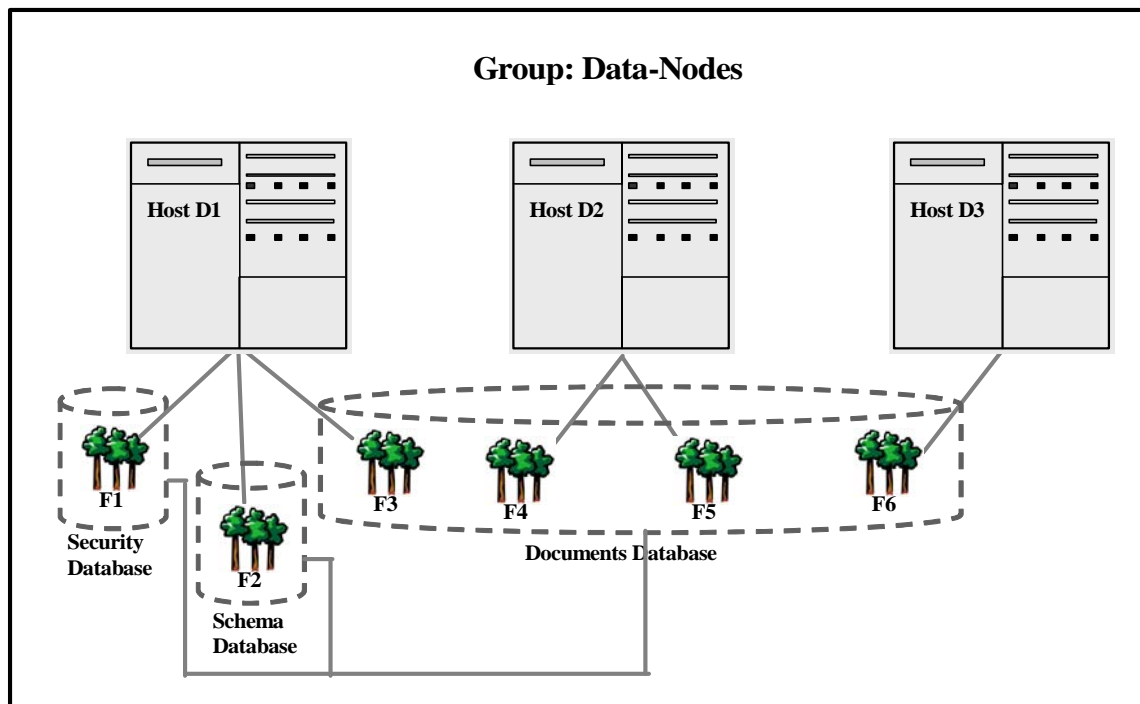
This section provides an example which demonstrates the concept of a database and the relationships between a database, a host and a forest in MarkLogic Server.

In the diagram below, Hosts D1, D2 and D3 belong to the Data-Nodes Group.

D1 is the first Host in Data-Nodes Group on which MarkLogic Server is loaded. Three **Databases** are created by default, **Security Database**, **Schema Database** and **Documents Database**. In the diagram below, 3 **Forests**, F1, F2 and F3 are configured on Host D1 and assigned to the **Security Database**, **Schema Database** and **Documents Database** respectively.

D2 is the second Host to join the Data-Nodes Group. **Forests** F4 and F5 are configured on D2 and attached to the **Documents Database**.

D3 is the third Host to join the Data-Nodes Group and has **Forest** F6, configured on it. F6 is also assigned to the **Documents Database**.



## 8.2 Creating a New Database

Follow the following steps to create a new database.

1. Click the Databases icon in the left tree menu.
2. Click the Create tab at the top right.

The Create Database page displays:

3. Enter the name of the database. This is the name the system will use to refer to this database.
4. Select a security database to be associated with this database. We recommend selecting *Security* as the security database.
5. Select a schema database to be associated with this database.
6. You may leave the rest of the parameters unchanged.
7. Click OK.

Your database is now created. You can now attach forests to the database. Creating a database is a “hot” admin task.

## 8.3 Attaching Forests to the Database

Forests can be moved from one database to another (detached from one and attached to another). However, to ensure correct query results, the two databases must have the same configuration. Otherwise, you should reload or reindex the forest data rather than simply attaching the forest.

Perform the following steps to attach a forest to a new database.

1. Click the database to which you want to attach forests.
2. Click the Forests icon.
3. Click the Attach tab.

The Attach Forest page displays.

4. Select from the list of available forests.
5. Click OK.

The forest is now attached to the database. Attaching a forest to a database is a “hot” admin task.

## 8.4 Viewing Database Settings

To view the settings for a particular database, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Locate the database for which you want to view settings, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to view the settings.
4. View the settings.
5. Click Forests, Triggers, Content Processing, Fragment Roots, Fragment Parents, Element-Word-Query-Throughs, Phrase-Throughs, Phrase-Arounds, Element Indexes and Attribute Indexes to view settings specific to those aspects of the database.


## 8.5 Loading Documents into a Database

You can use the Admin Interface to load documents into the database. The documents will be loaded with the default permissions and added to the default collections of the user with which you logged into the Admin Interface.

To load a set of documents into a database, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Click on the database into which you want to load the documents.

3. Click on the Load tab near the top right.



4. Enter the name of the directory in which the documents are located. This directory must be accessible by the host from which the Admin Interface is currently running.
5. Enter a filter for the names of the documents to be loaded (for example, \*.xml to load all files with an xml extension). For an exact match, enter the full name of the document.
6. Click OK to proceed.
7. The load confirmation screen will list all documents in the specified directory matching the specified filter. Click OK to complete the load.

The documents are loaded into the database.

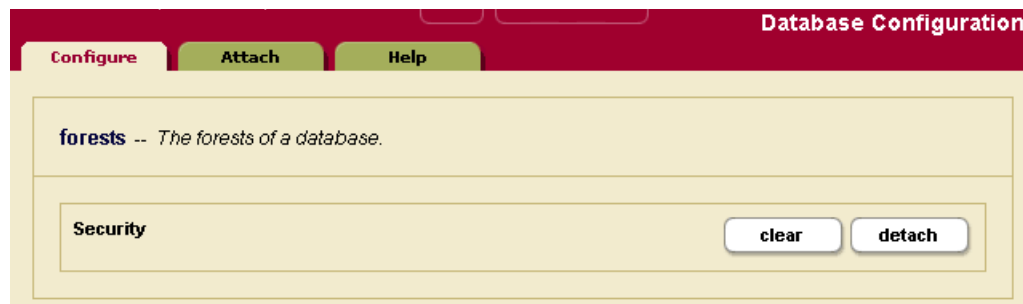
## 8.6 Detaching a Forest from a Database

Detaching a forest from a database does not delete the forest. The forest remains on the host on which it was created with the data intact. Forests can be moved from one database to another (detached from one and attached to another). However, before you attach the forest you are about to another database, you must ensure that the new database has the same configuration as the old database. Otherwise, you should reload or reindex the forest data rather than simply detaching the forest and attaching it to another database.

To detach a forest from a database, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Click the database from which you want to detach the forest.

3. Click Forests from the tree menu, underneath the database name. The database-forest configuration page displays:



4. Click the Detach button next to the forest you want to detach from the database.  
**Note:** If you click the Clear button, the forest is detached from the database *and* all of the data in the forest is deleted from disk. If you want to save the data in your forest, make sure to click the Detach button, not the Clear button.
5. Confirm that you want to detach the forest from the database. Click OK.

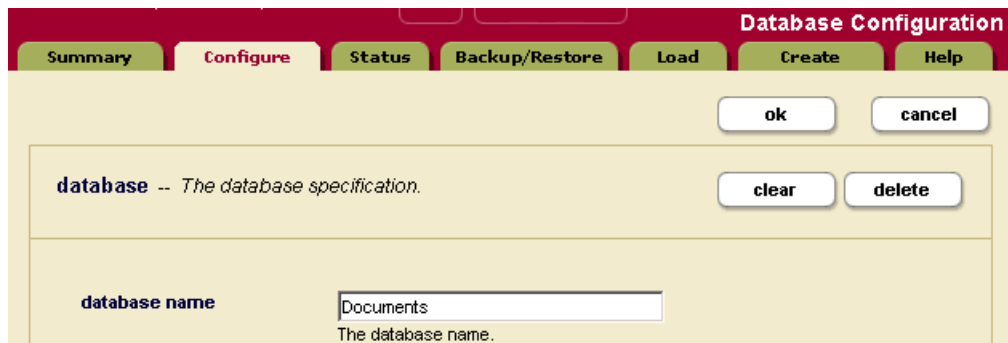
The forest is now detached from the database. Detaching a forest from a database is a “hot” task.

## 8.7 Deleting a Database

A database cannot be deleted if there are any HTTP, WebDAV, or XDBC servers that refer to the database. Deleting a database detaches the forests that are attached to it, but does not delete them. The forests remain on the hosts on which they were created with the data intact. Perform the following steps to delete a database:

1. Click the Databases icon on the left tree menu.
2. Locate the database which you want to delete, either in the tree menu or in the Database Summary table.

3. Click the name of the database which you want to delete.



4. Click on the Delete button near the top right.

**Note:** Clicking the Clear button clears all of the forests attached to this database, removing all of the data from the forests. Closing the Delete button removes the database configuration, but does not delete the data stored in the forests.

5. Assuming that there are not any HTTP, WebDAV, or XDBC servers referring to the database, a delete confirmation screen appears. Click OK.

The database is now permanently deleted. Deleting a database is a “hot” admin task.

## 9.0 Word Query Database Settings

This chapter describes how to configure a database to include or exclude elements, add index settings, and perform other configuration changes for `cts:word-query` operations. The following topics are included:

- [Understanding the Word Query Configuration](#)
- [Configuring Customized Word Query Settings](#)

### 9.1 Understanding the Word Query Configuration

Basic search of words and phrases in MarkLogic Server is based on the query constructor `cts:word-query`. You can control the behavior of these basic searches by changing the database configuration for word query. You can exclude and/or include elements from word queries, and you can add extra indexing options compared to the options configured in the database configuration. This section describes the options available in the word query configuration and includes the following parts:

- [Overview of Configuration Options](#)
- [Understanding Which Elements are Included and Excluded](#)
- [Adding a Weight to Boost or Lower the Relevance of an Included Element](#)
- [Specifying An Attribute Value for an Included Element](#)
- [Understanding the Index Option Configuration](#)

#### 9.1.1 Overview of Configuration Options

The following lists the main options you can set in the word query configuration to control how word queries are resolved in a database:

- By default, all elements are included in the word query configuration and the indexing options are the same as the database indexing options.
- All word query configurations are set on a per-database basis.
- The word query configuration controls the behavior of the `cts:word-query`, `cts:words`, and `cts:word-match` APIs. This includes controlling the words that get indexed as well as controlling the words that are returned from the filter (evaluator) portion of query evaluation.
- Word query inherits the database index settings as a starting point for its index settings.
- You can add extra index options for word query. These added index options will not affect other queries (for example, `cts:element-word-query`, `cts:element-attribute-word-query`).
- You cannot turn off indexing options that are enabled in the database settings.

- If you check index options in word query that are enabled in the database, it will not change any behavior. However, if you subsequently disable a database index setting that is checked in the word query settings, it will remain for the word query.
- You can include and/or exclude named elements from word queries.
- For any element you include, you can optionally constrain it by a value for a specified attribute.
- For any element you include, you can optionally specify a weight. The weight is used when determining relevance scores, where a weight greater than 1.0 will boost scores and a weight lower than 1.0 will lower scores for matches within the element.

### 9.1.2 Understanding Which Elements are Included and Excluded

You can include and/or exclude elements from word queries. This is useful if you know you will never want to search some element content. This section describes how MarkLogic Server determines what content is included in word queries and what is not when you include and/or exclude elements from the word query configuration.

**Note:** If you want to be able to search on everything in a word query, but also want a special view of the content that includes and/or excludes some elements, consider creating a field instead of modifying the word query configuration. For details on fields, see “Fields Database Settings” on page 68.

By default, all element content (all text node children of elements) is included in word queries. If you decide to include and/or exclude any elements from word queries, there are rules that govern which non-specified elements are indexed and which are not. The rules are based on inheriting the include state from the parent element. For example, if the parent element is marked as an included element (and is therefore indexed and evaluated for word query), then its children, if they do not appear on the exclude list, are also included.


When MarkLogic Server determines which elements to include/exclude, it walks the XML tree using the following rules:

1. Start at the root node of the document.
2. If the root node is included (either because it is explicitly included or because `include document root` is set to true), MarkLogic Server includes the immediate text node children of the document root element and then moves to its element children. If the root node is excluded, the text nodes are not included and MarkLogic Server moves down the XML tree to its element children.
3. If the parent element (the root element in this case) was included, MarkLogic Server keeps walking down the tree and including the text node children until it encounters an explicitly excluded element.

4. If the parent element (the root element in this case) was not included, MarkLogic Server keeps walking down the tree, not including the text node children, until it encounters an explicitly included element.
5. MarkLogic Server keeps walking down the tree, including or not according to the state inherited from the parent element, until it encounters the next included element (if it is in the *not included* state) or excluded element (if it is in the *included* state).
6. During this process, when an element is encountered that is neither included nor excluded, it inherits the included state (*not included* or *included*) from the parent element.
7. MarkLogic Server keeps walking down the XML tree using this logic to determine its included state, until it reaches the end of the document.

The only way to guarantee an element’s text node children will be included (assuming you have any elements included and/or excluded) is to add it to the included list, and the only way to guarantee an element is not included is to add it to the excluded list.

The following figure shows what is included for two configurations, one with the root node included and one with the root node excluded. Note that the includes and excludes are the same. The lines below the element names represent the text nodes, and the yes/no indicates whether the content in the text nodes is included in word queries. The `root` represents the root node of an XML structure, with elements `A` and `B` included and elements `C` and `D` excluded. Elements that are not explicitly included or excluded (for example, `E`, `F`, and `Z`) inherit from their parents.



The lines indicate text nodes, Yes is included, No is excluded

Notice that the `z` node, which is not explicitly included or excluded, sometimes is included and sometimes is not included, depending on the include state of its parent element.

### 9.1.3 Adding a Weight to Boost or Lower the Relevance of an Included Element

When you include an element, one of the options is to add a `weight` to the included element specification. When you add a weight, all text in this element (including any text in all text node descendants of the element) are weighted by the specified value, changing the relevance at query time. Specifying a weight greater than 1.0 will boost scores and a weight lower than 1.0 will lower scores for matches within the element.

When you specify a weight, the term frequency for any tokens in that element (including tokens in descendant text nodes) is multiplied by that number. This happens during document load, update, or reindexing. For example, if you specify a weight of 2.0, each term will have a term frequency of 2.0, making it as if each term appeared twice (for score calculation purposes). Similarly, if you specify a weight of 0.5, each term will have a term frequency of 0.5.

Adding a weight is useful to boost or lower scores on searches where the match occurs in a given element. For example, if you want matches in `TITLE` elements to contribute more towards the relevancy score than matches in other elements, you can specify a weight of 2.0 for the `TITLE` element. Conversely, if you want matches in `TITLE` elements to contribute less to the relevancy score than matches in other elements, you can specify a weight of 0.5 for the `TITLE` element. For details on how relevance is calculated, see the chapter [Composing cts:query Expressions](#) in the *Developer's Guide*.

### 9.1.4 Specifying An Attribute Value for an Included Element

When you include an element, one of the options is to specify an attribute value. This option allows you to only include elements with a particular attribute/value pair. The attribute/value pair acts as a predicate on which to constrain the content. For example, consider the following XML snippet:

```
<chapter class="history">some text here</chapter>
<chapter class="mathematics">some more text here</chapter>
<chapter class="english">some other text here</chapter>
<chapter class="history">some different text here</chapter>
<chapter class="french">other text here</chapter>
<chapter class="linguistics">still other text here</chapter>
```

For the element `chapter`, if you specify the attribute/value pair of `class` and `history`, then only the following elements will be included:

```
<chapter class="history">some text here</chapter>
<chapter class="history">some different text here</chapter>
```

You can only specify an attribute value for an included element; you cannot specify one for an excluded element.

### 9.1.5 Understanding the Index Option Configuration

The word query configuration allows you to add some extra indexing options from the ones that are currently set in the database configuration. Adding any index options to the word query configuration does not add those options to the element-based index options.

To add a particular index option to word query, you check the box corresponding to the index option. Adding any index options that are not enabled in the database configuration will cause new and updated documents to use the new indexing for word query, and will trigger a reindex operation if `reindex enable` is set to true in the database configuration.

Options that are enabled in the database configuration appear in bold on the word query configuration. If you check the box next to an option with bold-face type, it does not change your configuration. However, if you subsequently disable that index option in the database configuration, it will remain enabled for word query as long as the box is checked.

## 9.2 Configuring Customized Word Query Settings

This section provides the procedure for customizing the word query settings. For details on what the meaning of the various configuration options in fields, see “Understanding the Word Query Configuration” on page 60. The following is the procedure for modifying the word query configuration for your database:

**Note:** When you modify the word query settings, those modifications apply to all queries that use the `cts:word-query` constructor, which is the default constructor for `cts:search`. If you want to be able to search on everything in a word query, but also want a special view of the content that includes and/or excludes some elements, consider creating a field instead of modifying the word query configuration. For details on fields, see “Fields Database Settings” on page 68.

Use the Admin Interface to perform the following steps to add a new field configuration to a database.

1. Access the Admin Interface in a browser.
2. Navigate to and click the database for which you want to modify the word query configuration, either from one of the summary tables or in the left tree menu.
3. Under the database in which you want to create the field, click the Word Query link. The Word Query Configuration page appears.
4. If you want the word queries to include any extra index options from the database, check those index settings. Index settings shown in bold indicate the setting is inherited from the database setting. For details, see “Understanding the Index Option Configuration” on page 64.

5. If you want the word queries to include the root element of the document, even if it is not explicitly included, leave the default of `true` for include document root button. Note that if you set this to `false`, you will need to include elements in the word query configuration in order to get any results from word queries. Typically, you would leave this set to true and choose some elements to explicitly exclude and some to explicitly include (optionally adding a scoring weight and/or an attribute value constraint).
6. Click OK to save any changes you made. The configuration page refreshes with after the changes have been made to the MarkLogic Server configuration.
7. If you want to exclude any elements from word queries, click the Excludes tab.
8. Enter the namespace URI (if needed) and the localname for the excluded element.

**Add Word Query Exclude**

Configure Includes **Excludes** Help

ok cancel

**excluded element** -- *The element included in word query.*

**namespace uri**   
A namespace URI.

**localname**   
The localname of the excluded element.  
**Required. You must supply a value for localname.**

ok cancel

9. Click OK.
10. Repeat steps [7](#) through [9](#) for each element you want to exclude.

11. Click the Includes tab to specify elements to include in the word query.

**included element** -- *The element included in word query.*

**namespace uri**   
A namespace URI.

**localname**   
The localname of the included element.  
**Required. You must supply a value for localname.**

**weight**   
The weight, used to boost or lower relevance scores, of the included element.

**attribute namespace uri**   
Namespace of the child attribute.

**attribute localname**   
Localname of the child attribute.

**attribute value**   
Include only elements with the specified attribute having this value.

12. On the Included Element page, specify a localname for the element to include. If the element is in a namespace, specify the namespace URI for the element to include.
13. [OPTIONAL] If you want to boost or lower the relevance contribution for matches within this element, specify a weight other than the default of 1.0. Weights greater than 1.0 will boost the relevance contribution and weights lower than 1.0 will lower the contribution.
14. [OPTIONAL] If you want to only include elements that have an attribute with a specified value, enter the attribute namespace uri (if needed), the attribute localname, and a value for the attribute. Then only elements containing attributes with the specified value will be included. You must specify the exact value; no wildcard characters are used.
15. When you have specified everything for this element, click OK.
16. Repeat steps [11](#) through [15](#) for each element you want to include.

17. You can delete any included or excluded fields from the tables at the bottom of the field configuration page.



The screenshot displays a configuration window with two tables. The top table, titled "Included Elements", has columns for Localname, Namespace, Attribute, Attribute Namespace, Value, and Weight. It contains one entry: "ABSTRACT" with a weight of 2.0 and a "[delete]" link. The bottom table, titled "Excluded Elements", has columns for Localname and Namespace. It contains one entry: "script" with the namespace "http://www.w3.org/1999/xhtml" and a "[delete]" link. At the bottom of the window are "ok" and "cancel" buttons.

Included Elements					
Localname	Namespace	Attribute	Attribute Namespace	Value	Weight
ABSTRACT					2.0

Excluded Elements	
Localname	Namespace
script	http://www.w3.org/1999/xhtml

## 10.0 Fields Database Settings

This chapter describes how to configure fields in the database settings. Fields are used with the `cts:field-word-query`, `cts:field-words`, and `cts:field-word-match` APIs, and allow you to define a named field consisting of several elements over which you can search. The following topics are included:

- [Overview of Fields](#)
- [Understanding Field Configurations](#)
- [Field Word Lexicons](#)
- [Configuring Fields](#)

### 10.1 Overview of Fields

Fields provide a convenient mechanism for querying a portion of the database based on element QNames. Unlike collections or directories, which allow you to query portions of a database based on document URIs, fields allow you to query portions of a database based on elements. This offers extra convenience for the application developers, and also offers performance boosts over other methods of querying a portion of the database. Fields are extremely useful when you have content in one or more elements that you want to query simply and efficiently as a single unit.

Field query is similar to word query (in its default configuration, with everything included), but instead of querying everything in the database, fields query only what is configured for the specified field. Fields have their own set of indexes, independent of the database indexes. Because fields have their own indexes, and a field is typically a small subset of the whole database, querying a field is often more efficient than querying those same elements directly (with `cts:word-query`, for example).

Also, because fields have their own sets of indexes, relevance for fields is calculated based on the content in the field, not based on all of the content in the database. This provides finer-grain relevance for field searches than for other searches.

You can use fields to create portions of the content that you might want to query as a single unit. Additionally, you can configure a field with indexing options over and above the ones configured in the database. For example, consider a database containing many technical articles, each article containing an brief abstract. You might want to build an application that allows greater capabilities for searching through the abstracts than for searching through the rest of the articles. Assume your main content does not have wildcard indexes, but you want to be able to search through the abstracts using wildcard searches. You can create a field on the abstract, and then add wildcard indexes to that field. Because the field represents only a relatively small percentage of the content, the relative cost of the extra indexing is small.

## 10.2 Understanding Field Configurations

Field search of words and phrases in MarkLogic Server is based on the query constructor `cts:field-word-query`. You can control the behavior of these field searches by changing the database configuration for the field you query. You can exclude and/or include elements from fields, and you can add extra indexing options for some elements. This section describes the options available in the configuration and includes the following parts:

- [Overview of Field Configuration Options](#)
- [Understanding Which Elements are Included and Excluded](#)
- [Adding a Weight to Boost or Lower the Relevance of an Included Element](#)
- [Specifying An Attribute Value for an Included Element](#)
- [Understanding the Index Option Configuration](#)

### 10.2.1 Overview of Field Configuration Options

The following lists the main options you can set in the field query configuration to control how queries against the specified field are resolved:

- By default, no elements are included in the field query configuration and the indexing options are the same as the database indexing options. You must specify at least one element to include for the field to include anything.
- All field configurations are set on a per-database basis.
- The field configuration controls the behavior of the `cts:field-word-query`, `cts:field-words`, and `cts:field-word-match` APIs. This includes controlling the words that get indexed as well as controlling the words that are returned from the filter (evaluator) portion of query evaluation.
- Fields inherit the database index settings as a starting point for its index settings.
- You can add extra index options for each field. These added index options will not affect other queries (for example, `cts:word-query`, `cts:element-word-query`, `cts:element-attribute-word-query`).
- You cannot turn off indexing options that are enabled in the database settings.
- If you check index options in a field that are enabled in the database, it will not change any behavior. However, if you subsequently disable a database index setting that is checked in the field setting, it will remain for the field.
- You can include and/or exclude named elements from each field.
- For any element you include, you can optionally constrain it by a value for a specified attribute.
- For any element you include, you can optionally specify a weight. The weight is used when determining relevance scores, where a weight greater than 1.0 will boost scores and a weight lower than 1.0 will lower scores for matches within the element.

- Each field has its own set of indexes; it does not share the indexes with the word query indexes. Therefore, if you have a field with fewer elements than word query, there is a smaller amount of content to index and fewer I/O operations are needed to resolve the query from the indexes (index resolution phase of query processing).

## 10.2.2 Understanding Which Elements are Included and Excluded

You can include and/or exclude elements from a field. This is useful if you know you will never want to search some element content. This section describes how MarkLogic Server determines what content is included in the field and what is not when you include and/or exclude elements from the field configuration.

By default, no element content (all text node children of elements) is included in a field. When you include and/or exclude any elements from a field, there are rules that govern which non-specified elements are indexed and which are not. The rules are based on inheriting the include state from the parent element. For example, if the parent element is marked as an included element (and is therefore indexed and evaluated for field-based queries), then its children, if they do not appear on the exclude list, are also included.

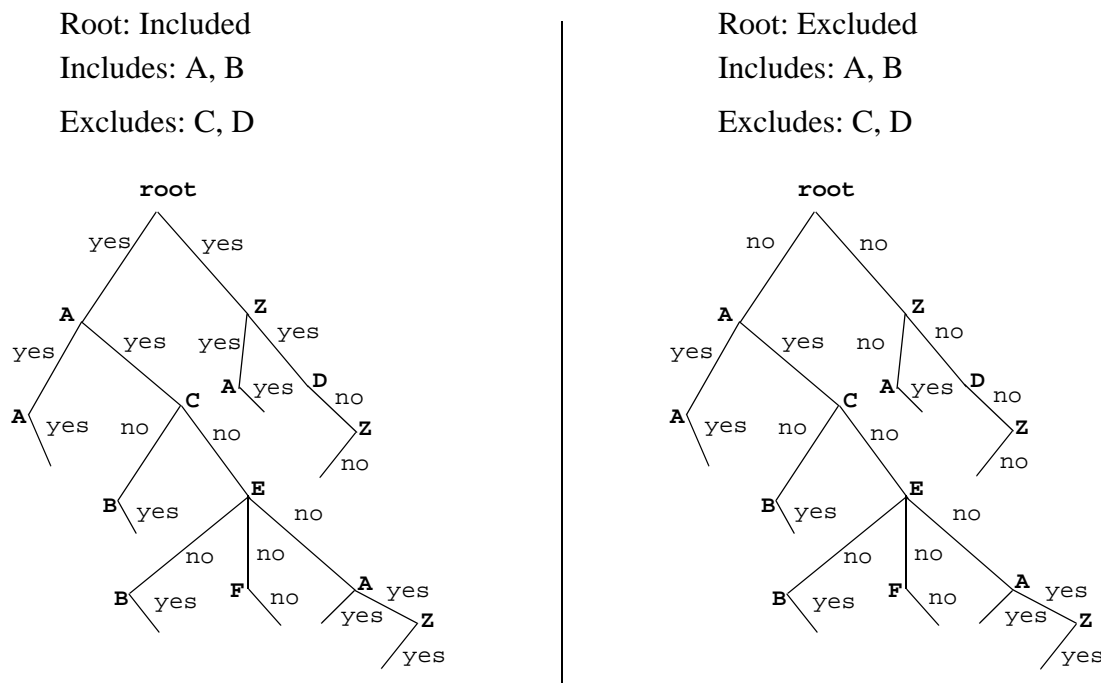
When MarkLogic Server determines which elements to include/exclude, it walks the XML tree using the following rules (note that these are the same rules used for including/excluding elements in the word query configuration):

1. Start at the root node of the document.
2. If the root node is included (either because it is explicitly included or because `include document root` is set to true), MarkLogic Server includes the immediate text node children of the document root element and then moves to its element children. If it is excluded, the text nodes are not included and MarkLogic Server moves down the XML tree to its element children.
3. If the parent element (the root element in this case) was included, MarkLogic Server keeps walking down the tree and including the text node children until it encounters an explicitly excluded element.
4. If the parent element (the root element in this case) was not included, MarkLogic Server keeps walking down the tree, not including the text node children, until it encounters an explicitly included element.
5. MarkLogic Server keeps walking down the tree, including or not according to the state inherited from the parent element, until it encounters the next included element (if it is in the *not included* state) or excluded element (if it is in the *included* state).
6. During this process, when an element is encountered that is neither included nor excluded, it inherits the included state (*not included* or *included*) from the parent element.

7. MarkLogic Server keeps walking down the XML tree using this logic to determine its included state, until it reaches the end of the document.

The only way to guarantee an element's text node children will be included (assuming you have any elements included and/or excluded) is to add it to the included list, and the only way to guarantee an element is not included is to add it to the excluded list.

The following figure shows what is included for two configurations, one with the root node included and one with the root node excluded. Note that the includes and excludes are the same. The lines below the element names represent the text nodes, and the yes/no indicates whether the content in the text nodes is included in word queries. The `root` represents the root node of an XML structure, with elements `A` and `B` included and elements `C` and `D` excluded. Elements that are not explicitly included or excluded (for example, `E`, `F`, and `Z`) inherit from their parents.



The lines indicate text nodes, Yes is included, No is excluded

Notice that the `z` node, which is not explicitly included or excluded, sometimes is included and sometimes is not included, depending on the include state of its parent element.

### 10.2.3 Adding a Weight to Boost or Lower the Relevance of an Included Element

When you include an element, one of the options is to add a `weight` to the included element specification. When you add a weight, all text in this element (including any text in all text node descendants of the element) are weighted by the specified value, changing the relevance at query time. Specifying a weight greater than 1.0 will boost scores and a weight lower than 1.0 will lower scores for matches within the element.

When you specify a weight, the term frequency for any tokens in that element (including tokens in descendant text nodes) is multiplied by that number. This happens during document load, update, or reindexing. For example, if you specify a weight of 2.0, each term will have a term frequency of 2.0, making it as if each term appeared twice (for score calculation purposes). Similarly, if you specify a weight of 0.5, each term will have a term frequency of 0.5.

Adding a weight is useful to boost or lower scores on searches where the match occurs in a given element. For example, if you want matches in `TITLE` elements to contribute more towards the relevancy score than matches in other elements, you can specify a weight of 2.0 for the `TITLE` element. Conversely, if you want matches in `TITLE` elements to contribute less to the relevancy score than matches in other elements, you can specify a weight of 0.5 for the `TITLE` element. For details on how relevance is calculated, see the chapter [Composing cts:query Expressions](#) in the *Developer's Guide*.

### 10.2.4 Specifying An Attribute Value for an Included Element

When you include an element, one of the options is to specify an attribute value. This option allows you to only include elements with a particular attribute/value pair. The attribute/value pair acts as a predicate on which to constrain the content. For example, consider the following XML snippet:

```
<chapter class="history">some text here</chapter>
<chapter class="mathematics">some more text here</chapter>
<chapter class="english">some other text here</chapter>
<chapter class="history">some different text here</chapter>
<chapter class="french">other text here</chapter>
<chapter class="linguistics">still other text here</chapter>
```

For the element `chapter`, if you specify the attribute/value pair of `class` and `history`, then only the following elements will be included:

```
<chapter class="history">some text here</chapter>
<chapter class="history">some different text here</chapter>
```

You can only specify an attribute value for an included element; you cannot specify one for an excluded element.

### 10.2.5 Understanding the Index Option Configuration

The field configuration allows you to add some extra indexing options from the ones that are currently set in the database configuration. Adding any index options to the field configuration does not add those options to the element-based index options.

To add a particular index option to a field, you check the box corresponding to the index option. Adding any index options that are not enabled in the database configuration will cause new and updated documents to use the new indexing for the field, and will trigger a reindex operation if `reindex enable` is set to true in the database configuration.

Options that are enabled in the database configuration appear in bold on the field configuration. If you check the box next to an option with bold-face type, it does not change your configuration. However, if you subsequently disable that index option in the database configuration, it will remain enabled for word query as long as the box is checked.

### 10.3 Field Word Lexicons

As with word lexicons, you can create a word lexicons for each field. A *field word lexicon* is a list of all of the unique words in the database that occur in the field. The list is ordered in the specified collation. You can create multiple field lexicons on the same field with different collations. The field word lexicons are accessed with the `cts:field-words` and `cts:field-word-match` APIs. For details about lexicons, see [Browsing With Lexicons](#) in the *Developer's Guide*.

### 10.4 Configuring Fields

This section provides procedures to create and modify field configurations in a database. For details on what the meaning of the various configuration options in fields, see “Understanding Field Configurations” on page 69. This section includes the following procedures:

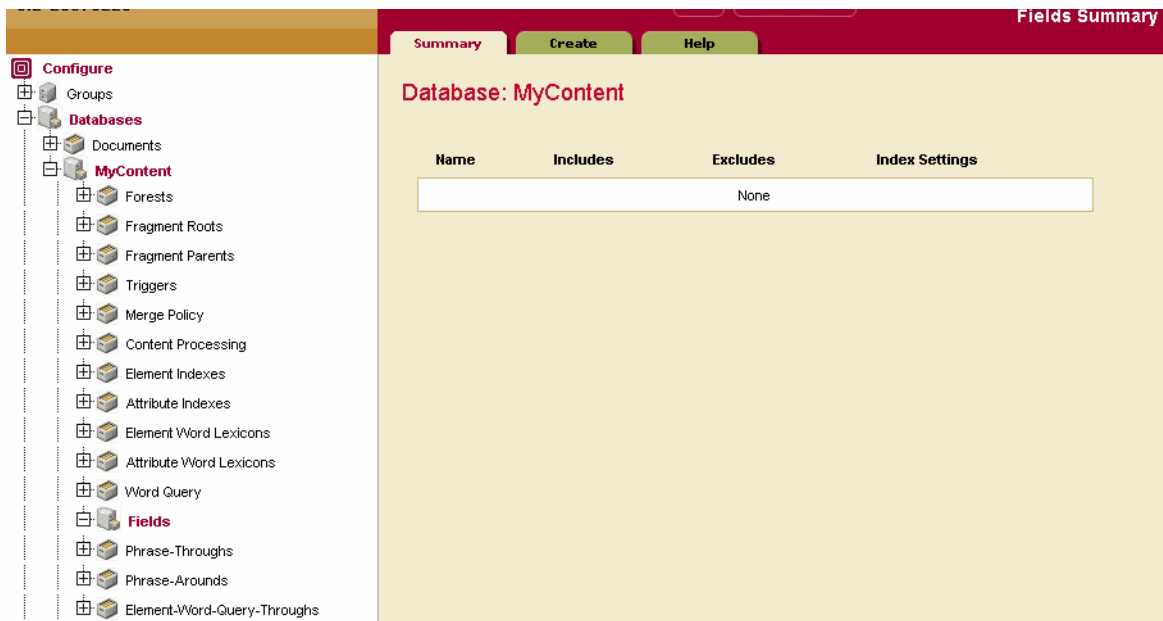
- [Configuring a New Field](#)
- [Modifying an Existing Field](#)

#### 10.4.1 Configuring a New Field

Use the Admin Interface to perform the following steps to add a new field configuration to a database.

1. Navigate to and click the database for which you want to create a field, either from one of the summary tables or in the left tree menu.

- Under the database in which you want to create the field, click the Fields link. The Field Summary page appears.



3. Click the Create tab. The Create Field in Database page appears.

**Database Fields Configuration**

Summary Create Help

**Create Field in Database**

**field name**   
The field name.  
**Required. You must supply a value for field-name.**

**index settings**

**stemmed searches:** basic

word searches

**fast phrase searches**

**fast case sensitive searches**

**fast diacritic sensitive searches**

trailing wildcard searches

trailing wildcard word positions

three character searches

three character word positions

two character searches

one character searches

*Options in bold inherited from database config*

**include document root**  true  false  
Includes elements starting at the document root

ok cancel

4. Enter a name for the field.
5. If you want the field to include any extra index options from the database, check those index settings. Index settings shown in bold indicate the setting is inherited from the database setting. For details, see “Understanding the Index Option Configuration” on page 72.
6. If you want the field to include the root element of the document, even if it is not explicitly included, click the `true` button for include document root. Typically, you leave this set to the default of `false`, unless your field will include most of the elements in the database.

7. Click OK. The configuration page with the field appears, adding the following parts to the bottom of the configuration page:

**word lexicons**      **[Keep] Collation URI**

**include document root**       true     false  
 Includes elements starting at the document root

**Included Elements**

Localname	Namespace	Attribute	Attribute Namespace	Value	Weight
None					

**Excluded Elements**

Localname	Namespace
None	

8. If you want to add a word lexicon for the field, enter the collation URI next in the add text box. The URI for the UCA Default Collation, <http://marklogic.com/collation/>, is useful for many applications. For details on collations, see the [Language Support in MarkLogic Server](#) chapter in the *Developer's Guide*. Click the OK button to add the field word lexicon (if you want to create one). If you want to create other field word lexicons with different collations, repeat this step specifying a different collation URI for the new lexicon.

- Click the Includes tab to specify elements to include in the field.

**Add Field Include**

Summary Configure **Includes** Excludes Create Help

ok cancel

**included element** -- *The element included in the field.*

**namespace uri**   
A namespace URI.

**localname**   
The localname of the included element.  
**Required. You must supply a value for localname.**

**weight**   
The weight, used to boost or lower relevance scores, of the included element.

**attribute namespace uri**   
Namespace of the child attribute.

**attribute localname**   
Localname of the child attribute.

**attribute value**   
Include only elements with the specified attribute having this value.

ok cancel

- On the Included Element page, specify a localname for the element to include. If the element is in a namespace, specify the namespace URI for the element to include.
- [OPTIONAL] If you want to boost or lower the relevance contribution for matches within this element, specify a weight other than the default of 1.0. Weights greater than 1.0 will boost the relevance contribution and weights lower than 1.0 will lower the contribution.
- [OPTIONAL] If you want to only include elements that have an attribute with a specified value, enter the attribute namespace uri (if needed), the attribute localname, and a value for the attribute. Then only elements containing attributes with the specified value will be included. You must specify the exact value; no wildcard characters are used.
- When you have specified everything for this element, click OK.
- Repeat steps [9](#) through [13](#) for each element you want to include.

15. If you want to exclude any elements from the field, click the Excludes tab.
16. Enter the namespace URI (if needed) and the localname for the excluded element.

**Add Field Exclude**

Summary Configure Includes **Excludes** Create Help

ok cancel

**excluded element** -- *The element excluded from the field.*

**namespace uri**   
A namespace URI.

**localname**   
The localname of the excluded element.  
**Required. You must supply a value for localname.**

ok cancel

17. Click OK.
18. Repeat steps 15 through 17 for each element you want to exclude.
19. You can delete any included or excluded fields from the tables at the bottom of the field configuration page.

**Included Elements**

Localname	Namespace	Attribute	Namespace	Value	Weight
ABSTRACT				1.0	[delete]

**Excluded Elements**

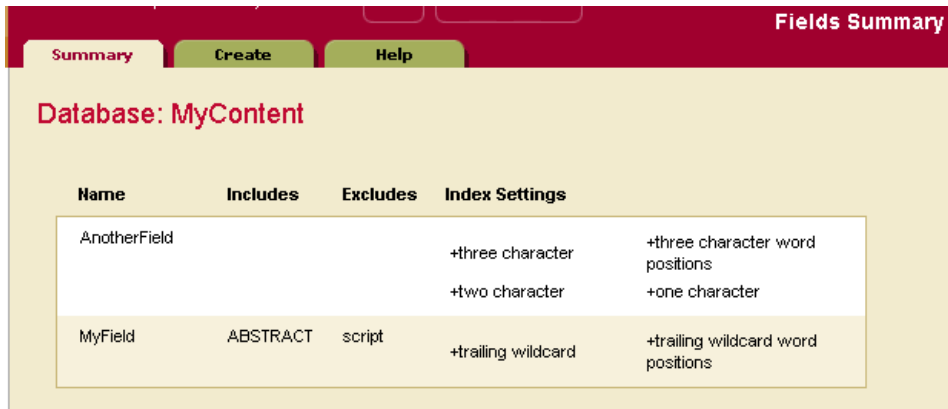
Localname	Namespace
script	http://www.w3.org/1999/xhtml

ok cancel

## 10.4.2 Modifying an Existing Field

Perform the following steps to modify an existing field:

1. To modify an existing field, click on the Fields link in the left tree menu. The Fields Summary page appears.



Name	Includes	Excludes	Index Settings
AnotherField			+three character +two character +three character word positions +one character
MyField	ABSTRACT	script	+trailing wildcard +trailing wildcard word positions

2. Click on the name of the field you want to edit. The Field Configuration page appears.
3. If you want to change any of the settings, make any desired modifications and click OK.
4. The remainder of the procedure is the same as the previous procedure for creating a field, starting with step 8 to create a field word lexicon, and continuing on to add/delete included and excluded elements.

## 11.0 Understanding and Controlling Database Merges

This chapter describes database merges and how you can control them. It includes the following sections:

- [Overview of Merges: Merges are Good](#)
- [Setting Merge Policy](#)
- [Blackout Periods for Merges](#)
- [Merges and Point-In-Time Queries](#)
- [Monitoring a Merge](#)
- [Explicit Merge Commands](#)
- [Configuring Merge Policy Rules](#)

### 11.1 Overview of Merges: Merges are Good

This section provides an overview of merges, and includes the following parts:

- [Dynamic and Self-Tuning](#)
- [What Happens During a Merge](#)
- [Dangers of Disabling Merges](#)
- [Merges Will Change Scores](#)

#### 11.1.1 Dynamic and Self-Tuning

Merges are a way of self-tuning the performance of the system, and MarkLogic Server continuously assesses the state of each database to see if it would benefit from self-tuning through a merge. In most cases, the default merge settings and the dynamic nature of merges will keep the database tuned optimally at all times. Because merges can be resource intensive (both disk I/O and CPU), however, some DBAs might need to control when merges occur and/or when they do not occur. You can do that by setting your merge policy as appropriate for your environment, as described in “Setting Merge Policy” on page 82.

Dynamic and self-tuning, merges are a “good thing”; they not only reclaim disk space, but improve the query and search performance of the system. Databases are made up of one or more forests, and forests are made up of one or more *stands*. The more stands there are in a forest, the more time it takes to resolve a query. Merges reduce the number of stands in each forest in a database, thereby improving the time it takes to resolve queries.

### 11.1.2 What Happens During a Merge

A database consists of one or more forests, and each forest consists of one or more stands. Each stand consists of one or more fragments. When a document is updated, new versions of all of the fragments associated with the document update are created in a new stand. Any old versions of the fragment remain in the old stand with a system timestamp that lets MarkLogic Server know that they are old versions of the fragments. Similarly, when a document is deleted, its fragments remain in the old stand with a system timestamp that lets MarkLogic Server know that they are old versions of the fragments.

Merges occur to move any unchanged fragments from an old stand into a new stand, deleting any old versions of fragments (including deleted fragments), thereby freeing up disk space and compacting the usable fragments so they are all together on disk. Additionally, merges combine index data for all of the fragments in a stand, thereby optimizing the indexes. Merges are a normal part of database operation, and they ensure that the system continues to perform at its best as updates and deletes occur.

To summarize, as part of merging, the following occurs:

- Multiple stands are combined into one for improved performance.
- Disk space is reclaimed.
- Indexes and lexicons are combined and re-optimized based on their new size.

The result is a database that is smaller and can resolve queries much faster than before the merge.

### 11.1.3 Dangers of Disabling Merges

MarkLogic Server is designed to periodically merge. Although there is a control to disable merges, it is dangerous to leave merges disabled on a database when there is a lot of updates occurring to the system. While disabling merges might eliminate some contention for resources during periods where merges and other requests are simultaneously occurring on the system, the performance of MarkLogic Server will degrade over time if merges not allowed to proceed when changes (inserts, updates, deletes) are made to the database.

Furthermore, disabling or eliminating merging may eventually lead to a condition in which the server is unable to make changes to the database. For example, when an in-memory stand fills up, it is written to an on-disk stand. MarkLogic Server has a fixed limit for the maximum number of stands (64), and eventually, that limit will occur and you will no longer be able to update your system.

In most cases where merges are causing disruptions to your system, you should be able to adjust the merge policy parameters to settings that will work in your environment. If you do need to disable merges, however, be sure to monitor the system and make sure the number of stands per forest does not grow too high. For details on setting merge controls, see “Description on Merge Parameters” on page 83 and “Configuring Merge Policy Rules” on page 89.

### 11.1.4 Merges Will Change Scores

When a database merges, it deletes old fragments that exist in the database, therefore changing (making it smaller) the total number of fragments in the database. Because the number of fragments in the database is used in determining the score for a `cts:search` operation, merges will have an impact on search scores, which in turn might impact the order of search results (which are ordered by relevance score).

The amount of impact that merges have on scores is dependent on how many old versions of fragments there are waiting to be merged, the content of the old fragments, and the overall size of the database. For large databases with relatively little amount of change, the difference in the scores will be very small. For smaller databases with large amount of change, the differences in scores can be significant before and after a merge completes.

## 11.2 Setting Merge Policy

This section describes the tools you can use to control merges, and has the following parts:

- [Overview of the Merge Policy Controls](#)
- [Description on Merge Parameters](#)

### 11.2.1 Overview of the Merge Policy Controls

If you determine that you need to manage your merges, there are several types of controls to help you manage the conditions in which merges occur:

- The following controls determine the conditions under which MarkLogic Server deems a merge is desirable:
  - `merge min size`
  - `merge min ratio`
- The following controls determine the conditions under which a merge will be allowed:
  - `merge max size`
  - `merge enable`
  - `merge blackout periods`
- The following control determines if multiple versions of fragments are preserved when a merge is performed:
  - `merge timestamp`
- The following controls explicitly initiate a merge (see “Manually Initiating a Merge” on page 87):
  - `xdmp:merge()`
  - The merge button in Admin Interface.

- The Admin Interface has controls for cancelling a merge (see “Cancelling a Merge” on page 88).

For more information on how set up your system to better control merges, see “Configuring Merge Policy Rules” on page 89.

### 11.2.2 Description on Merge Parameters

The following table describes the settings available on the Databases > *db\_name* > Merge Policy page of the Admin Interface. These parameters determine when automatic merges occur on a database, as well as other administrative functions.

Database Setting	Description
<code>merge enable</code>	Allows merges to occur. Set this to false to disable merges for the database. Use care when setting this to false, as merges are ultimately required for the system to maintain performance levels and to allow optimized updates to the system.
<code>merge max size</code>	The maximum size, in megabytes, of a stand that will result from a merge. If a stand grows beyond the specified size, it will not be merged. If two stands would be larger than the specified size if merged, they will not be merged together. If you set this to smaller sizes, large merges (which may require more disk and CPU resources) will be prevented. Set this to 0 (the default) to allow any sized stand to merge. Use care when setting this to a non-zero value, as this can prevent merges which are ultimately required for the system to maintain performance levels and to allow optimized updates to the system.
<code>merge min size</code>	The minimum number of fragments that a stand can contain. Two or more stands with fewer than this number of fragments are automatically merged.
<code>merge min ratio</code>	A positive integer indicating the minimum ratio between the number fragments in a stand and the number of fragments in all of the other smaller stands (that is stands with fewer fragments) in the forest. Stands with a fragment count below this ratio relative to all smaller stands are automatically merged with the smaller stands. For an example, see “If You Want to Avoid ‘Large’ Merges” on page 89.

Database Setting	Description
merge timestamp	<p>The timestamp stored on merged stands. This is used for point-in-time queries, and determines when space occupied by deleted fragments and old versions of fragments may be reclaimed by the database. If a fragment is deleted or updated at a time after the merge timestamp, then the old version of the fragment is retained for use in point-in-time queries. Set this to 0 (the default) to let the system reclaim the maximum amount of disk space during merge activities. A setting of 0 will remove all deleted and updated fragments when a merge occurs. Set this to 1 before loading or updating any content to create a complete archive of the changes to the database over time. Set this to the current timestamp to preserve all versions of content from this point on. The timestamp is a number maintained by MarkLogic Server that increments every time a change occurs in any of the databases in a system (including configuration changes from any host in a cluster). To set to the current timestamp, click the <code>current timestamp</code> button; the timestamp is displayed in red until you press OK to activate the timestamp for future merges. For details on point-in-time queries, see the <i>Developer's Guide</i>.</p>
merge blackout periods	<p>Specify times when merges are disabled. To specify a merge blackout period, click the Create tab and specify when you want the blackout to occur. You can make it a recurring blackout period, or specify a one-time blackout period. Use caution when setting large blackout periods when there are significant updates occurring on the system; merges are a normal part of the self-tuning mechanism of the database, and disabling them completely or for long periods of time can cause performance degradation.</p>

### 11.3 Blackout Periods for Merges

Although merges are a normal part of system behavior, there are times when it is inconvenient for a merge to start. Merge blackout periods allow you to specify times when a merge should not begin. This section describes merge blackouts and includes the following parts:

- [Understanding Merge Blackouts](#)
- [Configuring Merge Blackout Periods](#)
- [Deleting Merge Blackout Periods](#)

### 11.3.1 Understanding Merge Blackouts

A merge blackout is a predetermined time period in which automatic merges are disabled. A Merge that starts before a merge blackout period will continue until either it completes or until it is canceled, even if the merge continues into a blackout period. If you want to stop any merges at the beginning of a blackout period, you must cancel them manually as described in “Cancelling a Merge” on page 88. Because merges that start just before a blackout period will continue into the blackout period, if you want to be sure no merges occur during a time period you should make the blackout period start earlier. This is especially true for merges that might run a long time.

If the system determines that a merge is required and it is during a blackout period, the merge will not begin until the blackout period is past.

### 11.3.2 Configuring Merge Blackout Periods

Perform the following to configure merge blackout periods:

1. In the Admin Interface tree menu, click the Databases > *db\_name* link, where *db\_name* is the name of the database in which you want to specify merge blackout periods.
2. Click the Merge Controls menu item under your database. The Merge Control Configuration page appears.
3. Click the Create tab. The Add Merge Blackout page appears.

**Add Merge Blackout Periods to a Database**

**merge blackout type**       recurring     one time

**this blackout will**       disable merges completely     limit merges to:  MBs

**days**       Monday     Tuesday     Wednesday     Thursday     Friday     Saturday     Sunday  
The days this blackout is active.

**this blackout will last**       all day     for a time period

4. Fill in the form as needed for the blackout period you want to create. Clicking the radio buttons will bring up more forms to complete.
5. Click OK to create the blackout period.

The new blackout period will take effect immediately.

### 11.3.3 Deleting Merge Blackout Periods

Perform the following to delete a merge blackout period:

1. In the Admin Interface tree menu, click the Databases > *db\_name* link, where *db\_name* is the name of the database in which you want to delete a merge blackout period.
2. Click the Merge Controls menu item under your database. The Merge Control Configuration page appears.
3. In the area corresponding to the blackout period you want to delete, click the Delete button.
4. Click OK on the confirmation page to delete the blackout period.

The blackout period is deleted immediately.

## 11.4 Merges and Point-In-Time Queries

When a merge occurs, it deletes all fragments from the stands being merged that have a system timestamp older than the configured `merge_timestamp` (unless the `merge_timestamp` is set to 0, in which case it will delete all fragments older than the current timestamp). This can keep multiple versions of some fragments in the database. You can query the older fragments using point-in-time queries. For details, see the chapter on “Point-In-Time Queries” in the *Developer’s Guide*.

## 11.5 Monitoring a Merge

There are two main places to look for monitoring information about merges:

- [Messages in the ErrorLog.txt File](#)
- [Database Status Page](#)

### 11.5.1 Messages in the ErrorLog.txt File

MarkLogic Server logs INFO level messages to the `ErrorLog.txt` file whenever a merge begins, completes, or is canceled. Additionally, there are other log messages that are logged at more detail logging levels during a merge. The following are some sample log messages for a typical merge:

```
2006-04-20 13:43:11.151 Info: Merging /var/opt/MarkLogic/Forests/bill/00000004 and /var/opt/MarkLogic/Forests/bill/00000005 to /var/opt/MarkLogic/Forests/bill/00000006
2006-04-20 13:43:15.726 Debug: OnDiskStand /var/opt/MarkLogic/Forests/bill/00000006, disk=47MB, memory=20MB
2006-04-20 13:43:15.726 Info: Merged 81 MB in 4 s at 20 MB/s to /var/opt/MarkLogic/Forests/bill/00000006
2006-04-20 13:43:15.806 Debug: ~OnDiskStand /var/opt/MarkLogic/Forests/bill/00000004
2006-04-20 13:43:15.806 Debug: ~OnDiskStand /var/opt/MarkLogic/
```

```
Forests/bill/00000005
2006-04-20 13:43:15.859 Info: Deleted /var/opt/MarkLogic/Forests/bill/
000000042006-04-20 13:43:15.894 Info: Deleted /var/opt/MarkLogic/
Forests/bill/00000005
```

If you cancel a merge, you will see an message similar to the following in the `ErrorLog.txt` file:

```
2006-05-08 17:45:44.027 Error: PooledThread::run: XDMP-CANCELED:
Canceled merge of stands: 13419435601900621379, 6182944041533805976 to:
C:\Program Files\MarkLogic\Data\Forests\bill\0000009a
```

By examining the `ErrorLog.txt` file, you can determine when a merge started, when it completed, which stands were merged together, what stand they were merged into, the size of the merge, and other useful information.

**Note:** There must be sufficient disk space on the filesystem in which the forest data is stored for a merge to complete successfully; if a merge runs out of disk space, it will fail with an error message. Also, there must be sufficient disk space on the filesystem in which the log files reside to log any activity on the system. If there is no space left on the log file device, MarkLogic Server will abort. Additionally, if there is no disk space available to add messages to the log files, MarkLogic Server will fail to start.

## 11.5.2 Database Status Page

You can access the Database Status page by clicking the Databases > `db_name` link in the tree menu, then clicking the Status tab in the Admin Interface. The Database Status page lists the merge state, which indicates if a merge is going on, shows the size of the merge, and estimates how long it will take the merge to complete. Additionally, the Database Status page includes a link to cancel the current merge (for details, see “Cancelling a Merge” on page 88).

## 11.6 Explicit Merge Commands

This section describes how to manually perform the following operations:

- [Manually Initiating a Merge](#)
- [Cancelling a Merge](#)

### 11.6.1 Manually Initiating a Merge

You can manually initiate a merge, either by explicitly issuing the `xdmp:merge` command or by clicking the Merge button on the database configuration page of the Admin Interface. Either of these actions will immediately begin a merge on the database (if using `xdmp:merge`, on the database to which the App Server that responds to the request is connected, or if using the Admin Interface, the database being configured). Manually initiated merges continue even when merges are disabled for a database.

When you issue an `xdmp:merge` command or click the Merge button, it will ignore all of the merge control settings and merge all of the on-disk stands down to a single stand. This is different from automatic merges, where merges only occur if they meet the specifications set forth in the parameters for the database being merged.

**Note:** If you have updates occurring on the system while a merge is in progress, the new fragments will not be merged during the active merge operation; they will be merged during a subsequent merge.

Manually initiating a merge is useful when you have your merge controls set such that very large merges do not occur (for example, `merge min ratio` set to 1), but you want to run the large merges during a period of low activity on your system.

The `xdmp:merge` API also allows you to specify options to the merge to control the maximum merge size, the forests which are merged, as well as other options. For details, see the *Mark Logic Built-In and Module Functions Reference*.

## 11.6.2 Cancelling a Merge

You can cancel a merge in the Database Status page of the Admin Interface (Databases > `db_name` > Status tab). If you access the status page for a database during a merge, on the part of the status page for the stand(s) being merged, there is a cancel button (usually on the bottom right of the status page).

Forest	Stand	Merging	Stands	Size	Rate	Estimated Completion	
bill	00000063	00000062	1	52 MB	2.58MB/s	00:00:17	[cancel]
<b>Total</b>			<b>1</b>	<b>52 MB</b>	<b>n/a</b>	<b>n/a</b>	

When you cancel a merge, the new stand that has not completed its merge is discarded, leaving the unmerged stands as they were before the merge began. Note that if you cancel an automatic merge, it might start up a new merge as soon as it is canceled (if the merge controls are set such that a merge is triggered). To avoid this situation, you can change some of the merge control parameters before you cancel an automatic merge.

To cancel a merge:

1. Click the Databases menu item in the Admin Interface.
2. Click the name of the database, either from the tree menu or from the summary page.
3. Click the Status tab.

4. At the bottom right of the Database Status page, click the cancel button on the row for the stand being merged.
5. Click OK on the Cancel Merge confirmation page.

The merge is canceled and the Database Status page appears again.

## 11.7 Configuring Merge Policy Rules

By changing some of the merge policy parameters, you can effectively control certain aspects of your merges. The descriptions in “Description on Merge Parameters” on page 83 describes what each parameter does. This section describes some scenarios with suggestions for how to tune the merge control parameters to satisfy the conditions. It includes the following parts:

- [Determine the Baseline for Your Merges](#)
- [If You Want to Avoid ‘Large’ Merges](#)
- [Other Solutions](#)

### 11.7.1 Determine the Baseline for Your Merges

The merge characteristics of your system depend on many factors, including the size of your forests, the amount of update activity on the system, and the way your data is fragmented. If you feel you need to change the configuration of your merges, the first step is to determine the merge characteristics for your database. This requires running your system under normal loads, then analyzing the log files to determine the following about your merges:

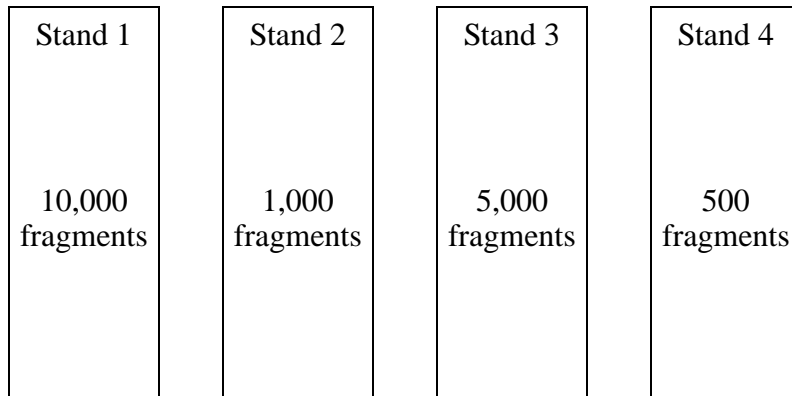
- average size of the merges
- average frequency of the merges
- average time it takes for the merges to complete

If it turns out that your merges are never taking more than a few minutes to complete, then there is probably no need to change any of your settings.

### 11.7.2 If You Want to Avoid ‘Large’ Merges

In most cases, MarkLogic Server will perform relatively small merges just often enough to keep the system properly optimized. Small merges are generally not very disruptive and reasonably fast. In some cases, however, you might find that your merges are too large and are taking too much time. Exactly how large constitutes a “Large” merge is difficult to measure, but if you determine that your merges are too large, then you might want to try and configure your settings to avoid a really large merge.

One way to accomplish this is to lower the value for `merge_min_ratio` to 1, which will effectively ensure that a merge will never be more than 1/2 the size of your forest. To illustrate this, consider the following scenario:



If the `merge_min_ratio` is set to 1, then a stand can merge if the following ratio is less than 1:

$$\frac{\text{\# of fragments in a stand}}{\text{total \# of fragments in all other smaller stands in the forest}}$$

Substituting in the values from the example for stand 1 yields:

$$10000 / (1000 + 5000 + 500) = 10000 / 6500 = 1.54$$

which is greater than 1. Therefore stand 1 is not merged. Next putting in the values for stand 2 yields:

$$1000 / (5000 + 500) = 1000 / 5500 = 0.18$$

which is less than 1. Note that stands 3 and 4 are smaller than stand 2, so the sum of the fragments in those stands appear in the denominator of the merge min ratio. Therefore stand 2 is merged. Similarly, stands 3 and 4 are merged. Therefore, a `merge_min_ratio` of 1 will cause this forest to be merged down to 2 stands, where stand 1 remains unmerged and stands 2, 3, and 4 are merged together into a new stand.

### 11.7.3 Other Solutions

In some cases, changing the merge parameters might not be the best solution for your system. For example, if your merges are taking a very long time due to slow disk drives or other system contention, addressing those issues might do more to help your merge times than any amount of tuning can do. Also, if your merges are extremely large, it could be that the forests are larger than optimal. There is no fixed maximum size for a forest, but experience in the field has shown that when forests grow over 200GB, query performance tends to start to decrease while merge times tend to start to increase. If your forests are larger than 200GB, consider breaking them into multiple forests.

## 12.0 Backing Up and Restoring a Database

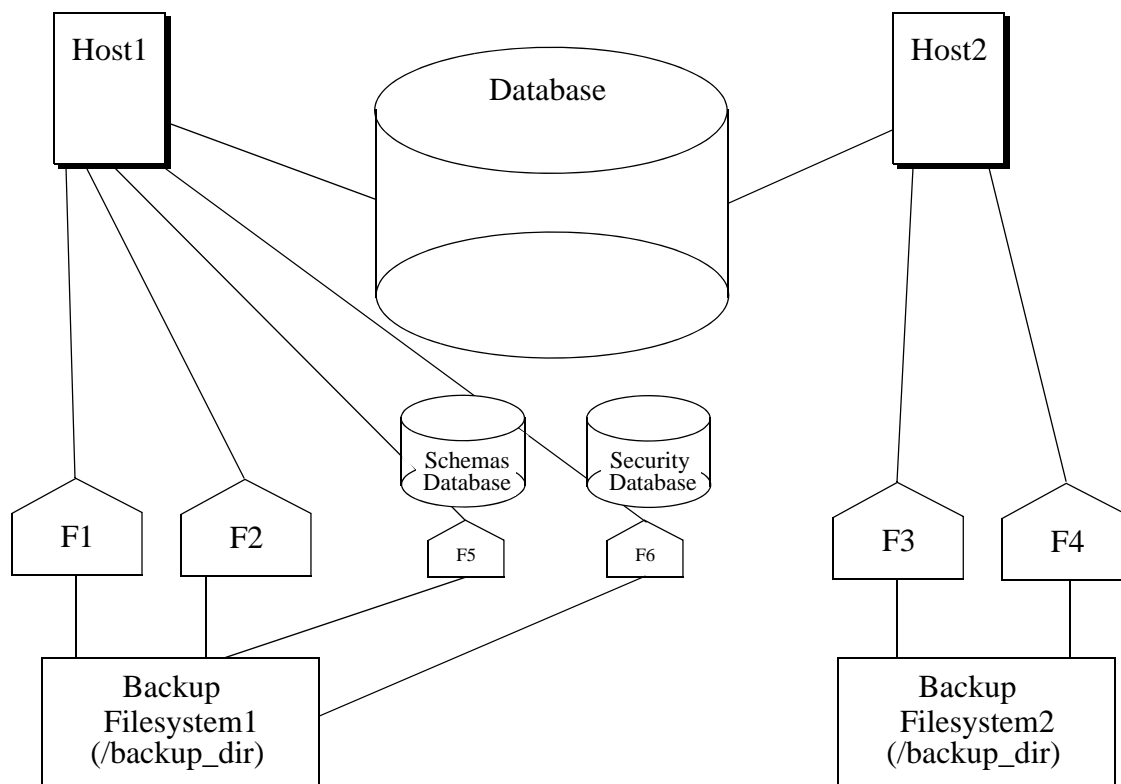
MarkLogic Server provides a facility to make a consistent backup of a database. This section describes the backup and restore architecture and provides procedures for backing up and restoring a database. The following topics are included:

- [Backup and Restore Overview](#)
- [Backing Up a Database](#)
- [Restoring a Database](#)

### 12.1 Backup and Restore Overview

Database backup and restore operations in MarkLogic Server are distributed over all of the data nodes in a cluster (that is, all of the nodes that contain forests), and provide consistent database-level backups and restores.

The directory you specify for a backup or restore operation must exist on each data node associated with the database (it can be either a shared or unshared directory). For example, if you have a data node on Host1 with forests F1 and F2, and another data node on Host2 with forests F3 and F4, then the backup directory you specify must exist on both Host1 and Host2. The following figure shows such a configuration, where the Schemas and Security databases have forests F5 and F6 respectively, and they are also attached to Host1.



### 12.1.1 Consistent, Database-Level Backup

By default, when you back up a database you backup everything associated with it, including the following:

- The configuration files.
- The Security database, including all of its forests.
- The Schemas database, including all of its forests.
- All of the forests of the database you are backing up.

If you choose to back up all forests, you will have a backup that you can restore to the exact same state as when the backup begins copying files.

You can also backup any individual forests that you choose, choosing only the ones you need to backup. These forest-level backups are consistent for the data in the forest and any other forests included in the backup, but might not be consistent with changes that occur in other forests not included in the backup.

You can also choose not to backup the Security and Schemas databases. While having backups of these databases that are synchronized with the database backups is important to get the exact same view of the system as when the backup began, you might have separate processes for backing up these databases that can ensure proper consistency. For example, if they do not change frequently, you may only need to back them up when they change.

The database-level backup and restore in MarkLogic Server provides the flexibility for you to decide how much or how little you want to backup or restore. The choices you make depend on the amount of change in your system and your unique backup and restore requirements.

**Note:** Restoring the Security database will require a restart of MarkLogic Server.

### 12.1.2 Admin Interface

You use the Admin Interface to initiate backup and restore operations. Use the Backup/Restore tab for each database configured in your system to initiate backup and restore operations. For specific procedures for backup and restore operations, see “Backing Up a Database” on page 96 and “Restoring a Database” on page 99.

### 12.1.3 Backup and Restore Transactions

Backup and restore operations are transactional and therefore guarantee a consistent view of the data. They do not lock the database, however. Therefore, if the data in a database changes after a backup or restore operation begins but before it completes, those changes are not reflected in the backup or restore operation. Similarly, changes to the Security and Schemas databases during a backup or restore operation are allowed, but will not be reflected in the backup or restore.

Database and Forest administrative tasks such as drop, clear, and delete cannot take place during a backup; any such operation is queued up and will initiate after the backup transaction has completed.

### 12.1.4 Backup Directory Structure

When you back up a database, you specify a backup directory. That directory must exist on each host in your configuration, and must be readable and by the user running MarkLogic Server (by default `daemon` on UNIX and the local System user on Windows). The backup directory structure for each host is the same, except that the forests are only backed up on the host from which they are served.

Below the specified backup directory, a subdirectory is created with a name based on the date when the backup begins. Each of these subdirectories contain one backup. The following is the basic backup directory structure.

```

<specified_backup_dir>/
  <date_1>-1/
    *.xml
    BackupTag.txt
    Forests/
      <security_forest_1>/
        <forest_files_and_directories>
      <security_forest_n>/
        <forest_files_and_directories>
      <schemas_forest_1>/
        <forest_files_and_directories>
      <schemas_forest_n>/
        <forest_files_and_directories>
      <database_forest_1>/
        <forest_files_and_directories>
      <database_forest_n>/
    <date_1>-n/
      <backup_directory structure>
  <date_n>-1/
    <backup_directory structure>
  <date_1>-n/
    <backup_directory structure>

```

For example, if you back up a database to the `/space/backups` directory on September 1, 2004, a directory structure similar to the following is created:

```

/space/backups
  20040901-1/
    *.xml
    BackupTag.txt
    Forests/
      Documents/
        Label
        000001e1/
          Journals/
            Schemas/
              Label
              000001e1/
                Journals/
                  Security/
                    Label
                    000001e1/
                      Journals/


```

### 12.1.5 Phases of Backup or Restore Operation

Backup and restore operations are divided into the following phases:

- Validation
- Copy
- Synchronization

The following figure shows the phases of a backup or restore operation:



#### 12.1.5.1 Validation Phase

The validation phase is where the backup directories are checked to make sure that all of the needed files exist and that all of the needed backup directories exist and are writable. For backup operations, they are checked for sufficient disk space. For restore operations, the configuration files are read and the other backup files are checked to make sure they appear to be valid. The validation phase does not actually write any data and is completely asynchronous.

### 12.1.5.2 Copy Phase

The copy phase is where the files are actually copied to or from the backup directory. The configuration files are copied at the beginning of the backup operation, and at this point a timestamp is written to the `BackupTag.txt` file. The copy phase that might take a significant amount of time, depending on the size of the database. The start of the copy phase starts a transaction; if the transaction fails on a restore operation, the database remains unchanged from its original state.

### 12.1.5.3 Synchronization Phase

During a backup or restore operation, the synchronization phase is where cleanup tasks such as deleting temporary files takes place, leaving the database in a consistent state. During a restore operation, the synchronization phase also takes the old version of the database offline and replaces it with the newly restored version.

**Note:** Any “cold” administrative tasks (tasks that require a server restart) will cause any backup or restore operations to fail. Do not perform any “cold” administrative tasks during a backup or restore operation. For a list of “hot” and “cold” operations, see “Appendix A: ‘Hot’ versus ‘Cold’ Admin Tasks” on page 188.

## 12.1.6 Notes about Backup and Restore Operations

This section provides notes and restrictions about backing up and restoring MarkLogic Server databases.

- The backup files are platform specific—backups on a given platform should only be restored onto the same platform. This is true for both database and forest backups.
- You can restore an individual forest using a database backup by unchecking all forests except the one you want to restore on the Confirm Restore screen (see step 7 in “Restoring a Database” on page 99).
- We recommend using the database-level backup/restore, not the forest-level backup/restore. If you do use the forest-level backup/restore, note that you cannot restore a backup created with the forest-level backup as a database-level restore operation; forest-level backups created with the forest backup/restore utility must be restored from the forest restore utility. For details, see “Restoring a Forest” on page 109.
- The restore operation is designed to restore into a database that has the same configuration settings as the one that was backed up, but it neither requires nor checks that the configurations are the same. The restore operation must occur on a database that has its configuration defined. Also, the restore operation does not change the database configuration files. Because the configuration files hold all of the database configuration information such as index options, fragmentation, range indexes, and so on, the restored database will take on the configuration information of the database to which it is restored. If this configuration information is different from the database that was backed up, and if reindexing is enabled, the database will reindex to the new configuration after the restore completes.

## 12.2 Backing Up a Database

Perform the following steps to backup a database:

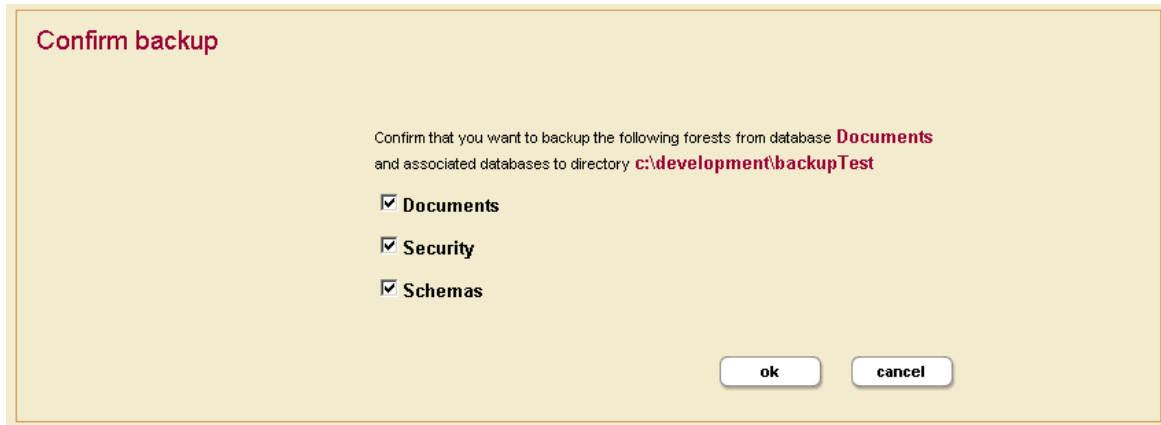
1. Log into the Admin Interface as a user with the Admin role.
2. Click the Databases link in the left menu of the Admin Interface.
3. Click the database name for the database you want to back up, either on the tree menu or on the summary page.
4. Click the Backup/Restore tab. The Backup/Restore screen appears.
5. Enter the directory to which you want the database backed up in the Backup to directory field.

**Note:** The directory path must exist on all hosts that serve any forests in the database.

The image shows two dialog boxes for the 'Documents' database. The top dialog is titled 'Backup the Documents database.' It has a label 'Backup to directory' and a text input field containing 'c:\development\backupTest'. Below the input field is the text 'The backup directory pathname.' and 'Required.' in red. At the bottom are 'ok' and 'cancel' buttons. The bottom dialog is titled 'Restore the Documents database.' It has a label 'Restore from directory' and an empty text input field. Below the input field is the text 'The backup directory pathname.' and 'Required.' in red. At the bottom are 'ok' and 'cancel' buttons.

6. Click OK.
7. If an invalid directory error appears, then the directory does not exist or is not writable. Create the directory with the proper permissions (readable and writable by the user running MarkLogic Server, by default `daemon` on UNIX and the local System user on Windows) and click OK again.

8. The Confirm Backup screen appears and lists all the forest selected for back up.



9. Click OK to begin the backup immediately, or deselect forests that you do not want to back up.

**Note:** If you deselect any of the forests to backup, you might not have a completely consistent view of the database to restore. Only deselect any forests if you are sure you understand the implications of what you are backing up. To guarantee the exact same view of the database, backup all of the forests associated with the database, including the Schemas and Security database forests.

- After the backup is underway, the Admin Interface redirects you to the Database Status page.

**Database: Documents** show more

database status -- A detailed view of this database's status.

<b>Database</b>	Documents
<b>Mount State</b>	Online (2005/04/19 13:45:60)
<b>Size</b>	4 MB
<b>Forests</b>	1
<b>Merge State</b>	0 merges in progress
<b>Reindexing/Refragmending State</b>	Not reindexing/refragmending
<b>Backup/Recovery State</b>	Backup in progress (see below for details)

Forest	Host	State	Documents	Fragments	Deleted Fragments	Stands	Size	Free Space
Documents	dsokolsky-ll.marklogic.com	open	57	125	0	1	4 MB	21,963 MB
<b>Total</b>			<b>57</b>	<b>125</b>	<b>0</b>	<b>1</b>	<b>4 MB</b>	

**List Cache**

Forest	Hits	Misses	Ratio	Hit Rate	Miss Rate	Ratio	
Documents	24	2	92%	0.8	0	100%	
<b>Total</b>			<b>24</b>	<b>2</b>	<b>92%</b>	<b>0.8</b>	<b>100%</b>

**Compressed Tree Cache**

Forest	Hits	Misses	Ratio	Hit Rate	Miss Rate	Ratio
Documents	0	0	n/a	0	0	n/a
<b>Total</b>			<b>0</b>	<b>0</b>	<b>n/a</b>	<b>n/a</b>

**Backups**

Forest	Path	Start Time	Estimated Completion In	Current Size	Final Size
Documents	c:\development\backupTest\20050419-3\Forests\Documents	4:58 PM April 19, 2005	unknown	0 MB	82 MB

- You can refresh the Database Status screen to view the progress of the backup.

**Backups**

Forest	Path	Start Time	Estimated Completion In	Current Size	Final Size
Documents	c:\development\backupTest\20050419-5\Forests\Documents	5:18 PM April 20, 2005	00:00:08	44 MB	82 MB

The Backups table lists when the backup was started, provides an estimate of the amount of time left, and lists other status information about the backup operation. When the backup is complete, the entry in the backup table disappears.

If the status for any of the forests was something besides “completed,” then an error occurred during the backup operation. Check the `Mark_Logic_Data/Logs/ErrorLog.txt` file for any errors, correct them, and try the backup operation again.

### 12.3 Restoring a Database

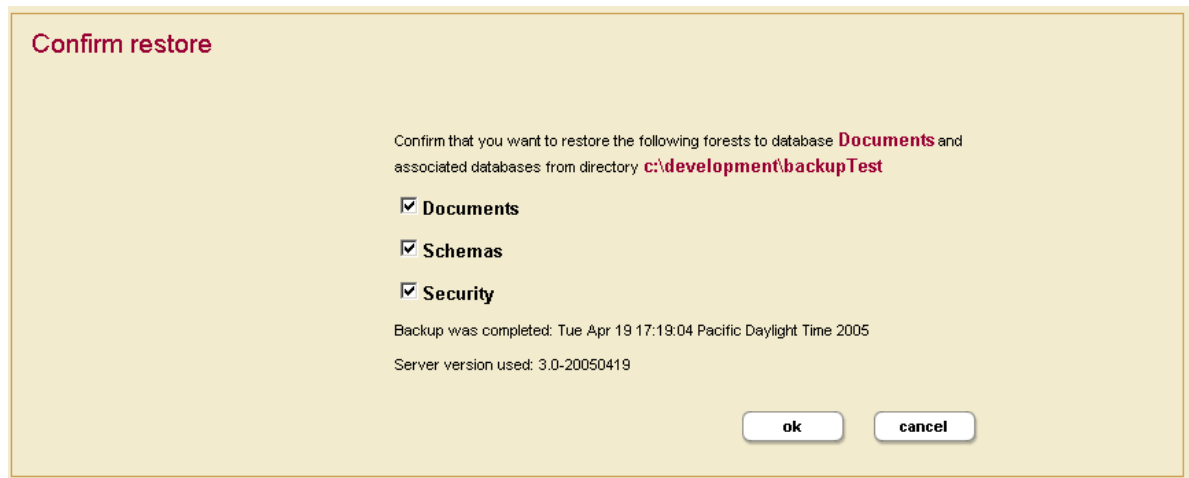
To restore an entire database from a backup, perform the following steps:

1. Log into the Admin Interface as a user with the Admin role.
2. Click the Databases link in the left menu of the Admin Interface.
3. Click the database name for the database you want to restore, either on the tree menu or on the summary page. This database should have the same configuration settings (index options, fragmentation, range indexes) as the one that was backed up.
4. Click the Backup/Restore tab. The Backup/Restore screen appears.
5. Enter the directory in which the back up exists in the Restore from directory field.

The image shows two dialog boxes from the MarkLogic Admin Interface. The top dialog is titled "Backup the Documents database." It has a label "Backup to directory" and an empty text input field. Below the field is the text "The backup directory pathname." and "Required." in red. At the bottom are "ok" and "cancel" buttons. The bottom dialog is titled "Restore the Documents database." It has a label "Restore from directory" and a text input field containing the path "c:\development\backupTest". Below the field is the text "The backup directory pathname." and "Required." in red. At the bottom are "ok" and "cancel" buttons.

**Note:** If you enter a directory that contains multiple backups of the same database, the latest one is used. If you want to choose a particular backup to restore, enter the `date_stamp` subdirectory corresponding to the backup you want to restore. For details of the directory structure, see “Backup Directory Structure” on page 93.

6. Click OK.
7. The Confirm restore screen appears and lists all the forest selected for restoring.



The Confirm restore screen also lists the date the backup was performed and the server version used for the backup you selected.

8. By default, all of the forests associated with a database are checked to restore. If you do not want to restore all of the forests, deselect any forests you do not want to restore.

**Note:** If you deselect any of the forests to restore, you might not be restoring a completely consistent view of the database. Only deselect any forests if you are sure you understand the implications of what you are restoring. To guarantee the exact same view of the database, restore all of the forests associated with the database, including the Schemas and Security database forests.

9. Click OK to begin the restore operation.

The Restores table lists when the restore was started, provides an estimate of the amount of time left, and lists other status information about the restore operation. When the restore is complete, the entry in the backup table disappears.

If the status for any of the forests was something besides “completed,” then an error occurred during the restore operation. Check the *Mark\_Logic\_Data/Logs/ErrorLog.txt* file for any errors, correct them, and try the restore operation again.

## 13.0 Hosts

A host is an instance of MarkLogic Server. A host is not configured individually but as a member of a group. A host is added to the *Default* group if it is not joined to another group during the installation process. For example, in cases of Standard Edition or Enterprise Edition running in a single host environment, the host is added to the *Default* group.

Forests are created on hosts and added to a database to interact with HTTP servers and XDBC Servers running on the same or other hosts.

See the chapters “Groups” on page 16 and “Databases” on page 44 for more details on hosts as they relate to groups and databases.

A host is managed from both the Group and Hosts configuration screens. Use the following procedures to administer your hosts

- [Adding a Host to a Cluster](#)
- [Changing the Group of the Host](#)
- [Shutting Down or Restarting a Host](#)
- [Clearing a Forest on a Host](#)
- [Deleting a Forest on a Host](#)
- [Leaving the Cluster](#)

### 13.1 Adding a Host to a Cluster

This only applies for MarkLogic Server running Enterprise Edition in a distributed environment. For information about installing MarkLogic Server and a more detailed procedure about joining a cluster, see the *Installation Guide*.

To add a host to a cluster, perform the following steps:

1. Install MarkLogic Server on the host if it is not already installed.
2. Start MarkLogic Server.
3. Access the Admin Interface on the host in which you want to add to the cluster and accept the license agreement.

4. After the server restarts, you will be prompted to join a cluster.

**Join a Cluster**

Now that MarkLogic Server is installed on this host, you can join an existing cluster. In order to do so, enter the host name of one of the cluster's hosts and provide the port number of that host's administration interface.

Press skip if you do not wish to join a cluster.

**Host Name**   
One of the target cluster's hosts  
**Required.**

**Admin Port**   
Port for admin interface on server  
**Required.**

5. Enter the DNS name or the IP address of one of the machines in the cluster. For example, if this is the second host you are installing, you can enter the DNS name of the first host you installed.
6. You will be prompted for an admin username and password. Enter the admin username and password for the security database used by the cluster. Click OK.
7. Select a Group to assign this host. Click OK.
8. Click OK to confirm that you are joining the cluster.
9. Click OK for the confirmation message that indicates that you have joined the cluster.

## 13.2 Changing the Group of the Host

To change the group to which a host belongs, perform the following steps:

1. Click the Hosts icon in the left frame.
2. Click the name of the host you want to change, either on the tree menu or the summary page.
3. Select from the available groups in the Group drop-down menu.
4. Click OK to confirm the change.

Changing the group to which a host belongs is a “cold” task; the server restarts to reflect the changes.

### 13.3 Shutting Down or Restarting a Host

To shut down or to restart a host, perform the following steps:

1. Click the Hosts icon on the left tree menu.
2. Click the name of the host you want to shut down or restart, either on the tree menu or the summary page.
3. Click the Status tab at the top right.
4. Click the Shutdown or the Restart button as appropriate.
5. Click OK to confirm to confirm the shutdown or restart operation.

**Note:** The restart operation normally completes within a few seconds. It is possible, however, for it to take longer under some conditions (for example, if the Security database needs to run recovery or if the connectivity between hosts in a cluster is slow). If it takes longer than a few seconds for MarkLogic Server to restart, than the Admin Interface might return a 503: *Service Unavailable* message. If you encounter this situation, wait several seconds and then reload the Admin Interface.

### 13.4 Clearing a Forest on a Host

Clearing a forest on a host permanently deletes the data in the forest. The configuration information of the forest will be preserved. For example, you may want to clear the forest if you want to load new data into the same configuration.

To clear the data from a forest, perform the following steps:

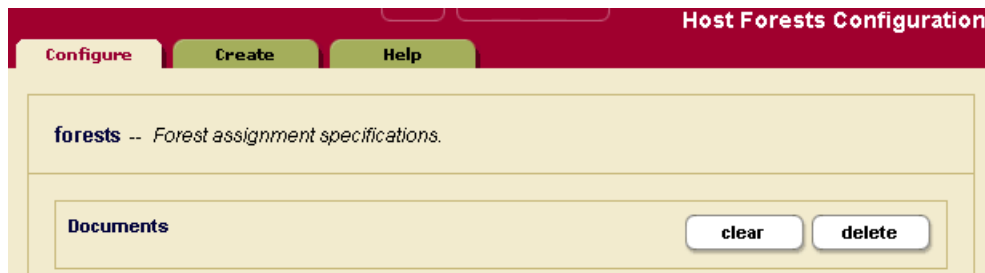
1. Click the Hosts icon on the left tree menu.
2. Click the name of the host which contains the forest you want to clear, either on the tree menu or the summary page.
3. Click the Forests icon under the selected host.
4. Click the Clear button corresponding to the forest you want to clear.
5. Click OK to confirm clearing the data from the forest.

### 13.5 Deleting a Forest on a Host

Deleting a forest on a host permanently deletes the data in the forest as well as the configuration information. A forest cannot be deleted if it is still attached to a database. You must first detach the forest from the database before you can delete from a host.

Assuming that the forest is not attached to any database, perform the following steps to delete a forest from a host.

1. Click the Hosts icon on the left tree menu.
2. Click the name of the host which contains the forest you want to delete, either on the tree menu or the summary page.
3. Click the Forests icon under the selected host.
4. Click on the Delete button corresponding to the forest you want to delete.
5. Click OK to confirm deleting the forest from the host.



6. Click the Delete button.
7. Click OK to confirm dropping the host.

Deleting a host is a “hot” admin task for the other hosts in the group.

### 13.6 Leaving the Cluster

A host has to leave a cluster first to be moved to another cluster. Leaving a cluster is also a way to switch a host from a single host environment to a multi-host environment or vice versa. In host cannot leave a cluster if there are still forests assigned to it. In a single-host environment, a host cannot leave a cluster because it will always have forests assigned to it.

Perform the following steps to make a host leave a cluster:

1. Access the Admin Interface from any host in the cluster.
2. Click on the Hosts icon in the left frame.

3. Click on the name of the host you want to remove from the cluster. The host configuration screen appears:

The screenshot shows a 'Host Configuration' dialog box. At the top, there is a title bar with the text 'Host Configuration' and four tabs: 'Summary', 'Configure', 'Status', and 'Help'. The 'Configure' tab is selected. Below the tabs, there are two buttons: 'ok' and 'cancel'. The main area of the dialog is titled 'host -- The host specification.' and contains three input fields: 'host name\*' with the value 'pubs.marklogic.com', 'group\*' with a dropdown menu set to 'Default', and 'bind port\*' with the value '7999'. Below these fields is a note: '\* -- requires restart of one or more hosts'. At the bottom of the dialog are two buttons: 'ok' and 'cancel'.

4. Click on the Leave button.
5. Click OK to confirm leaving the cluster.
6. The host restarts to load the new configuration.
7. Click OK to self-install initial databases and application servers.
8. You will be prompted to join a cluster.
9. To join another cluster, enter the name of one of the hosts in that cluster and click OK. Otherwise, click Skip.
10. Set up an admin user name and password if prompted.
11. Log in with the admin user name and password if prompted.

You should see the Admin Interface.

## 14.0 Forests

A forest is a collection of documents. They can be XML, text (CLOB), or binary (BLOB) documents. Forests are created on hosts and attached to databases to appear as a contiguous set of content for query purposes. A forest can only be attached to one database at a time.

Use the following procedures to create, manage and maintain your forests:

- [Creating a Forest](#)
- [Making Backups of a Forest](#)
- [Restoring a Forest](#)
- [Clearing Data in a Forest](#)
- [Deleting a Forest from a Host](#)

**Note:** You cannot load data into a forest that is not attached to a database.

### 14.1 Creating a Forest

To create a new forest, complete the following procedure:

1. Click the Forests icon in the left frame.
2. Click the Create tab at the top right. The Create Forest page displays:

The screenshot shows the 'Forest Add' dialog box. At the top, there is a title bar 'Forest Add' and a set of tabs: 'Summary', 'Configure', 'Status', 'Backup/Restore', 'Create', and 'Help'. The 'Create' tab is selected. Below the tabs, there are 'ok' and 'cancel' buttons. The main area of the dialog is titled 'forest -- The forest assignment specification.' and contains three input fields:

- forest name:** A text input field. Below it, the text reads: 'The forest name. Required. You must supply a value for forest-name.'
- host:** A dropdown menu showing 'dsokolsky-it.marklogic.com'. Below it, the text reads: 'The primary host to which the forest is assigned.'
- data directory:** A text input field. Below it, the text reads: 'The optional public directory for forests.'

At the bottom of the dialog, there are 'ok' and 'cancel' buttons.

3. Enter the name of your forest in the Forest Name textbox.

4. Select the host on which you want the forest to be created.
5. Enter the path to the Data Directory, which specifies where the forest data is stored.

The name of the forest is used by the system as a directory name. Therefore, the forest name must be a legal directory name and cannot contain any of the following 9 characters: \ \* ? / : < > | " . Additionally, the name cannot begin or end with a space or a dot (.). Mark Logic recommends that you use an absolute path if you specify a data directory. If you do not specify an absolute path for the data directory, your forest will be created in the default data directory.

The Forests directory is either a fully-qualified pathname or is relative to the Forests directory, set at installation time based on the directory in which MarkLogic Server is installed. The following table shows the default location Forest directory for each platform:

Platform	Program Directory
Microsoft Windows	C:\Program Files\MarkLogic\Data\Forests
Red Hat Linux	/var/opt/MarkLogic/Forests
Sun Solaris	/var/opt/MARKlogic/Forests

6. Click OK.

Creating a forest is a “hot” admin task; the changes take effect immediately.

## 14.2 Making Backups of a Forest

MarkLogic Server backs up forest data by transactionally creating an image copy of a specified forest. You can back up data at the granularity of a forest or of a database. Use the Admin Interface to back up a forest.

Forest-level backups only back up the data in a forest, and are not guaranteed to have a consistent database state to restore. The data in the forest is consistent, but other parts of the database (other forests, the schema database, and so on) might be different when you restore the data. For a guaranteed consistent backup, perform a complete database backup. For information on backing up a database, see “Backing Up and Restoring a Database” on page 91.

To back up a forest using the Admin Interface, complete the following procedure:

1. Click the Forests icon on the left tree menu.
2. Decide which forest to back up.

3. Click the icon for this forest name.
4. Click the Backup/Restore tab at the top right. The Forest Backup screen appears.

The screenshot shows the 'Forest Backup' configuration interface. At the top, there is a red header with the text 'Forest Backup' on the right. Below the header is a navigation bar with several tabs: 'Summary', 'Configure', 'Status', 'Backup/Restore' (which is highlighted in red), 'Create', and 'Help'. The main content area has a light yellow background. It features a 'Backup directory' label followed by a text input field. Below the input field, there is a red message that says 'The backup directory pathname. Required.' Underneath this, there are two radio button options: 'Backup Documents to the backup directory.' and 'Restore Documents from the backup directory.' At the bottom of the form, there are two buttons labeled 'ok' and 'cancel'.

5. Enter the name of the directory in which you want the backup copy of the forest. You must provide an absolute path.

**Warning** The software deletes *all* the files in this directory before writing the new backup. To retain multiple generations of backup, specify a different backup directory for each backup.

6. Select Backup.
7. Click OK.
8. A confirmation message appears. Click OK again to confirm the backup.

Your data in the selected forest is now backed up to the specified directory. Backing up your data is a “hot” admin task; the changes take effect immediately.

**Warning** When performing backups on the Windows platform, ensure that no users have the Forests or Data directories (or any subdirectories within them) open while the backup is being made.

### 14.3 Restoring a Forest

You can restore a forest from a backup made earlier either using the Admin Interface. Backups are restored at the forest granularity only.

**Note:** You can restore a forest in Version 3.0 from a backup made in Version 2.2 as long as the Version 2.2 backup completed cleanly. If the 2.2 backup did not complete cleanly (journal files are present), then you must restore the forest in Version 2.2 and then upgrade to Version 3.0.

To restore a forest from a backup made previously, complete the following procedure:

1. Click the Forests icon on the left tree menu.
2. Decide which forest to restore.
3. Click the icon for this forest name.
4. Click the Backup/Restore tab on the top right.
5. Enter the name of the directory that contains the backup copy of the forest.
6. Select Restore.
7. Click OK.

A confirmation message displays.

8. Confirm that you want to restore data from this backup directory and click OK.

Restoring data from your backup is a “hot” admin task; the changes take effect immediately.

**Warning** When performing restores on the Windows platform, ensure that no users have the Forests or Data directories (or any subdirectories within them) open while the restore process is executing.

### 14.4 Clearing Data in a Forest

You can clear the document data from a forest using the Admin Interface. Clearing a forest does not remove its configuration information and is useful if you want to reload document data. For example, if you move a forest from one database to another, you may want to clear and reload the documents to regenerate the indexes.

To clear all data from a forest, complete the following procedure:

1. Click the Forests icon on the left tree menu.

2. Decide which forest you want to clear.
3. Click the forest name, either on the tree menu or the summary page.

The Forest Configuration page displays:

4. Click the Clear button on the Forest Configuration page.
- A confirmation message displays.
5. Confirm that you want to clear the document data from this forest and click OK.

Clearing data in a forest is a “hot” admin task; the changes take effect immediately.

## 14.5 Deleting a Forest from a Host

You can use the Admin Interface to delete a forest, which completely removes the document data and the configuration information for the forest. The forest cannot be deleted if it is still attached to a database.

To delete a forest permanently, complete the following procedure:

1. Click the Forests icon on the left tree menu.
2. Decide which forest to delete.

3. Click the forest name, either on the tree menu or the summary page.

The Forest Configuration page displays:

The screenshot shows the Forest Configuration page. At the top, there is a red header with the text "Forest Configuration". Below the header is a navigation bar with tabs: "Summary", "Configure", "Status", "Backup/Restore", "Create", and "Help". The "Configure" tab is selected. The main content area is yellow and contains a form for configuring a forest. The form has a title "forest -- The forest assignment specification." and a "delete" button. Below the title are three fields: "forest name" with a text input containing "Documents", "host" with a dropdown menu showing "dsokolsky-ll.marklogic.com", and "data directory" with an empty text input. There are "ok" and "cancel" buttons at the bottom of the form.

4. Click the Delete button on the Forest Configuration page.  
A confirmation message displays.
5. Confirm that you want to delete this forest permanently and click OK.

Deleting a forest is a “hot” task; the changes take effect immediately.

## 15.0 Security Administration

MarkLogic Server uses a role-based security model. A user's privileges and permissions are based on the roles assigned to the user. For background information on understanding the security model in MarkLogic Server, see *Understanding and Using Security*. This section describes administration tasks related to security, and includes the following sections:

- [Security Entities](#)
- [Users](#)
- [Roles](#)
- [Execute Privileges](#)
- [URI Privileges](#)
- [Amps](#)
- [Collections](#)
- [Realm](#)

### 15.1 Security Entities

The key entities in MarkLogic Server's security model are:

- User  
A User within the model has a set of roles. A user has privileges and permissions within the system based on the roles he is given.
- Role  
A Role gives privileges and permissions to a user. A role may inherit from multiple roles. Role inheritance is an "is-a" relationship. Hence, an inherited role also has the privileges and permissions of its parent(s).
- Execute Privilege  
An Execute Privilege grants authorization to perform a protected action. Only roles (and their inherited roles) specified in the execute privilege can perform the action.
- URI Privilege  
A URI Privilege grants authorization to create a document within a protected base URI. Only roles (and their inherited roles) specified in the URI privilege can create the document within the protected base URI.

- Permission

A Permission protects a document or a collection. Each permission associates a single role with a capability (Read, Update, Insert). A protected document or collection has a set of associated permissions.

- Collection

A Collection groups a set of documents that are related. A document may belong to any number of collections. A collection exists in the system when a document in the system states that it is part of that collection. However, an associated collection object is not created and stored in the *Security* database unless it is protected.

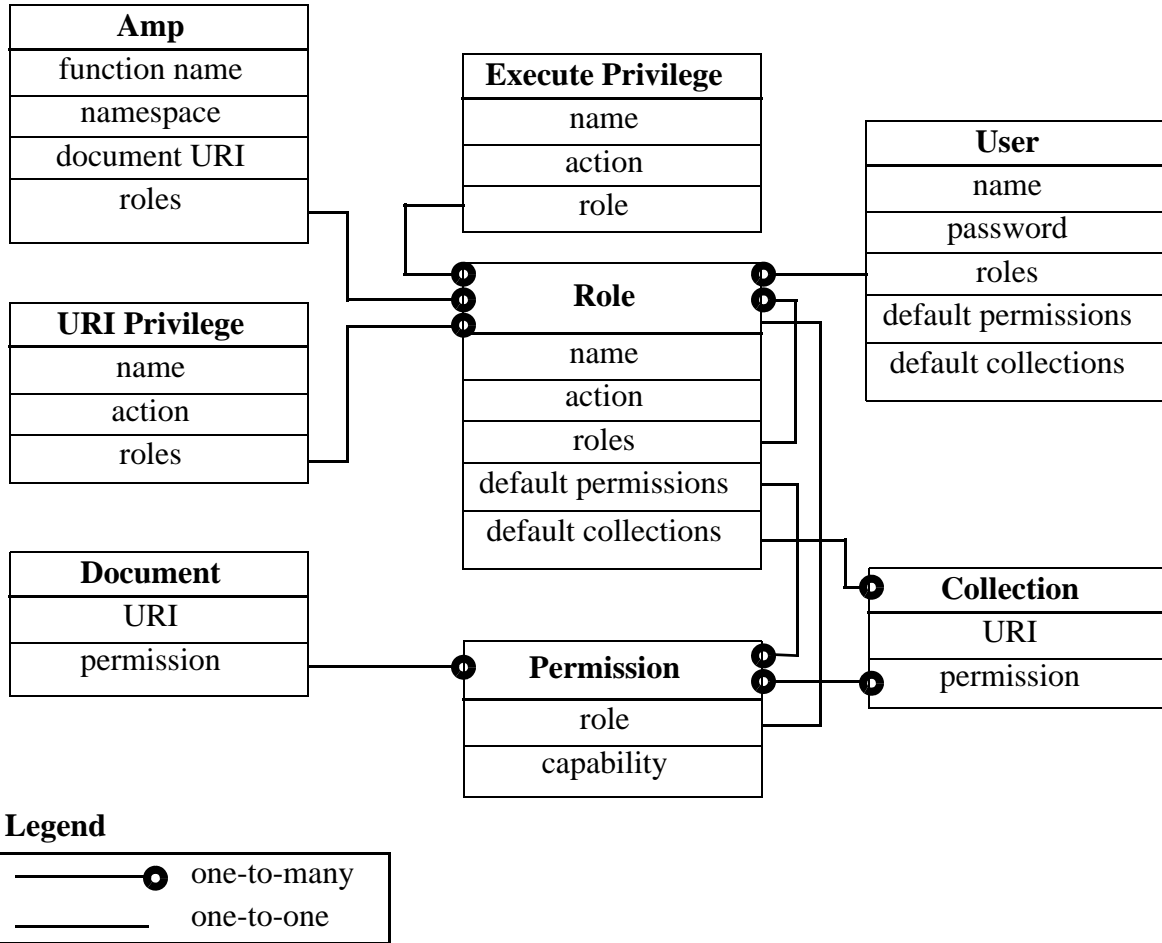
Permissions created at the collection level apply to the collection but not to documents within the collection. A user needs to have permissions at the both the collection and document level to be able to add documents to or remove documents from a collection.

- Amp

An Amp gives the User additional roles temporarily while the user is performing a certain task (executing a function).

- Security Entity Relationships

The following diagram illustrates the relationships between the different entities in the MarkLogic Server security model.



The remaining sections of this chapter detail the procedures to administer MarkLogic Server security entities. All security administrative tasks are “hot”—the changes take effect immediately without a server restart.

Permissions are not administered through the administrative interface and are not described in detail in this document. For more information on using permissions in MarkLogic Server, see the *Mark Logic Built-In and Module Functions Reference*.

## 15.2 Users

A User has a set of roles. A user has privileges and permissions within the system based on the roles he is given. A user can perform tasks (execute functions) based on his privileges and access data based on his permissions.

Each user has an associated user name and password. A user also has default collections. When a user creates a document but does not explicitly associate the document with a set of collections, the document is automatically added to the user's default collections. Default permissions can be created for a user. When a user creates a document but does not explicitly set the permissions for the document, the document will be given the user's default permissions.

If security is turned on for an HTTP or XDBC server, all users in the security database will have access to the server. Finer granularity security control to functions in XQuery programs running on the HTTP or XDBC servers are accomplished through the use of `xmdp:security-assert()` within the code. Granular secured access to documents is achieved through the use of permissions associated with each protected document.

Use the following procedures to create, manage and maintain users:

- [Creating a User](#)
- [Viewing a User's Configuration](#)
- [Deleting a User](#)

### 15.2.1 Creating a User

Follow these steps to create a user:

1. Click the Security icon in the left tree menu.
2. Click the Users icon.

3. Click the Create tab. The User Configuration page appears:

**User Configuration**

Summary Configure **Create** Help

**New User** ok cancel

**user** -- *A database user.*

**user name**   
User login name (unique)  
**Required. You must supply a value for user-name.**

**description**   
An object's description.

**password**   
Encrypted Password.  
**Required.**

**confirm password**   
Encrypted Password.  
**Required.**

**roles** -- *The roles assigned.*

admin

admin-builtins

4. Enter a name for the user in the username field.
5. Enter the description for the user (optional).
6. Enter a password for the user.
7. Re-enter the password to confirm it.
8. Under the roles section, check the roles to assign the user.
9. Create default permissions for this user (optional). Select a role and pair the role with the appropriate capability (read, insert, update). If there are more than 3 default permissions you want to add for this user, you can do so on the next screen after you click OK.

10. Create default collections for this user (optional). Type in the collection URI for each collection you want to add to the user's default collection. If there are more than 3 default permissions you want to add for this user, you can do so on the next screen after you click OK.
11. Click OK.

The user is now added to the system and the user configuration page appears. If you want to add more default permissions or collections to the user, scroll down to the section for default permissions or collections.

## 15.2.2 Viewing a User's Configuration

Perform the following steps to view a user's configuration:

1. Click the Security icon in the left tree menu.
2. Click the Users icon.
3. Locate the name of the user whose settings you want to view, either on the tree menu or on the summary page.
4. Click the name. The user configuration page appears where you can view the user's configuration:

The screenshot shows the 'User Configuration' page for a user named 'testuser'. The page has a red header with the title 'User Configuration' and several tabs: 'Summary', 'Configure', 'Describe', 'Create', and 'Help'. The 'Configure' tab is selected. Below the header, the user's name 'User: testuser' is displayed, along with 'ok' and 'cancel' buttons. The main content area shows the user's configuration details:

- user** -- *A database user.* (with a 'delete' button)
- user name**: testuser (User/login name (unique))
- description**: This is a test user (An object's description.)
- password**: [encrypted] (Encrypted Password.)
- confirm password**: [encrypted] (Encrypted Password.)

At the bottom, there is a section for **roles** -- *The roles assigned.* (inherited roles in **Bold**).

### 15.2.3 Deleting a User

Perform the following steps to delete a user from the security database:

1. Click the Security icon in the left tree menu.
2. Click the Users icon.
3. Locate the user you want to delete, either on the tree menu or on the summary page.
4. Click the user name.
5. Click on the Delete button.
6. Click OK to confirm deleting the user.

The user is permanently deleted from the security database.

## 15.3 Roles

MarkLogic Server implements a role-based security model. Therefore, the Role is a central security concept in MarkLogic Server. A role gives a user privileges (both Execute and URI) to perform certain actions in a system. An Execute Privilege allows a user to perform a protected action. A URI Privilege allows a user to create a document under a protected URI. A role also gives a user the permissions to access protected documents.

A role may inherit from multiple roles. The inheritance relationship for roles is an “is-a” relationship. Therefore, a role gets the privileges and permissions of the roles from which they inherit.

MarkLogic Server is installed with the following default roles:

Role	Description
admin	This role has the privileges and permissions needed to perform administrative tasks. This role has the highest level of access in the system.
admin-builtins	This role has the privileges needed to call the admin-builtins functions.
filesystem-access	This role has the privileges to access the filesystem.
merge	This role has the privileges needed to force a merge in the system.
security	This role has the privileges to perform all the security-related administrative functions.

While you are able to change the configuration settings of these default roles (except for the admin role) or delete any of them, we strongly recommend that you proceed with caution.

A role has default collections. When a user of a role creates a document but does not explicitly associate the document with a set of collections, the document is automatically added to a set of default collections. This set of default collections is the union of the default collections defined for the user, the roles the user has, and the roles from which the user's directly assigned roles inherit.

A role has default permissions. When a user of a role creates a document but does not explicitly set the permissions for the document, the document will be given a set of default permissions. This set of default permissions is the union of the default permissions defined for the user, the roles the user has, and the roles from which the user's directly assigned roles inherit.

For more details about the role-based security model in MarkLogic Server, see *Understanding and Using Security*.

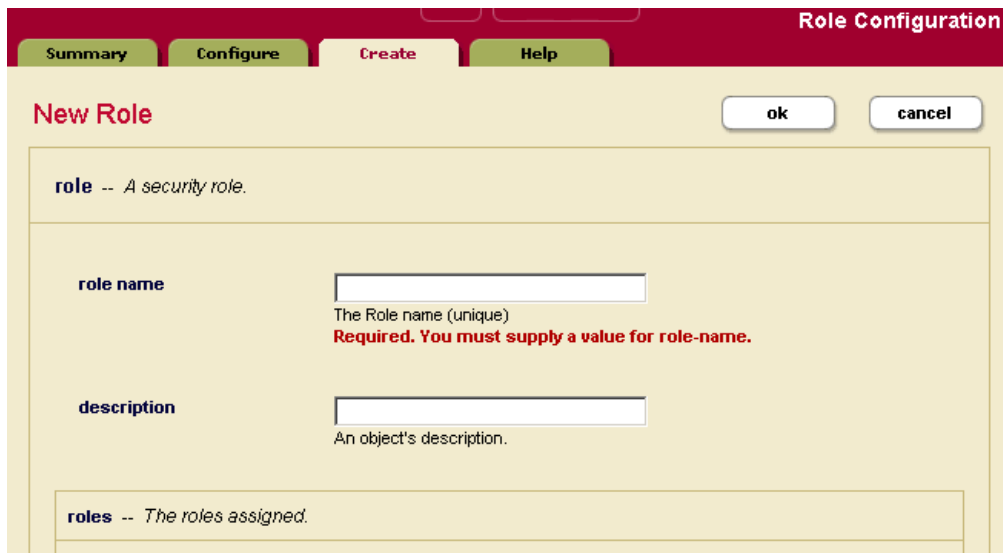
Use the following procedures to create, manage and maintain roles:

- [Creating a Role](#)
- [Viewing a Role](#)
- [Deleting a Role](#)

### 15.3.1 Creating a Role

Perform the following steps to create a user.

1. Click the Security icon in the left tree menu.
2. Click the Roles icon.
3. Click the Create tab. The Role Configuration page appears:



The screenshot shows the 'Role Configuration' page with the 'Create' tab selected. The page title is 'New Role'. There are two buttons: 'ok' and 'cancel'. The form contains the following fields:

- role** -- *A security role.*
- role name**: A text input field. Below it, the text reads: 'The Role name (unique)' and '**Required. You must supply a value for role-name.**'
- description**: A text input field. Below it, the text reads: 'An object's description.'
- roles** -- *The roles assigned.*

4. Type in a name for role in the role name field.
5. Type in a description for the role (optional).
6. Under the roles section, select the roles from which this role will inherit.
7. Under the execute privileges section, select from the available execute privileges to be associated with the role.
8. Under the URI privileges section, select from the available URI privileges to be associated with the role.
9. Create default permissions for this role (optional). Select a role and pair the role with the appropriate capability (read, insert, update). If there are more than 3 default permissions you want to add for this role, you can do so on the next screen after you click OK.
10. Create default collections for this role (optional). Type in the collection URI for each collection you want to add to the role's default collections. If there are more than 3 default permissions you want to add for this user, you can do so on the next screen after you click OK.

11. Click OK.

The role is now added to the system and the Role Configuration page appears. If you want to add more default permissions or collections to the role, scroll down to the section for default permissions or collections.

### 15.3.2 Viewing a Role

Perform the following steps to create a user.

1. Click the Security icon in the left tree menu.
2. Click the Roles icon.
3. Click the name of the role you want to view, either on the tree menu or on the summary page. The Role Configuration page appears.

The screenshot shows the 'Role Configuration' page for a role named 'security'. The page has a red header with the title 'Role Configuration' and several tabs: 'Summary', 'Configure', 'Describe', 'Create', and 'Help'. The 'Configure' tab is active. Below the header, there is a section for the role 'security' with 'ok' and 'cancel' buttons. The role is described as 'A security role.' and has a 'delete' button. Below this, there are two input fields: 'role name' with the value 'security' and a description 'The Role name (unique)', and 'description' with the value 'security role' and a description 'An object's description.'. At the bottom, there is a section for 'roles' described as 'The roles assigned. (inherited roles in **Bold**)'. There are two checkboxes: 'admin' and 'admin-builtins', both of which are unchecked.

View the configuration for the role.

### 15.3.3 Deleting a Role

You can delete a role from the security database. The system does not check to see if there are any users with that role before deleting it. A deleted role is automatically removed from all users still assigned to that role. Users who were assigned to the deleted role lose the permissions and privileges given by that role.

Perform the following steps to delete a role.

1. Click the Security icon in the left tree menu.
2. Click the Roles icon.
3. Click the name of the role you want to delete, either on the tree menu or on the summary page.
4. Click the Delete button.
5. Click OK to confirm deleting the role.

The role is now deleted from the security database.

### 15.4 Execute Privileges

An Execute Privilege grants authorization to perform a protected action. An execute privilege specifies a protected action, and the roles that can perform the action. Roles that inherit from the specified roles can also perform the protected action. The protected action is represented as a URI.

Once an execute privilege is created, it is enforced in XQuery programs through the use of `xdrm:security-assert(<protected-action-uri>, "execute")` in the code. That is, `xdrm:security-assert(<protected-action-uri>, "execute")` can be added at the entrance to function or a section of code that has been protected. If the system is executing as a user without the appropriate roles as specified by the execute privilege, an exception is thrown. Otherwise, system satisfies the security-assert condition and proceeds to execute the protected code.

Use the following procedures to create, manage and maintain execute privileges:

- [Creating an Execute Privilege](#)
- [Viewing an Execute Privilege](#)
- [Deleting an Execute Privilege](#)

### 15.4.1 Creating an Execute Privilege

Perform the following steps to create an execute privilege:

1. Click the Security icon in the left tree menu.
2. Click on the Execute Privileges icon.
3. Click the Create tab. The Execute Privilege Configuration page appears:

The screenshot shows the 'New Execute Privilege' configuration page. It features a red header with the title 'Execute Privilege Configuration' and three tabs: 'Summary', 'Create', and 'Help'. The main content area is titled 'New Execute Privilege' and includes 'ok' and 'cancel' buttons. The form is divided into three sections: 'execute privilege -- Privilege representation.', 'privilege name' with a text input field and a red error message 'Required. You must supply a value for privilege-name.', and 'action' with a text input field and a red error message 'Required. You must supply a value for action.'. Below these is a 'roles -- The roles assigned.' section with two checkboxes: 'admin' and 'admin-builtins'.

4. Enter the name of the execute privilege. Use a name that is descriptive of the action this execute privilege will protect. For example, `create-user` is the name of an execute privilege that gives a role the authorization to create a user.
5. Enter a protected action, represented as a URI. You can use any URI but we recommend you follow the conventions for your company. For example, the URI for the `create-user` execute privilege is `http://marklogic.com/xdmp/privileges/create-user`.
6. Under the roles section, select the roles that are allowed to perform the protected action.
7. Click OK.

The execute privilege is now added to the security database.

## 15.4.2 Viewing an Execute Privilege

Perform the following steps to view an execute privilege:

1. Click the Security icon in the left tree menu.
2. Click the Execute Privileges icon. The Execute Privileges Summary page appears:



Privilege	Action	Roles
amp-add-roles	http://marklogic.com/xdmp/privileges/amp-add-roles	security
amp-get-roles	http://marklogic.com/xdmp/privileges/amp-get-roles	security
amp-remove-roles	http://marklogic.com/xdmp/privileges/amp-remove-roles	security
amp-set-roles	http://marklogic.com/xdmp/privileges/amp-set-roles	security
any-collection	http://marklogic.com/xdmp/privileges/any-collection	security
any-uri	http://marklogic.com/xdmp/privileges/any-uri	security
collection-add-permissions	http://marklogic.com/xdmp/privileges/collection-add-permissions	security

3. Click on the name of the execute privilege that you want to view.
4. View the configuration for the execute privilege.

## 15.4.3 Deleting an Execute Privilege

You can delete an execute privilege from the security database. However, an exception will be thrown when a `security-assert()` on the protected action specified in the deleted execute privilege is encountered. That is, a deleted execute privilege behaves like an execute privilege for which no role has been given access to the protected action. Follow these steps to delete an execute privilege:

1. Click the Security icon in the left tree menu.

- Click the Execute Privileges icon. The Execute Privileges Summary page appears:

Privilege	Action	Roles
amp-add-roles	http://marklogic.com/xdmp/privileges/amp-add-roles	security
amp-get-roles	http://marklogic.com/xdmp/privileges/amp-get-roles	security
amp-remove-roles	http://marklogic.com/xdmp/privileges/amp-remove-roles	security
amp-set-roles	http://marklogic.com/xdmp/privileges/amp-set-roles	security
any-collection	http://marklogic.com/xdmp/privileges/any-collection	security
any-uri	http://marklogic.com/xdmp/privileges/any-uri	security
collection-add-permissions	http://marklogic.com/xdmp/privileges/collection-add-permissions	security

- Click the name of the execute privilege that you want to delete.
- On the Execute Privileges page for the given privilege, click the Delete button.
- Click OK to confirm deleting the execute privilege.

The execute privilege is now deleted from the security database.

## 15.5 URI Privileges

A URI Privilege grants authorization to create documents under a protected URI. That is, a URI privilege specifies the roles that are allowed to create documents with the protected URI as the base URI (prefix) in the document URI. Roles that inherit from the specified roles can also create the documents under the protected URI.

Unlike an execute privilege, where `xdmp:security-assert()` needs to be called explicitly to protect a function, a URI privilege is automatically enforced. When `xdmp:document-insert()` is called, the system checks the base URIs (prefix) of the document URI specified to see if they might be protected by a URI privilege. If the base URI has an associated URI privilege, it checks the roles of the user to see if any of the user's roles gives the user authorization to create the document within the protected base URI. If the user has the requisite authorization, the document is inserted into the database. Otherwise, an exception is thrown.

Use the following procedures to create, manage and maintain URI privileges:

- [Creating a URI Privilege](#)
- [Viewing a URI Privilege](#)
- [Deleting a URI Privilege](#)

### 15.5.1 Creating a URI Privilege

Perform the following steps to create a URI privilege:

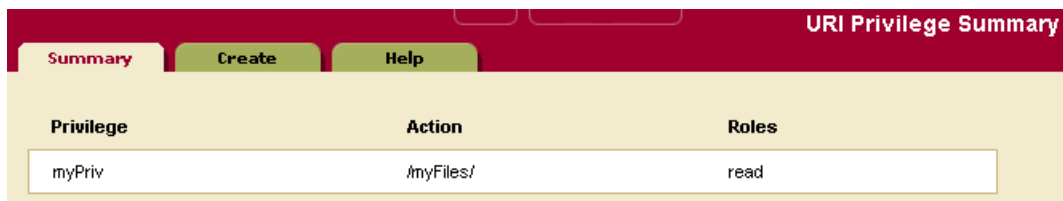
1. Click the Security icon in the left tree menu.
2. Click the URI Privileges icon.
3. Click on the Create tab. The URI Privilege Configuration page appears:
4. Enter the name of the URI privilege. Use a name that is descriptive of the base URI to be protected. For example, to restrict the creation of documents under a base URI reserved for the accounting group, you might use the name “accounting\_files”.
5. In the action field, enter the base URI to be protected. While the base URI does not have to map to an actual directory, it should follow the directory structure convention (for example, /myfiles/accounting\_files). In this example, only the user with this URI privilege can create a file with the URI /myfiles/accounting\_files/account1.xml.
6. Under the roles section, select the roles that are allowed to create documents under the base URI.
7. Click OK.

The URI privilege is created and added to the security database.

### 15.5.2 Viewing a URI Privilege

Perform the following steps to view a URI privilege:

1. Click the Security icon in the left tree menu.
2. Click the URI Privileges icon. The URI Privileges Summary Page appears:



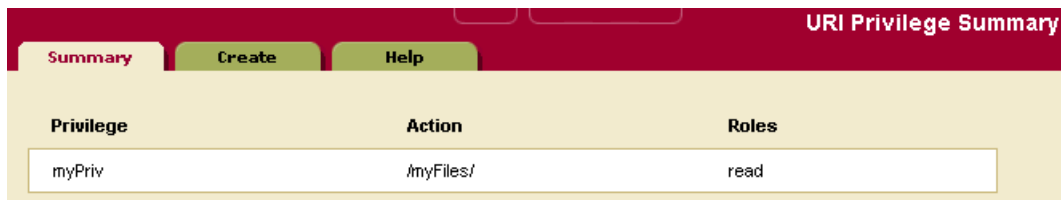
URI Privilege Summary		
Summary	Create	Help
Privilege	Action	Roles
myPriv	/myFiles/	read

3. Click the name of the URI privilege you want to view.
4. View the URI privilege.

### 15.5.3 Deleting a URI Privilege

You can delete a URI privilege from the security database. Perform the following steps to delete a URI privilege:

1. Click the Security icon in the left tree menu.
2. Click the URI Privileges icon. The URI Privileges Summary Page appears:



Privilege	Action	Roles
myPriv	/myFiles/	read

3. On the URI Privileges page for the given privilege, click the Delete button.
4. Click OK to confirm deleting the URI privilege.

The URI privilege is now deleted from the security database.

### 15.6 Amps

An Amp gives the user additional roles temporarily while the user is performing a certain task (executing a function). While the user is executing the “amp-ed” function, the user receives additional privileges and permissions given by the additional roles. An amp is useful when a user needs additional privileges and permissions only while the user is executing a certain function.

Giving the user additional roles permanently could compromise the security of the system. On the other hand, an amp enables granular security control by limiting the effect of the additional roles (privileges and permissions) to a specific function. For example, a user may need a count of all the documents in the database when the user is creating a report. However, the user does not have read permissions on all the documents in the database, and hence does not know the existence of all the documents in the database. An amp can be created for the `document-count()` function to elevate the user to an admin role temporarily while the user is executing the function to count the documents in the system.

An amp is defined by the local name of the function, the namespace and the document URI. The document URI must begin with a forward slash “/” and is treated as being rooted relative to the *Modules* directory in the installation path. When resolving an amp, MarkLogic Server looks for the file using a path rooted relative to the *Modules* directory in the installation path. If it finds a function that matches the local name and namespace using the specified path, it applies the amp to the function.

For more details about amps, see *Understanding and Using Security*. For examples of amps, look at one of the amps created during installation. To view an amp, follow the instructions in the section “Viewing an Amp” on page 129.

Use the following procedures to create, manage and maintain amps:

- [Creating an Amp](#)
- [Viewing an Amp](#)
- [Deleting an Amp](#)

### 15.6.1 Creating an Amp

To create an amp, Perform the following steps:

1. Click the Security icon in the left tree menu.
2. Click the Amps icon.
3. Click on the Create tab. The Amp Configuration page appears:

The screenshot shows the 'New Amp' configuration page within the 'Amp Configuration' window. The window has a red header with 'Amp Configuration' and three tabs: 'Summary', 'Create', and 'Help'. The 'Create' tab is active. The page title is 'New Amp' with 'ok' and 'cancel' buttons. Below the title is a description: 'amp -- A role amplification.' The form contains four fields: 'local name' (text input), 'namespace' (text input), 'document uri' (text input), and 'database' (dropdown menu). Each text input field has a red error message: 'Required. You must supply a value for local-name.', 'Required. You must supply a value for namespace.', and 'Required. You must supply a value for document-uri.' The 'database' dropdown is set to '(filesystem)'. Below the fields is a section titled 'roles -- The roles assigned.' with two checkboxes: 'admin' and 'admin-builtins', both of which are unchecked.

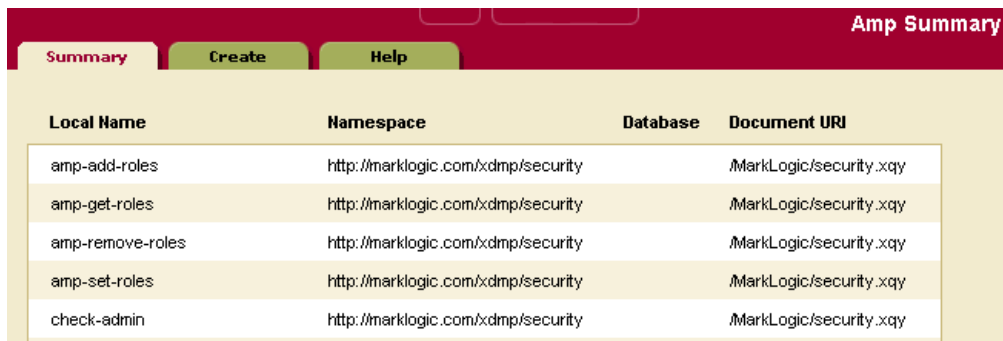
4. Enter the database in which the function is stored. If the function is stored in the *Modules* directory on the filesystem, set the database to `filesystem` (which is the default value).
5. Enter the local name of the function (without parenthesis) in which the amp takes effect. For example: `my-function`.
6. Enter the namespace in which the function is defined.
7. Enter the document URI for the document in which the function is defined. This document URI must begin with a forward slash (for example, `/amped-functions.xqy`). The specified document must be placed in the *Modules* directory within the installation path. For example, if `/mydir/my-amps.xqy` is specified in the document uri, `my-amps.xqy` must be placed in `<installation-directory>/Modules/mydir`.
8. Under the roles section, select the additional roles that will be given to the user while the user is executing the function.
9. Click OK.

The amp is now added to the security database.

## 15.6.2 Viewing an Amp

Perform the following steps to view an amp:

1. Click the Security icon in the left tree menu.
2. Click the Amps icon. The Amps Summary page appears:



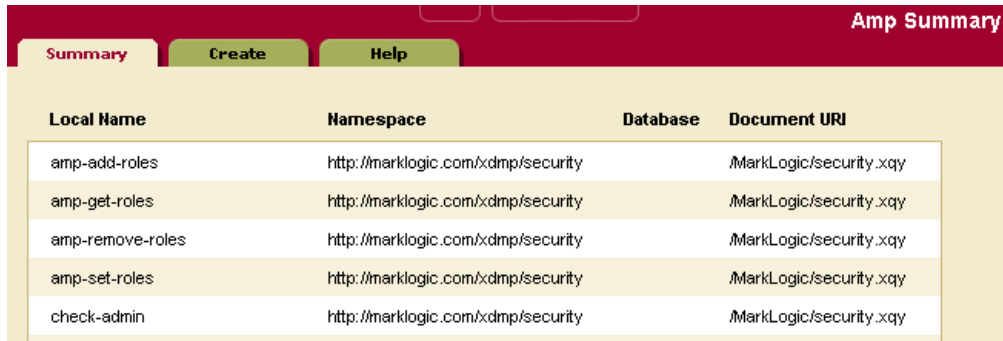
Local Name	Namespace	Database	Document URI
amp-add-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
amp-get-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
amp-remove-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
amp-set-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
check-admin	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy

3. Click on the name of the amp you want to view.
4. View the amp.

### 15.6.3 Deleting an Amp

You can delete an amp from the security database. Perform the following steps to delete an amp:

1. Click the Security icon in the left tree menu.
2. Click the Amps icon. The Amps Summary page appears:



Local Name	Namespace	Database	Document URI
amp-add-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
amp-get-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
amp-remove-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
amp-set-roles	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy
check-admin	http://marklogic.com/xdmp/security		/MarkLogic/security.xqy

3. Click on the name of the amp you want to delete.
4. On the Amp page for the given amp, click the Delete button.
5. Click OK to confirm deleting the amp.

The amp is now deleted from the security database.

## 15.7 Collections

A Collection groups a set of documents that are related and enables queries to target subsets of documents within a database efficiently. A document may belong to any number of collections simultaneously. A collection exists in the system when a document in the system states that it is part of that collection. However, an associated collection object is not created and stored in the security database unless it is protected. A collection created through the Admin Interface is a protected collection and is stored in the security database.

Read, Insert, Update, and Execute capabilities apply for permissions on a collection. A user needs to have read permissions for both the collection and the documents to be able to see the documents in the collection by collection. A user needs to have Update permissions for both the collection and the document to be able to add documents to or remove documents from a collection.

Use the following procedures to create, manage and maintain collections:

- [Creating a Collection](#)
- [Viewing a Collection](#)
- [Removing a Permission from a Collection](#)
- [Deleting a Collection](#)

### 15.7.1 Creating a Collection

Perform the following steps to create a collection:

1. Click the Security icon in the left tree menu.
2. Click the Collections icon.
3. Click the Create tab, The Collection Configuration page appears:

The screenshot shows the 'New Collection' dialog box. At the top, there are tabs for 'Summary', 'Configure', 'Create', and 'Help'. The 'Create' tab is selected. The dialog has a title bar 'Collection Configuration' and two buttons: 'ok' and 'cancel'. The main content area is titled 'New Collection' and contains a text input field for 'uri' with a red error message: 'Required. You must supply a value for uri.' Below this is a section for 'permissions' with three rows of dropdown menus for 'Role Name + Capability' and 'read'.

4. Enter the URI for the collection.

5. In the permissions section, add permissions (role-capability pair) to the collection. Select from the available roles and pick Read, Insert, Update, or Execute capability for the role. If you want to add more than 3 permissions to the role, you can do so from the next screen after you click OK.
6. Click OK.

The protected collection is added to the database.

### **15.7.2 Viewing a Collection**

Perform the following steps to view a collection:

1. Click the Security icon in the left tree menu.
2. Click the Collections icon. The Collection Summary page appears.
3. Click the name of the collection you want to view, either on the tree menu or on the summary page. The Collection Configuration page appears.
4. View the collection.

### **15.7.3 Removing a Permission from a Collection**

Perform the following steps to remove a permission from a collection:

1. Click the Security icon in the left tree menu.
2. Click the Collections icon. The Collection Summary page appears.

3. Click the name of the collection from which you want to remove a permission, either on the tree menu or on the summary page. The Collection Configuration page appears.

The screenshot shows the 'Collection Configuration' page for a collection named 'test'. The page has a red header with the title 'Collection Configuration' and a navigation bar with tabs for 'Summary', 'Configure', 'Describe', 'Create', and 'Help'. The 'Configure' tab is active. Below the header, the collection name 'test' is displayed, along with 'ok' and 'cancel' buttons. The main content area is divided into sections: 'collection -- A collection object.' with a 'delete' button; 'uri' with a text input field containing 'test' and the label 'The collection uri.'; and 'permissions -- Permissions to the collection'. The permissions section contains a table with columns for '[Keep]', 'Role Name (capability)', and an '[add]' button. The first row shows a checked checkbox, 'read (read)', and a dropdown menu with 'read' selected.

4. In the permissions section, uncheck the box next to the permission you want to remove.
5. Click OK.

The permission is removed from the collection.

#### 15.7.4 Deleting a Collection

Perform the following steps to remove delete a collection:

1. Click the Security icon in the left tree menu.
2. Click the Collections icon.

3. Click the name of the collection you want to delete, either on the tree menu or on the summary page. The Collection Configuration page appears.

The screenshot shows the 'Collection Configuration' page for a collection named 'test'. The page has a red header with the title 'Collection Configuration' and a navigation bar with tabs for 'Summary', 'Configure', 'Describe', 'Create', and 'Help'. The 'Configure' tab is active. Below the header, the collection name 'test' is displayed, along with 'ok' and 'cancel' buttons. The main content area is divided into sections: 'collection -- A collection object.' with a 'delete' button; 'uri' with a text input field containing 'test' and the label 'The collection uri.'; and 'permissions -- Permissions to the collection'. The permissions section includes a table with columns for '[Keep]', 'Role Name (capability)', and a checkbox. One entry is shown: a checked checkbox, 'read (read)'. Below this is an '[add]' section with an empty text input field and a dropdown menu set to 'read'.

4. Click on the Delete button near the top right.
5. Click OK to confirm deleting the collection.

The collection is deleted from the security database.

## 15.8 Realm

MarkLogic Server stores the realms for application servers in the security database. Each application server takes its realm from the security database to which it is connected. Realms are used in computing digest passwords.

## 15.8.1 Setting the Realm

The realm stored in the security database to which the Admin Interface is connected is set at installation time:

**Security Setup**

MarkLogic Server has detected that Administration has not been secured. Please supply a user name and password for the Administrative user to set up security.

You also need to specify a realm for this security database. This is the realm that will be displayed to clients authenticating against this database. Since this value is used in password hashes it is recommended that you not change this value once it is set. Please read the further documentation about realms.

<b>Admin</b>	<input type="text" value="admin"/> User/Login name (unique) <b>Required. You must supply a value for user-name.</b>
<b>Password</b>	<input type="password" value="*****"/> Encrypted Password. <b>Required.</b>
<b>Confirm Password</b>	<input type="password" value="*****"/> Encrypted Password. <b>Required.</b>
<b>Realm</b>	<input type="text" value="public"/> The authentication realm.

## 15.8.2 Changing the Realm

Changing the realm in the security database invalidates all user digest passwords. This only affects application servers in digest or digestbasic mode.

In digest mode, you need to re-enter all user passwords in the security database. Changing the passwords in the security database will cause the server to recalculate the digest passwords. In digestbasic mode, the first time a user logs into the server after the realm is changed, the user will be prompted to enter their passwords multiple times before they are logged into the system. However, the server will automatically recalculate their digest password with the new realm at that time, and they will have a normal login process for future access.

To change the realm after installation, Perform the following steps:

1. Click Security in the left tree menu.

2. Click the Configure tab. The Security Configuration page appears.



3. Change the realm to the desired value.
4. Click OK.
5. Click OK again on the confirmation page.

## 16.0 Text Indexing

Before loading documents into a database, you have the option of specifying a number of parameters that will impact how the text components of those documents will be treated. This chapter describes those parameters and includes the following sections:

- [Text Indexes](#)
- [Phrasing and Element-Word-Query Boundary Control](#)
- [Query Behavior with Reindex Settings Enabled and Disabled](#)

Text indexes and phrasing parameters are set on a per-database basis.

### 16.1 Text Indexes

MarkLogic Server allows you to configure, at the database level, which types of text indexes are constructed and maintained during document loading and updating. Each type of index accelerates the performance of a certain type of query. You can specify whether or not each different type of index is maintained for a given database.

**Note:** The index settings are designed to apply to an entire database. If you change any index settings on a database in which documents are already loaded, you must reindex your existing data, either by setting the `reindexer enable` setting to `true` for that database or by reloading the data.

Understanding your likely query set will help you determine which of these index types to maintain. The cost of supporting additional indexes is increased disk space and document load times. As more and more indexes are maintained, document load speed decreases. By default, MarkLogic Server builds a set of indexes that is designed to yield the fast query performance in general usage scenarios.

Text index types are configured on a per-database basis. This configuration should be completed before any documents are loaded into the specified database, although it can be changed later. If you change any index settings on a database in which documents are already loaded, you must reindex your existing data, either by setting the `reindexer enable` setting to `true` for that database or by reloading the data.

In addition to the standard indexes, you can configure indexes on individual elements and attributes in a database. You can create range indexes and/or lexicons on individual elements or attributes in a database. For information on these indexes, see “Element and Attribute Range Indexes and Lexicons” on page 155. You can also create named fields which can explicitly include or exclude specified elements. For details on fields, see “Fields Database Settings” on page 68.

This section describes the text indexes in MarkLogic Server and includes the following subsections:

- [Understanding the Text Index Settings](#)
- [Viewing Text Index Configuration](#)
- [Configuring Text Indexes](#)

### 16.1.1 Understanding the Text Index Settings

The following table describes the different types of indexes available. The indexes are not mutually independent. If both the word search and stemmed search indexes are disabled, the configuration of the remaining indexes is irrelevant, as they all depend on the existence of the word and/or stemmed-search index.

Index	Default Setting	Description
language	en	Specifies the default language for content in this database. Any content without an <code>xml:lang</code> attribute will be indexed in the language specified here. You should have a license key if you specify a non-English language; if you specify a non-english language and do not have a license for that language, the stemming and tokenization will be generic.

Index	Default Setting	Description
stemmed searches	Basic (index is built, each word stems to a single stem)	<p>Enables searches to return relevance ranked results by matching word stems. A word <i>stem</i> is the word that has the same meaning as the specified word, and other words can also have that same stem; therefore, stemmed searches will return more matching results than the exact words specified in the query. A stemmed search for a word finds the exact same terms as well as terms that derive from the same meaning and part of speech as the search term. For example, a stemmed search for <code>run</code> returns results containing <code>run</code>, <code>running</code>, <code>runs</code>, and <code>ran</code>. For details on stemming, see the chapter <a href="#">Understanding and Using Stemmed Searches</a> in the <i>Developer's Guide</i>.</p> <p>There are three types of stemming: basic (one stem per word), advanced (one or more stems per word), and decompounding (advanced plus smaller component words of large compound words).</p> <p>Without either this index or the word searches index, MarkLogic Server is unable to perform relevance ranking and will refuse to execute any <code>cts:word-query()</code>-related built-in function.</p> <p>If both the stemmed search and word search indexes are enabled, MarkLogic Server defaults to performing stemmed searches (unless an unstemmed search is explicitly specified).</p> <p>Turn this index off if you want to disable stemmed searches. If word and stemmed search indexes are both off, then full-text searches are effectively disabled.</p>
word searches (unstemmed)	Off (index is not built)	<p>Enables MarkLogic Server to return relevance ranked results which match exact words in text elements. Either this index or the stemmed search index is needed for MarkLogic Server to execute any <code>cts:word-query()</code>-related function.</p> <p>For many applications, keeping this word search index off and the stemmed search index on is sufficient to return the desired results for queries.</p> <p>Turn this index on if you want to do exact word-only matches. If word and stemmed search indexes are both off, then full-text searches are effectively disabled.</p>
word positions	Off (index is not built)	<p>Speeds up the performance of proximity queries that use the <code>cts:near-query</code> function and of multi-word phrase searches.</p> <p>Turn this index off if you are not interested in proximity queries or phrase searches and if you want to conserve disk space and decrease loading time. If you turn this option on, you might find that you no longer need <code>fast phrase searches</code>, as they have some overlapping functionality.</p>

Index	Default Setting	Description
fast phrase searches	On (index is built)	Accelerates phrase searches by building additional indexes that describe sequences of words at load (or reindex) time. Without this index, MarkLogic Server will still perform phrase searches, just more slowly. Turn this index off if only a small percentage of your queries will contain phrase searches, and if conserving disk space and enhancing load speed is more important than the performance of those queries.
fast case sensitive searches	On (index is built)	Accelerates case sensitive searches by building both case sensitive and case insensitive indexes at load time. Without this index, MarkLogic Server will still perform case sensitive searches, just more slowly. Turn this index off if only a small percentage of your text searches will be case sensitive, and if conserving disk space and enhancing load speed is more important than the performance of those queries.
fast diacritic sensitive searches	On (index is built)	Speeds up diacritic-sensitive searches by eliminating some false positive results. Turn this option off if you do not want to do diacritic-sensitive searches.
fast element word searches	On (index is built)	Accelerates searches that look for words in specific elements by building additional indexes at load time. Without this index, MarkLogic Server will still perform these searches, just more slowly. Turn this index off if only a small percentage of your queries rely on finding words within specific document elements, and if conserving disk space and enhancing load speed is more important than the performance of those queries.
element word positions	Off (index is not built)	Speeds up the performance of proximity queries that use the <code>cts:near-query</code> function in an element and of multi-word element phrase searches. Turn this index off if you are not interested in proximity queries and if you want to conserve disk space and decrease loading time.

Index	Default Setting	Description
fast element phrase searches	On (index is built)	Accelerates phrase searches on elements by building additional indexes that describe sequences of words in elements at load (or reindex) time. Without this index, MarkLogic Server will still perform phrase searches, just more slowly. Turn this index off if only a small percentage of your queries will contain phrase searches at the element level, and if conserving disk space and enhancing load speed is more important than the performance of those queries.
element value positions	Off (index is not built)	Speeds up the performance of proximity queries that use the <code>cts:element-value-query</code> function. Turn this index off if you are not interested in proximity queries and if you want to conserve disk space and decrease loading time.
trailing wildcard searches	Off (index is not built)	Speeds up wildcard searches where the wildcard is at the end of the search pattern (for example, <code>abc*</code> ). The trailing wildcard index is more efficient than the three character index, but it does not speed up queries where the wildcard character is at the beginning of the term.
trailing wildcard word positions	Off (index is not built)	Speeds up the performance proximity queries that use trailing-wildcard word searches, such as wildcard queries that use the <code>cts:near-query</code> function and multi-word phrase searches that contain one or more wildcard terms. Turn this index on if you are using trailing wildcard searches and proximity queries together in the same search.
fast element trailing wildcard searches	Off (index is not built)	Faster wildcard searches with the wildcard at the end of the search pattern within a specific element, but slower document loads and larger database files.
three character searches	Off (index is not built)	Enables wildcard searches where the search pattern contains three or more consecutive non-wildcard characters (for example, <code>abc*x</code> , <code>*abc</code> , <code>a?bcd</code> ). For details on wildcard characters, see the chapter on wildcard searches in the <i>Developer's Guide</i> . When character indexing is turned on in the database, the system will also deliver higher performance for <code>fn:contains()</code> , <code>fn:matches()</code> , <code>fn:starts-with()</code> and <code>fn:ends-with()</code> for most query expressions. Turn this index on if you want to enable wildcard searches that match three or more characters. If you need your wildcard searches to match only two or one characters, then you should enable two character searches and/or one character searches.

Index	Default Setting	Description
three character word positions	Off (index is not built)	Speeds up the performance of proximity queries that use three-character word searches, such as queries that use the <code>cts:near-query</code> function and multi-word phrase searches that contain one or more wildcard terms. Turn this index on if you are using wildcard searches and proximity queries together in the same search.
two character searches	Off (index is not built)	Enables wildcard searches where the search pattern contains two or more consecutive non-wildcard characters. For details on wildcard characters, see the chapter on wildcard searches in the <i>Developer's Guide</i> . When character indexing is turned on in the database, the system will also deliver higher performance for <code>fn:contains()</code> , <code>fn:matches()</code> , <code>fn:starts-with()</code> and <code>fn:ends-with()</code> for most query expressions. Turn this index on if you want to enable wildcard searches that match two or more characters (for example, <code>ab*</code> ).
one character searches	Off (index is not built)	Enables wildcard searches where the search pattern contains only a single non-wildcard character. For details on wildcard characters, see the chapter on wildcard searches in the <i>Developer's Guide</i> . When character indexing is turned on in the database, the system will also deliver higher performance for <code>fn:contains()</code> , <code>fn:matches()</code> , <code>fn:starts-with()</code> and <code>fn:ends-with()</code> for most query expressions. Turn this index on if you want to enable wildcard searches that match one or more characters (for example, <code>a*</code> ).
fast element character searches	Off (index is not built)	Enables searches to return results which match the wildcard characters. Also, speeds up element-based wildcard searches. For details on wildcard characters, see the chapter on wildcard searches in the <i>Developer's Guide</i> . Turn this index on if you want to enable wildcard searches.

Index	Default Setting	Description
word lexicons	Off (index is not built)	Maintains a lexicon of all of the words in a database, with uniqueness determined by a specified collation. If you specify the collation <code>http://marklogic.com/collation/</code> , that specifies the UCA root collation, which is useful for many locales. The lexicon is capitalization-sensitive and diacritic-sensitive (therefore, there is a different entry for <code>Ford</code> and <code>ford</code> ). For details on lexicons, see “Element and Attribute Range Indexes and Lexicons” on page 155 and the chapter on lexicons in the <i>Developer’s Guide</i> . For details on collations, see the <a href="#">Language Support in MarkLogic Server</a> chapter in the <i>Developer’s Guide</i> .
uri lexicon	Off (index is not built)	Maintains a lexicon of all of the URIs used in a database. The URI lexicon speeds up queries that constrain on URIs. It is like a range index of all of the URIs in the database. To access values from the URI lexicon, use the <code>cts:uris</code> or <code>cts:uri-match</code> APIs.
collection lexicon	Off (index is not built)	Maintains a lexicon of all of the collection URIs used in a database. The collection lexicon speeds up queries that constrain on collections. It is like a range index of all of the collection URIs in the database. To access values from the collection lexicon, use the <code>cts:collections</code> or <code>cts:collection-match</code> APIs.

### 16.1.2 Viewing Text Index Configuration

To view text index configuration for a particular database, complete the following procedure:

1. Click on the Databases icon on the left tree menu.
2. Locate the database for which you want to view text index configuration settings, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to view the settings.

4. Scroll down until the text index settings are visible. The following screen shots show the default configuration of text indexing for a database:

<b>language</b>	<input type="text" value="en"/>	The default language assumed for content (if xml:lang encoding is absent)
<b>stemmed searches</b>	<input type="text" value="basic"/>	Enable stemmed word searches (slower document loads and larger database files).
<b>word searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable unstemmed word searches (slower document loads and larger database files).
<b>word positions</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Index word positions for faster phrase and near searches (slower document loads and larger database files).
<b>fast phrase searches</b>	<input checked="" type="radio"/> true <input type="radio"/> false	Enable faster phrase searches (slower document loads and larger database files).
<b>fast case sensitive searches</b>	<input checked="" type="radio"/> true <input type="radio"/> false	Enable faster case sensitive searches (slower document loads and larger database files).
<b>fast diacritic sensitive searches</b>	<input checked="" type="radio"/> true <input type="radio"/> false	Enable faster diacritic sensitive searches (slower document loads and larger database files).

<b>fast element word searches</b>	<input checked="" type="radio"/> true <input type="radio"/> false	Enable faster element-word searches (slower document loads and larger database files).
<b>element word positions</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Index element word positions for faster element-based phrase and near searches (slower document loads and larger database files).
<b>fast element phrase searches</b>	<input checked="" type="radio"/> true <input type="radio"/> false	Enable faster element phrase searches (slower document loads and larger database files).
<b>element value positions</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Index element value positions for faster near searches involving element-value-query (slower document loads and larger database files).

<b>trailing wildcard searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable trailing wildcard searches (slower document loads and larger database files).
<b>trailing wildcard word positions</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Index word positions for trailing-wildcard searches only when trailing-wildcard-searches are enabled (slower document loads and larger database files).
<b>fast element trailing wildcard searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable element trailing wildcard searches (slower document loads and larger database files).
<b>three character searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable wildcard searches and faster character-based XQuery predicates using three or more characters (slower document loads and larger database files).
<b>three character word positions</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Index word positions for three-character searches only when three-character-searches are enabled (slower document loads and larger database files).
<b>two character searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable wildcard searches and faster character-based XQuery predicates using two character (slower document loads and larger database files).
<b>one character searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable wildcard searches and faster character-based XQuery predicates using one character (slower document loads and larger database files).
<b>fast element character searches</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Enable element wildcard searches and element-character-based XQuery predicates (slower document loads and larger database files).

---

<b>word lexicons</b>	<b>[Keep]</b>	<b>Collation URI</b>
	[add]	<input type="text"/>

<b>uri lexicon</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Maintain a lexicon of document URIs (slower document loads and larger database files).
<b>collection lexicon</b>	<input type="radio"/> true <input checked="" type="radio"/> false	Maintain a lexicon of collection URIs (slower document loads and larger database files).

### 16.1.3 Configuring Text Indexes

To configure text indexes for a particular database, complete the following procedure:

1. Click on the Databases icon on the left tree menu.
2. Locate the database for which you want to view text index configuration settings, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to view the settings.
4. Scroll down until the text indexing controls are visible.

5. Configure the text indexes for this database by selecting the appropriate radio buttons for each index type.

Click on the `true` radio button for a particular text index type if you want that index to be maintained. Click on the `false` radio button for a particular text index type if you do not want that index to be maintained.

**Note:** If word searches and stemmed searches are disabled (that is, the `false` radio button is selected for `word searches` and `off` is selected for `stemmed searches`), the settings for the other text indexes are ignored, as explained above.

6. Leave the rest of the parameters unchanged.
7. Scroll to the top or bottom of the right frame and click OK.

The database now has the new text indexing configurations.

## 16.2 Phrasing and Element-Word-Query Boundary Control

MarkLogic Server allows you to specify how XML element constructors impact text phrasing and element-word-query boundaries for searches. This section has the following parts:

- [Phrasing Control](#)
- [Element Word Query Througths](#)
- [Procedures](#)

### 16.2.1 Phrasing Control

By default, MarkLogic Server assumes that any XML element constructor acts as a phrase boundary. This means that phrase searches (for example, searches for sequences of terms) will not match a sequence of terms that contains one or more XML element constructors. Phrasing control lets you specify which XML elements should be transparent to phrase boundaries (for example, a bold or italic element), and which XML elements should be ignored for phrase purposes (for example, footnotes or graphic captions).

For example, consider the following sample XML fragment:

```
<paragraph>
  These two words <italic>are italicized</italic>. The italic element
  <footnote>Elements are defined in the W3C XML standard.</footnote>
  is a standard part of this document's schema.
</paragraph>
```

By default, MarkLogic Server would extract the following five sequences of text for phrase matching purposes (ignoring punctuation and case for simplicity):

- “these two words”
- “are italicized”
- “the italic element”
- “elements are defined in the w3c xml standard”
- “is a standard part of this document's schema”

If you then attempted to match the phrases “words are italicized” or “element is a standard part” against this XML fragment, no matches would be found, because of the embedded XML element constructors.

In fact, a human looking at this XML fragment would realize that the `italic` element should be transparent for phrasing purposes, and that the `footnote` element is a completely independent text container. Seen from this viewpoint, the XML fragment shown above contains only two text sequences (again, ignoring punctuation and case for simplicity):

- “these two words are italicized the italic element is a standard part of this document's schema”
- “elements are defined in the w3c xml standard”

In this case, “words are italicized” and “element is a standard part” would each properly generate a match. But a search for “the w3c xml standard is a standard” would not result in a match.

MarkLogic Server lets you achieve this type of phrasing control by specifying particular XML element names as `phrase-through`, `phrase-around`, and `element-word-query-through` elements:

Type	Definition
<code>Phrase-through</code>	Elements that should not create phrase boundaries (as in the example above, <code>italic</code> should be specified as a <code>phrase-through</code> element).
<code>Phrase-around</code>	Elements whose content should be completely ignored in the context of the current phrase (as in the example above, <code>footnote</code> should be specified as a <code>phrase-around</code> element).

Phrase controls are configured on a per-database basis. You should complete this configuration before loading any documents into the specified database; otherwise, in order for the changes to take effect with your existing content, you must either reload the content or reindex the database after changing the configuration.

## 16.2.2 Element Word Query Throughs

Element-word-query-throughs allow you to specify elements that should be included in text searches that use `cts:element-word-query` on a parent element. For example, consider the following XML fragment:

```
<a>
  <b>hello</b>
  <c>goodbye</c>
</a>
```

If you perform a `cts:element-word-query` on `<a>` searching for the word `hello`, the search does not find any matches in this fragment. The following query shows this pattern:

```
cts:search(fn:doc(), cts:element-word-query(xs:QName("a"), "hello"))
```

This query does not find any matches because `cts:element-word-query` only searches for text nodes that are immediate children of the element `<a>`, not text nodes that are children of any child nodes of `<a>`. Because `hello` is in a text node that is a child of `<b>`, it does not satisfy the `cts:element-word-query`.

If you add an `element-word-query-through` for the element `<b>`, however, then the `cts:element-word-query` on `<a>` searching for the word `hello` returns a match. The `element-word-query-through` on `<b>` causes the text node children of `<b>` behave like the text node children of its parent (in this case, `<a>`).

**Note:** If an element is specified as a phrase-through, then it also behaves as an element-word-query-through, and therefore you do not need to specify it as an element-word-query-through.

## 16.2.3 Procedures

Use the following procedures to configure phrase controls for a particular database:

- [Viewing Phrasing and Element-Word-Query Settings](#)
- [Configuring Phrasing and Element-Word-Query Settings](#)
- [Deleting a Phrasing or Element-Word-Query Setting](#)

### 16.2.3.1 Viewing Phrasing and Element-Word-Query Settings

To view `element-word-query-through`, `phrase-through`, and `phrase-around` settings for a particular database, complete the following procedure in the Admin Interface:

1. Click on the Databases icon on the left tree menu.
2. Locate the database for which you want to view `element-word-query-through`, `phrase-through`, or `phrase-around` settings, either in the tree menu or in the Database Summary table.

3. Click the name of the database for which you want to view the settings.
4. Click the Element-Word-Query-Throughs, Phrase-Throughs, or Phrase-Arounds icon, depending on which one you want to view.
5. The configuration page displays.

The following example shows that the Documents database has been configured with a number of phrase-through elements, including the `<abbr>`, `<acronym>`, `<b>`, `<big>`, `<br>` and `<center>` elements of the XHTML namespace:

The screenshot shows a web-based configuration interface titled "Phrase-Throughs Configuration" for the "Documents" database. It features a navigation bar with "Configure", "Create", and "Help" buttons. The main area is titled "Database: Documents" and contains a "phrase throughs" section with a description: "The phrase-through specifications." Below this is a "phrase through" section with the description "Phrases may cross these markup boundaries." and a "drop" button. The "phrase through" section is expanded to show two fields: "namespace uri" with the value "http://www.w3.org/1999/xhtml" and a note "A namespace URI.", and "localname" with the value "abbr,acronym,b,big,br,center,cite,code,dl" and a note "One or more localnames." At the bottom of the dialog are "ok" and "cancel" buttons.

### 16.2.3.2 Configuring Phrasing and Element-Word-Query Settings

To configure element-word-query-through, phrase-through, and phrase-around settings for a particular database, perform the following procedure in the Admin Interface:

1. Click the Databases icon in the left tree menu.
2. Locate the database for which you want to configure element-word-query-through, phrase-through, or phrase-around settings, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to configure the settings.

- Click the Element-Word-Query-Throughs, Phrase-Throughs, or Phrase-Arounds icon, depending on which one you want to configure.

**Note:** The remainder of this procedure will assume that you have chosen to configure phrase-through settings. If you wish to configure phrase-around or element-word-query-through settings, the steps are completely analogous, once you have clicked on the corresponding icon.

- Click the Create tab at the top right.

The Phrase-Throughs Configuration page displays:

- Enter the namespace URI of the XML element that you are specifying as a phrase-through element.

Every XML element is associated with a namespace. For the phrase-through setting to be precise, you must specify the namespace of the XML element. Leaving the namespace URI field blank specifies the universal unnamed namespace.

Alternatively, you can specify that the element is namespace independent by putting an asterisk (\*) in the namespace URI field.

- Enter the element name in the localname field.

The local name is the name of the XML element that you are specifying as a phrase-through element. If you want to specify more than one element that is associated with the specified namespace, you can provide a comma-separated list of element names.

8. Repeat step [5](#) – step [6](#) for each phrase-through element as needed.
9. Scroll to the top or bottom and click OK.

The new phrase-through is added.

**Note:** If you change the element-word-query-through, phrase-through, or phrase-around settings for a particular database after documents have already been loaded, you should reindex your existing data, either by setting the `reindexer enable` setting to `true` for that database or by reloading the data.

### 16.2.3.3 Deleting a Phrasing or Element-Word-Query Setting

To delete an element-word-query-through, phrase-through, or phrase-around setting for a particular database, perform the following procedure in the Admin Interface:

1. Click the Databases icon in the left tree menu.
2. Locate the database for which you want to delete element-word-query-through, phrase-through, or phrase-around settings, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to delete the settings.
4. Click the Element-Word-Query-Throughs, Phrase-Throughs, or Phrase-Arounds icon, depending on which one you want to delete.

The appropriate configuration page displays.

5. Scroll down to the element that you want to delete.
6. Click the Drop button next to the element that you want to delete.

A confirmation message displays.

7. Confirm the delete operation and click OK.

The Phrase-Through or Phrase-Around element is deleted from the database.

**Note:** If you change the element-word-query-through, phrase-through, or phrase-around settings for a particular database after documents have already been loaded, you should reindex your existing data, either by setting the `reindexer enable` setting to `true` for that database or by reloading the data.

## 16.3 Query Behavior with Reindex Settings Enabled and Disabled

When you load a document into a database, it is indexed based on the index settings at the time of the load. When you issue a query to a database, it is evaluated based on a consistent view of the index settings. This consistent view might not include all of the index features that are enabled in the database. This section describes the behavior of queries at various index-setting states of the database, and includes the following parts:

- [Understanding the Reindexer Enable Settings](#)
- [Query Evaluation According to the Lowest Common Denominator](#)
- [Reindexing Does Not Apply to Point-In-Time Versions of Fragments](#)
- [Example Scenario](#)

### 16.3.1 Understanding the Reindexer Enable Settings

At the database level, you can enable or disable automatic reindexing by setting the `reindexer enable` setting to `true` or `false` for that database. When the reindexer is enabled, any index or fragment changes to the database settings will cause all documents in the database that are not indexed/fragmented according to the settings to initiate a reindex operation. Note the following about the database settings and the reindex operation:

- When reindexing is enabled, the reindex operation runs as a background task. You can set a higher or lower priority on the reindexing task by increasing or decreasing the setting of the `reindexer throttle`.
- Any new documents added to or updated in the database will get the new database settings. This is true both with reindexing enabled and with reindexing disabled.
- After changing index or fragmentation settings in a database, because new or modified documents get the new settings, the database can get into a state where some documents are indexed/fragmented differently from other documents in the database.
- After changing index or fragmentation settings in a database in which reindexing is enabled, the old documents are reindexed according to the new settings, but the new settings do not take effect for queries until the reindex operation has completed and all documents are indexed to the state matching the database settings.
- After changing index or fragmentation settings in a database in which reindexing is disabled, new and changed documents get the current settings, but queries will not take advantage of the new settings until all documents in the database match the database settings.

### 16.3.2 Query Evaluation According to the Lowest Common Denominator

When queries are evaluated, they use the index settings that are calculated for the database at a given time. The current index settings for a query are determined at the time of query evaluation, and are based on the lowest common denominator of (that is, the index/fragmentation settings that are the least of) the following:

- The index/fragmentation settings defined in the database configuration.
- The actual index/fragmentation of documents/fragments in the database.

At any given time, the current lowest common denominator is invalidated upon the following events:

- system startup
- a change to the database configuration settings
- when a reindexing operation completes

If the lowest common denominator is invalidated, it is recalculated the next time a query is issued against the database.

The net impact is that, when index/fragmentation settings have changed on a database after any data is loaded, queries cannot take advantage of the new settings until the new settings meet the lowest common denominator criteria. Depending on the types of index setting changes you make, this can cause queries that behaved one way before index settings were changed to behave differently after the changes. The next section provides a sample scenario to help illustrate this behavior.

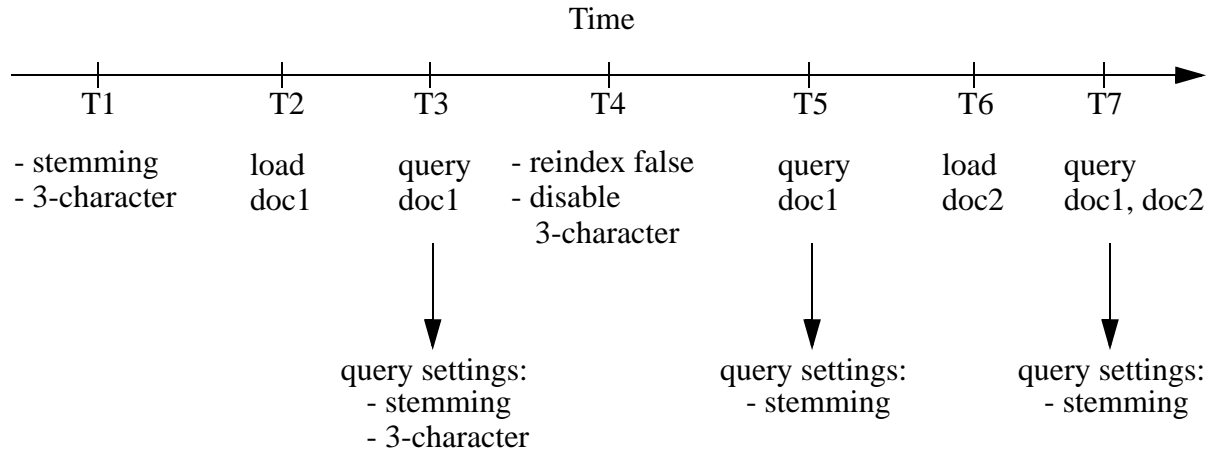
### 16.3.3 Reindexing Does Not Apply to Point-In-Time Versions of Fragments

If you have set a `merge timestamp` on the database to retain older versions of fragments for point-in-time queries, the older versions of the fragments will retain the indexing properties of the database at the time when they were updated. Because of this, reindexing a database that uses point-in-time queries can cause unpredictable query results. Mark Logic recommends that you do not reindex a database that has the `merge timestamp` parameter set to anything but 0. For details on point-in-time queries, see the “Point-In-Time Queries” chapter in the *Developer’s Guide*. For details on setting the `merge timestamp` parameter, see “Merges and Point-In-Time Queries” on page 86.

### 16.3.4 Example Scenario

This section describes a simple scenario showing the effect of changing index settings on query behavior over time.

The following figure shows how changing the index settings can effect queries that initiate after index setting changes occur.



In this scenario, the query issued at time T3 sees the `doc1` document with stemming and 3-character wildcard indexes enabled. Wildcard queries such as `abc*` will be successful. The same wildcard query at time T5, however, will not be successful, because the 3-character index (which is required for the `abc*` query) was disabled at time T4. Note that the document `doc1` is actually indexed with 3-character and stemming, but the query at time T5 only is able to use the stemming index. At time T7, the database has `doc1` indexed with both stemming and 3-character indexes, but `doc2` only has the stemming index. With reindexing disabled, the query at T7 will use the lowest common denominator, which is in this case stemming.

## 17.0 Element and Attribute Range Indexes and Lexicons

MarkLogic Server allows you to create, at the database level, indexes and lexicons on elements and attributes according to their QNames. This chapter describes these range indexes and lexicons. The following sections are included:

- [Understanding Element and Attribute Range Indexes](#)
- [Using Range Indexes for Element and Attribute Value Lexicons](#)
- [Understanding Element and Attribute Word Lexicons](#)
- [Viewing Element Range Index or Element Word Lexicon Settings](#)
- [Defining Element Range Indexes or Element Word Lexicons](#)
- [Viewing Attribute Range Index and Attribute Word Lexicon Settings](#)
- [Defining Attribute Range Indexes or Attribute Word Lexicons](#)
- [Defining Element or Attribute Value Lexicons](#)
- [Deleting Range Indexes or Lexicons](#)

### 17.1 Understanding Element and Attribute Range Indexes

MarkLogic Server maintains a universal index for every database to rapidly search the text, structure, and combinations of the text and structure that are found within collections of XML documents.

In some cases, however, XML documents can incorporate numeric or date information. Queries against these documents may include search conditions based on inequalities (for example, `price < 100.00` or `date ≥ thisQtr`). Specifying range indexes for these elements and/or attributes will substantially accelerate the evaluation of these queries.

Defining a range index on an element or attribute also allows you to use the range query constructors (`cts:element-range-query` and `cts:element-attribute-range-query`) in `cts:search` operations, making it easy to compose complex range-query expressions to use in searches. For details, see the [Using Range Queries in cts:query Expressions](#) chapter in the *Developer's Guide*.

Similarly, you can create range indexes on elements or attributes of type `xs:string`, and these indexes can accelerate the performance of queries that sort by the string values, and are also used for lexicon queries (see “Understanding Element and Attribute Word Lexicons” on page 158).

If you specify a range index on an element, and if you have elements of that name that have complex content (for example, elements with child elements), the content is indexed based on a casting of the element to the specified type of the range index. For example, if you specify a range index of type `xs:string` on an element named `h1`, then the following element:

```
<h1>This is a <b>bold</b> title.</h1>
```

is indexed with the value of `This is a bold title`, which is the value returned by casting the `h1` element to `xs:string`. This behavior allows you to index values of complex elements without pre-processing the content.

Also, range indexes can improve the performance of queries that sort the results using an `order by` clause and return a subset of the data (for example, the first ten items). For details on this order by optimization using range indexes, see [Optimizing Order By Expressions With Range Indexes](#) in the *Query Performance and Tuning* guide.

MarkLogic Server supports range indexes for both elements and attributes across a wide spectrum of XML data types. For the most part, this list conforms to the XML totally ordered data types:

Type	Description
<code>int</code>	Positive and negative integers
<code>unsignedInt</code>	Positive integers (including 0)
<code>long</code>	Large positive and negative integers
<code>unsignedLong</code>	Large positive integers (including 0)
<code>float</code>	32-bit floating point numbers
<code>double</code>	64-bit floating point numbers
<code>decimal</code>	Large floating point numbers
<code>dateTime</code>	Combined date and time
<code>time</code>	Time (including timezone)
<code>date</code>	Full date (year, month, day)
<code>gYearMonth</code>	Year and month only
<code>gYear</code>	Year only
<code>gMonth</code>	Month only
<code>gDay</code>	Day only
<code>yearMonthDuration</code>	Duration of years and months
<code>dayTimeDuration</code>	Duration of days and time
<code>string</code>	String character data
<code>anyURI</code>	A URI string

It is important to note that the date and time types listed above adhere to the XML specification for dates and times. At present, other date and time formats are not supported by MarkLogic Server range indexes. For a more detailed description of the definition of these data types, consult the W3C XML Schema documents.

Range indexes must be specified manually using the Admin Interface. Specifying a range index requires an element name, a namespace for that element name, and the data type found in that element.

Range indexes are constructed during the document loading process, and are automatically kept in sync through subsequent updates to indexed data. Consequently, element and attribute range indexes should be specified for a database before any XML documents containing those elements and/or attributes are loaded into that database, otherwise the content must be either reindexed or reloaded to take advantage of the new range indexes.

When creating range index with a scalar type of string (`xs:string`), you specify a collation as well as the element/attribute QNames. The collation specifies the unique ordering for the string values. You can have multiple range indexes on the same element or attribute with different collations; that is, the collation is part of the unique identifier for the string range index. For details about collations, see the [Language Support in MarkLogic Server](#) chapter in the *Developer's Guide*.

Range indexes use disk space and consume memory. That is the trade-off for improved performance. Additionally, if you have a large amount of range index data and if your system is updated regularly, you might need to increase the size of your journals. For details on the database journal settings, see “Memory and Journal Settings” on page 51.

## 17.2 Using Range Indexes for Element and Attribute Value Lexicons

In addition to speeding up sorting and comparison queries, MarkLogic Server uses range indexes to resolve element and attribute value lexicon queries. These are queries that use the following search APIs:

- `cts:element-attribute-values`
- `cts:element-attribute-value-match`
- `cts:element-values`
- `cts:element-value-match`

In order to use any of these APIs, you must create a range index of type `xs:string` on the element(s) and/or attribute(s) specified in the query.

### 17.3 Understanding Element and Attribute Word Lexicons

MarkLogic Server allows you to create a word lexicon that is restricted to a particular element or attribute. The element word lexicon and the attribute word lexicon store all of the unique words that are stored in the specified element or attribute. The words are stored case-sensitive and diacritic sensitive, so the words `Ford` and `ford` would be separate entries in the lexicon. To use the element or attribute word lexicons, use the following search APIs:

- `cts:element-attribute-words`
- `cts:element-attribute-word-match`
- `cts:element-words`
- `cts:element-word-match`

### 17.4 Viewing Element Range Index or Element Word Lexicon Settings

To view the range index or lexicon that will be applied to documents as they are loaded or reindexed, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Locate the database whose range index or lexicon you want to view, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to view the range index or lexicon.
4. Click the Element Range Indexes icon or the Element Word Lexicons icon.

The Element Index or the Element Word Lexicon Configuration page displays.

### 17.5 Defining Element Range Indexes or Element Word Lexicons

To define a range index or lexicon for an element, perform the following steps:

1. Click the Databases icon on the left frame.
2. Locate the database for which you want to create a range index or lexicon, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to create a range index or lexicon.
4. Click the Element Range Indexes icon or the Element Word Lexicons icon in the tree menu, under the selected database.

- Click the Add tab. The Element Range Index Configuration page or the Element Word Lexicon Configuration page displays:

- If you are creating a Range index, select the type of the XML element for which you want to build a range index (this field does not appear if you are creating an element word lexicon).
- If the index is of type `xs:string`, a collation box appears with a default collation. If you want the index to use a different collation than the default, enter the collation URI. For details about collations, see the [Language Support in MarkLogic Server](#) chapter in the *Developer's Guide*.
- Enter the namespace URI of the XML element.
 

Every XML element is associated with a namespace. For the description of the element to be precise, you must specify the namespace of the XML element. The asterisk (\*) cannot be used to indicate namespace independence. Leaving the namespace URI field blank specifies the universal unnamed namespace.
- Enter the element name in the localname field.
 

The local name is the name of the XML element to be indexed. If you have more than one element of the same type in the same namespace that you want to index, you can provide a comma-separated list of element names.
- Repeat step 5 – step 9 for each index as needed.

11. Scroll to the top or bottom and click OK.

The new element range index or element word lexicon is added to the database. These rules are applied to XML documents loaded into the specified database from this point on.

**Note:** If you have reindexing enabled for the database and you specify an element that exists in a document, reindexing will run in the background. When the reindexing is complete, the new index will become available to queries.

## 17.6 Viewing Attribute Range Index and Attribute Word Lexicon Settings

To view the range index or lexicon that will be applied to documents as they are loaded or reindexed, perform the following steps:

1. Click the Databases icon on the left frame.
2. Locate the database for which you want to view a range index or lexicon, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to view a range index or lexicon.
4. Click the Attribute Range Indexes icon or the Attribute Word Lexicons icon in the tree menu, under the selected database.

The Attribute Range Index Configuration page or the Element-Attribute Word Lexicon page displays.

## 17.7 Defining Attribute Range Indexes or Attribute Word Lexicons

To define a range index or lexicon for an attribute of a particular element, perform the following steps:

1. Click the Databases icon on the left frame.
2. Locate the database for which you want to create an index or lexicon, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to create an index or lexicon.
4. Under the selected database, click the Attribute Range Indexes icon in the tree menu for an attribute range index, or the Attribute Word Lexicon icon for an attribute word lexicon.

- Click the Add tab. The appropriate Add Index Configuration page displays:

- If you are creating a Range index, select the type of the XML attribute for which you want to build an attribute range index (this field does not appear if you are creating an attribute word lexicon).
- If the index is of type `xs:string`, a collation box appears with a default collation. If you want the index to use a different collation than the default, enter the collation URI. For details about collations, see the [Language Support in MarkLogic Server](#) chapter in the *Developer's Guide*.
- Enter the namespace URI of the XML element that contains the attribute you want to index into the parent namespace URI field.

Every XML element is associated with a namespace. For the description of the element to be precise, you must specify the namespace of the XML element. The asterisk (\*) cannot be used to indicate namespace independence. Leaving the namespace URI field blank specifies the universal unnamed namespace.

- Enter the element name in the parent localname field.

The local name is the name of the XML element that contains the attribute to be indexed. If you have more than one element in the same namespace that contains the attribute you want to index, you can provide a comma-separated list of element names.

10. Enter the namespace URI of the attribute that you want to index into the namespace URI field.

Every XML attribute is associated with a namespace. For the description of the attribute to be precise, you must specify the namespace of the XML attribute. The asterisk (\*) cannot be used to indicate namespace independence. Leaving the namespace URI field blank specifies the universal unnamed namespace.

11. Enter the attribute name in the localname field.

The local name is the name of the XML attribute to be indexed. If you have more than one attribute in the same namespace within the specified parent element(s) that you want to index, you can provide a comma-separated list of attribute names.

12. Repeat step [5](#) – step [10](#) for each attribute index as needed.
13. Scroll to the top or bottom and click OK.

The new attribute index or attribute word lexicon is added to the database. These rules are applied to XML documents loaded into the specified database from this point on.

**Note:** If you have reindexing enabled for the database and you specify an element-attribute pair that exists in a document, reindexing will run in the background. When the reindexing is complete, the new index will become available to queries.

## 17.8 Defining Element or Attribute Value Lexicons

Element and Attribute value lexicons are implemented using range indexes of type `xs:string` on the element(s) and/or attribute(s) specified in the query. Therefore, to create a value lexicon, you create a range index (element index or attribute index) of type `xs:string` for the specified element(s) and/or attribute(s).

## 17.9 Deleting Range Indexes or Lexicons

To delete element or attribute indexes or lexicons for a specific database, perform the following steps:

1. Click the Databases icon on the left frame.
2. Locate the database for which you want to delete a range index or lexicon, either in the tree menu or in the Database Summary table.
3. Click the name of the database for which you want to delete a range index or lexicon.
4. Determine whether you need to delete an element range index, an attribute range index, an element word lexicon, or an attribute word lexicon.

5. Click the Element Range Index icon, Attribute Range Index icon, Element Word Lexicon icon, or the Attribute Word Lexicon icon. The configuration page for the appropriate index appears.
6. Locate the index you want to delete and click Delete.
7. A confirmation message displays. Confirm the delete and click OK.

The index or lexicon is deleted from the database.

## 18.0 Fragments

Before loading data into a database, you have the option of specifying how XML documents will be partitioned for storage into smaller blocks of information called fragments. For large source XML documents, size can become an issue, and using fragments help manage performance of your system. In general, fragments for XML documents should be sized between 10K and 100K. Fragments set too small or too big can slow down performance, so proper fragment sizing is important.

The actual fragmentation of an XML document is completely transparent to the application programmer. At the XQuery program level, the document appears to be a single integral structure, regardless of how it is stored and managed as fragments on disk. Fragmentation is an application-transparent tuning mechanism.

However, fragmentation *does* impact relevance ranking. The relevance-ranking algorithm considers both term frequency within a target piece of content and overall term frequency within the database to rank results by relevance. Rather than consider term frequency across the entire XML document for ranking purposes, MarkLogic Server considers term frequency within the individual fragment (and its descendants) being ranked. Consequently, different fragmentation strategies may impact relevance rankings—particularly in situations when a single fragment may straddle multiple XML structures that you are trying to differentiate on a relevance basis.

With MarkLogic Server, you specify fragmentation *rules* that are used to partition your XML documents. These rules are applied one document at a time. However, fragmentation rules are specified at the database level—on the assumption that databases contain many documents with similar structures where the same fragmentation rules should be applied.

Fragmentation rules are applied to documents during document loads, updates, and database reindexing. Specifying additional fragmentation rules after documents have been loaded causes future updates and/or reindexing of those documents to use the new fragmentation rules, but does not change the fragmentation of existing documents (if `reindex enable` is set to `true`, however, the documents will eventually be reindexed and take on the new fragmentation policy). As a result, if you want to change the fragmentation rules for already loaded content, you will have to reload your documents or reindex the database so that your new fragmentation rules can take effect.

Use the following procedures for managing fragmentation rules:

- [Choosing a Fragmentation Strategy](#)
- [Defining Fragment Roots](#)
- [Defining Fragment Parents](#)
- [Viewing Fragment Rules](#)
- [Deleting Fragment Rules](#)

## 18.1 Choosing a Fragmentation Strategy

Proper fragmentation is important to performance. Before you specify how to fragment the XML data being loaded, you need to plan your fragmentation strategy. Apply the following guidelines:

- Fragments are described generically using XML element names.
- Fragments for XML documents should be between 10K and 100K in size (these are just general guidelines; in some situations, larger or smaller fragment sizes can work fine, and there are many factors that will affect performance for a given fragment size including disk block size, how many fragments are in the database, how often fragments are accessed, the types of queries used in the application, and so on).
- Fragments can be (and in many cases, should be) nested hierarchically.
- Smaller fragment sizes allow more efficient element-level updates in the database, but excessively small fragments can slow down both loading speed and query performance.
- Larger fragment sizes can also slow down query performance by requiring excessive loading of data from disk in resolving queries.
- In general, within the size range set above, larger fragment sizes deliver higher-performance overall than smaller fragment sizes.
- Binary and text documents must fit in a single fragment. Therefore, set the database `in memory tree size` parameter to 1 to 2 MB larger than your largest binary or text file.

After you decide how to fragment your data, you can use either of the following methods:

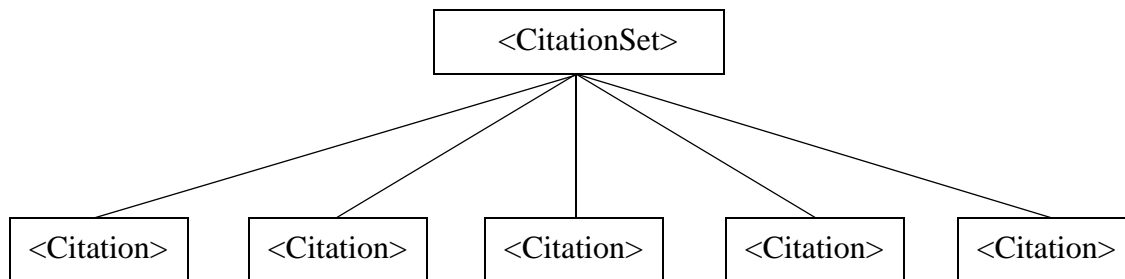
- [Fragment Roots](#)
- [Fragment Parents](#)

Both methods turn your fragmentation strategy into concrete rules for the system.

### 18.1.1 Fragment Roots

If a document contains many instances of an XML structure that share a common element name, then these structures make sensible fragments. With MarkLogic Server, you can use this common element name as a fragment root.


The following diagram shows an XML document rooted at `<CitationSet>` that contains many instances of a `<Citation>` node. Each `<Citation>` node contains further XML and averages between 15K and 20K in size. Based on this information, `<Citation>` is a sensible element to use as a fragment root:



### 18.1.2 Fragment Parents

If your document contains many different XML substructures, each of which is a good candidate to be a fragment, then it would be time consuming to specify each substructure as a fragment root. Instead, you can specify fragments by setting the parent of these substructures to be a fragment parent—so that every substructure under this parent becomes a separate fragment, regardless of its name.

The following diagram shows a document with substructures of different names:

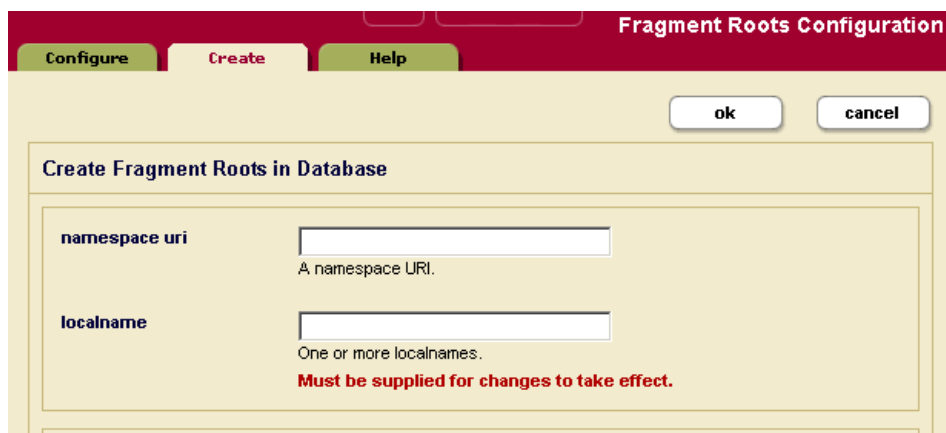


In this case, you can use the `<Products>` element as a fragment parent, and the `<Books>`, `<Movies>`, `<Music>`, `<Games>` and `<Toys>` children automatically become fragments.

## 18.2 Defining Fragment Roots

To define a rule for a fragment root, complete the following procedure:

1. Click the Databases icon on the left tree menu.
2. Determine the database for which you are specifying a new fragment rule.
3. Click the icon for this database, either in the tree menu or the Database Summary page.
4. Click the Fragment Roots icon.
5. Click the Create tab. The Fragment Roots Configuration page displays:



6. Enter the namespace URI of the XML element that you are using as a rule for the fragment root.

Every XML element is associated with a namespace. For the fragment rule to be precise, you must specify the namespace of the XML element. Leaving the namespace URI field blank specifies the universal unnamed namespace.

Alternatively, you can specify that the rule for the fragment root is namespace independent by putting an asterisk (\*) in the namespace URI field.

7. Enter the element name in the localname field.

The local name is the name of the XML element used as the root of a fragment. If you have more than one fragment root rule associated with the specified namespace, you can provide a comma-separated list of element names.

8. Repeat step [6](#) – step [7](#) for each rule for a fragment root as needed.

9. Scroll to the top or bottom and click OK.

The new fragment root rules are added to the database. These rules are applied to XML documents loaded into the specified database from this point on.

### 18.3 Defining Fragment Parents

To define a rule for a fragment parent, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Determine the database for which you are specifying a new fragment parent.
3. Click the icon for this database, either in the tree menu or the Database Summary page.
4. Click the Fragment Parents icon.
5. Click the Create tab. The Fragment Parents Configuration page displays:

6. Enter the namespace URI of the XML element that you are using as a rule for the fragment parent.

Every XML element is associated with a namespace. For the fragment rule to be precise, you must specify the namespace of the XML element. Leaving the namespace URI field blank specifies the universal unnamed namespace.

Alternatively, you can specify that the rule for the fragment root is namespace independent by putting an asterisk (\*) in the namespace URI field.

7. Enter the element name in the localname field.

The local name is the name of the parent XML element whose children will be fragment roots. If you have more than one fragment parent rule associated with the specified namespace, you can provide a comma-separated list of element names.

8. Repeat step [6](#) – step [7](#) for each fragment parent as needed.
9. Scroll to the top or bottom and click OK.

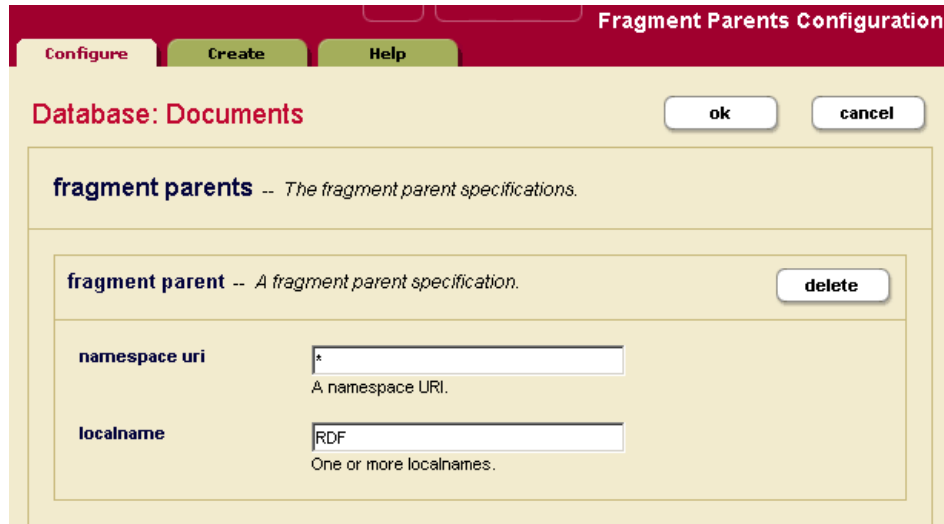
The new fragment rules are added to the database. These rules are applied to XML documents loaded into the specified database from this point on.

## 18.4 Viewing Fragment Rules

To view fragment rules that are in effect, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Locate the database whose fragment rules you want to view, either in the tree menu or the Database Summary page.
3. Click the icon for this database.
4. Determine whether to view the rules for the fragment root or fragment parent.
5. Click either the Fragment Roots icon or Fragment Parents icon, under the specified database.

The following example shows that the Documents database has only one rule defined for a fragment parent. The rule states that any direct child of an `<RDF>` element, regardless of the namespace for the `<RDF>` element, should form the root of a fragment:



## 18.5 Deleting Fragment Rules

To delete fragment rules for a specific database, perform the following steps:

1. Click the Databases icon on the left tree menu.
2. Locate the database that contains the fragment rules you want to delete, either in the tree menu or the Database Summary page.
3. Click the icon for this database.
4. Determine whether you need to delete a rule for a fragment root or fragment parent.
5. Click either the Fragment Roots icon or Fragment Parents icon, under the specified database.
6. Locate the fragment rule you want to delete and click Delete.
7. A confirmation message displays. Confirm the delete and click OK.

The fragment rule is dropped from the database.

**Note:** Deleting fragment rules has no impact on the fragmentation that has already been applied to documents loaded into the database, unless reindexing is enabled for the database.

## 19.0 Namespaces

Namespaces are a powerful mechanism used to differentiate between potentially ambiguous XML elements. Namespaces can be defined within individual XQuery programs. They can also be defined using the Admin Interface.

Namespaces can be defined for a group to apply to all HTTP, XDBC, and WebDAV servers in a group or for a particular HTTP, XDBC, or WebDAV server. Namespaces cannot be defined through the Admin Interface only to apply to a particular forest, database, or XQuery program.

For more information about namespaces, see the “Namespaces” chapter in *Introduction to XQuery*, which provides a detailed description of XML namespaces and their use. Be sure to review this information before using the Admin Interface to manage your namespaces.

Use the following procedures for managing namespaces in the Admin Interface:

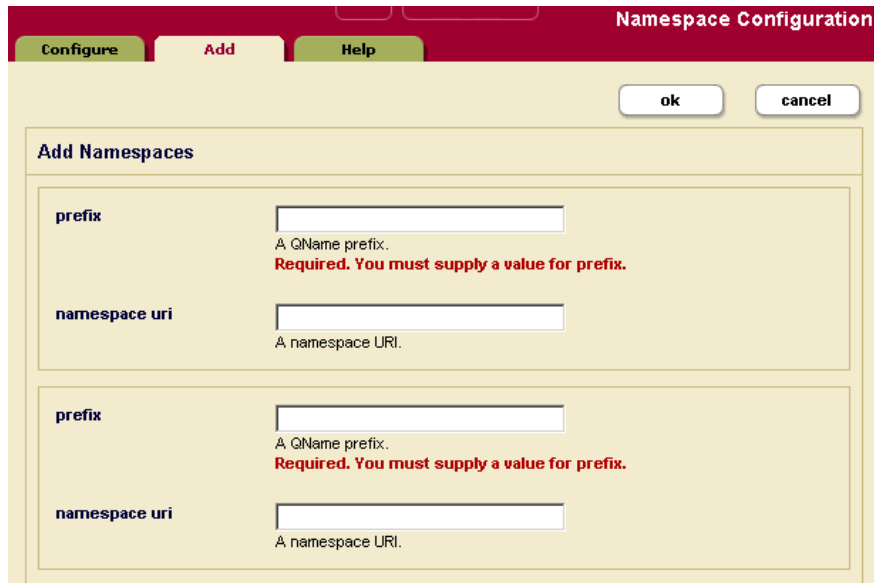
- [Defining Namespaces for a Group](#)
- [Defining Namespaces for an HTTP or XDBC Server](#)
- [Viewing Namespace Settings for a Group](#)
- [Viewing Namespace Settings for an HTTP or XDBC Server](#)
- [Deleting Namespaces for a Group](#)
- [Deleting Namespaces for an HTTP or XDBC Server](#)

### 19.1 Defining Namespaces for a Group

To define namespaces using the Admin Interface for a group, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click the group in which you want to define the namespace, either in the tree menu or the Groups Summary page.
3. Click the Namespaces icon on the left tree menu, under the group name.

4. Click the Add tab. The Namespaces Configuration page displays:



The screenshot shows the 'Namespace Configuration' window with the 'Add' tab selected. The window title is 'Namespace Configuration'. At the top, there are three tabs: 'Configure', 'Add', and 'Help'. Below the tabs are 'ok' and 'cancel' buttons. The main content area is titled 'Add Namespaces' and contains two identical form sections. Each section has a 'prefix' label and a text input field. Below the 'prefix' field, it says 'A QName prefix.' and 'Required. You must supply a value for prefix.' in red text. Below the 'prefix' field is a 'namespace uri' label and a text input field. Below the 'namespace uri' field, it says 'A namespace URI.'

5. Enter a prefix for your namespace.
6. Enter a URI for your namespace.  
  
If you are defining a prefix for the universal unnamed namespace, leave the URI blank.
7. Repeat step [5](#) – step [6](#) for each namespace as needed.
8. Scroll to the top or bottom and click OK.

The namespace is now defined in the group.

## 19.2 Defining Namespaces for an HTTP or XDBC Server

To define namespaces using the Admin Interface for an HTTP or XDBC Server, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click the group which contains the HTTP or XDBC server for which you want to define the namespace, either in the tree menu or the Groups Summary page.
3. Click the HTTP Servers or XDBC Servers icon as appropriate.
4. Click on the name of the HTTP or XDBC server for which you want to define the namespace.

- Click on the Namespaces icon on the left tree menu, under the specified HTTP or XDBC server.
- Click the Add tab at the top right. The Namespaces Configuration page displays:

The screenshot shows a dialog box titled "Namespace Configuration" with three tabs: "Configure", "Add", and "Help". The "Add" tab is selected. At the top right of the dialog are "ok" and "cancel" buttons. The main content area is titled "Add Namespaces" and contains two identical forms. Each form has a "prefix" field with a text input, a label "A QName prefix.", and a red error message "Required. You must supply a value for prefix." Below it is a "namespace uri" field with a text input and a label "A namespace URI."

- Enter a prefix for your namespace.
- Enter a URI for your namespace.  
If you are defining a prefix for the universal unnamed namespace, leave the URI blank.
- Repeat step [7](#) – step [8](#) for each namespace as needed.
- Scroll to the top or bottom and click OK.

The namespace is now defined for the HTTP or XDBC Server.

### 19.3 Viewing Namespace Settings for a Group

To view namespaces you have defined in the Admin Interface, perform the following steps:

- Click the Groups icon on the left tree menu.
- Click the group which contains the namespace you want to view, either in the tree menu or the Groups Summary page.

3. Click the Namespaces icon on the left tree menu, under the specified group. The Namespace Configuration page appears.

The screenshot shows a dialog box titled "Namespace Configuration". At the top, there are three tabs: "Configure" (selected), "Add", and "Help". To the right of the tabs are "ok" and "cancel" buttons. Below the tabs, the text reads "namespaces -- The namespace binding specifications." Below this is a section for editing a namespace, titled "namespace -- A namespace binding specification." with a "delete" button. This section contains two input fields: "prefix" with the value "ml" and a description "A QName prefix.", and "namespace uri" with the value "http://marklogic.com/ml" and a description "A namespace URI." At the bottom of the dialog are "ok" and "cancel" buttons.

## 19.4 Viewing Namespace Settings for an HTTP or XDBC Server

To view namespaces you have defined in the Admin Interface, perform the following steps:

1. Click the Groups icon on the left menu tree.
2. Click the group which contains the HTTP or XDBC server for which you want to view the namespace, either in the tree menu or the Groups Summary page.
3. Click the HTTP Servers or XDBC Servers icon as appropriate.
4. Click on the name of the HTTP or XDBC server for which you want to view the namespace.

5. Click the Namespaces icon on the left tree menu, under the specified HTTP or XDBC server. The Namespace Configuration page appears.

The screenshot shows a 'Namespace Configuration' dialog box. At the top, there are three tabs: 'Configure' (selected), 'Add', and 'Help'. Below the tabs are 'ok' and 'cancel' buttons. The main area contains a list of namespaces. The first entry is 'namespaces -- The namespace binding specifications.' Below it is a single namespace entry: 'namespace -- A namespace binding specification.' To the right of this entry is a 'delete' button. Below the list, there are two input fields: 'prefix' with the value 'ml' and a description 'A QName prefix.', and 'namespace uri' with the value 'http://marklogic.com/ml' and a description 'A namespace URI.' At the bottom of the dialog are 'ok' and 'cancel' buttons.

## 19.5 Deleting Namespaces for a Group

To delete namespaces that you defined in the Admin Interface, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click the group from which you want to delete the namespace, either in the tree menu or the Group Summary page.
3. Click the Namespaces icon on the left tree menu, under the specified group.
4. Locate the namespace to be deleted and click Delete.
5. A confirmation message displays. Confirm the delete and click OK.

The namespace is deleted from the group.

## 19.6 Deleting Namespaces for an HTTP or XDBC Server

To delete namespaces that you defined in the Admin Interface for an HTTP or XDBC server, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click on the group which contains the HTTP or XDBC server from which you want to delete the namespace, either in the tree menu or the Group Summary page.

3. Click on the App Servers icon.
4. Click on the name of the HTTP or XDBC server from which you want to delete the namespace, either in the tree menu or the App Server Summary page.
5. Click the Namespaces icon on the left tree menu, under the specified HTTP or XDBC server. The namespace configuration screen appears.
6. Locate the namespace to be deleted and click Delete.
7. A confirmation message displays. Confirm the delete and click OK.

The namespace is deleted from the App Server.

## 20.0 Understanding and Defining Schemas

This chapter describes schemas and lists procedures for defining them. The following topics are included:

- [Understanding Schemas](#)
- [Procedures For Defining Schemas](#)

For more information on the Schema database, loading schemas into MarkLogic Server, and using schemas in your applications, see the “Loading Schemas” chapter of the *Developer’s Guide*.

### 20.1 Understanding Schemas

A schema is a data dictionary for your XML content. To specify a schema, you need to define the namespace to which the schema applies as well as the location of the schema file.

Schemas define the types of elements within XML documents. When knowing the type of an XML element would be beneficial to evaluating an XQuery program, MarkLogic Server will look for the relevant schema document (based on that element's namespace) using the following strategy:

1. If the XQuery program explicitly references a schema for the namespace in question, MarkLogic Server uses this reference.
2. Otherwise, MarkLogic Server searches the schema database for an XML schema document whose target namespace is the same as the namespace of the element that MarkLogic Server is trying to type.
3. If no matching schema document is found in the database, MarkLogic Server looks in its `Config` directory for a matching schema document.
4. If no matching schema document is found in the `Config` directory, no schema is found.

Problems can arise in step 2 above when there are multiple schema documents in the schema database whose target namespace matches the namespace of the element that MarkLogic Server is trying to type. In this case, it is convenient to be able to use the Admin Interface to specify a default mapping.

Schema mappings can be specified for the HTTP or XDBC servers individually or for the group to apply to all HTTP or XDBC servers in the group. If the schema mapping defined for an HTTP or XDBC server conflicts with the schema mapping defined for the group, the former mapping is used.

When you specify a schema mapping in the Admin Interface, MarkLogic Server uses the following strategy to locate the schema:

1. First, MarkLogic Server searches the schema database for a document with the exact URI you specified in the schema mapping.

**Note:** If the schema mapping for the HTTP or XDBC server conflicts with the schema mapping for the group, the former mapping is used.

2. If no matching schema document is found in the schema database, MarkLogic Server looks in its `config` directory for a schema document whose filename matches the filename portion of the URI you specified.
3. If no matching schema document is found in the `config` directory, no schema is found.

If a namespace is invoked by one or more data elements stored in a particular database, and the schema for that namespace is defined for the group or HTTP server or XDBC server, MarkLogic Server applies the schema to the storage, indexing, and retrieval of that data.

**Note:** The schema database in this case is the schema database for the database in which the data is located.

## 20.2 Procedures For Defining Schemas

Use the following procedures for defining schemas:

- [Adding a Schema Definition for a Group](#)
- [Adding a Schema Definition for an HTTP or XDBC Server](#)
- [Viewing Schema Definitions for a Group](#)
- [Viewing Schema Definitions for an HTTP or XDBC Server](#)
- [Deleting a Schema Definition for a Group](#)
- [Deleting a Schema Definition for an HTTP or XDBC Server](#)

### 20.2.1 Adding a Schema Definition for a Group

To make a schema available to all HTTP or XDBC servers in a group, complete the following procedure:

1. Click the Groups icon on the left tree menu.
2. Click the group in which you want to define the schema.
3. Click the Schemas icon on the left tree menu, under the specified group.

- Click the Add tab. The Schema Configuration page displays:

- Enter a namespace URI and corresponding schema location.

If you are planning to store the schema in your `Config` directory, the following table lists the default location of the `Config` directory on each platform:

Platform	Schema Directory
Microsoft Windows	C:\Program Files\MarkLogic\Config
Red Hat Linux	/opt/MarkLogic/Config
Sun Solaris	/opt/MARKlogic/Config

- Repeat step [5](#) for other schemas as needed.
- Scroll to the top or bottom and click OK.

The schema is added to the group.

### 20.2.2 Adding a Schema Definition for an HTTP or XDBC Server

To make a schema available to a particular HTTP or XDBC server, perform the following steps:

- Click the Groups icon on the left tree menu.
- Click the name of the group which contains the HTTP or XDBC server to which you want to add a schema.

3. Click the App Servers icon.
4. Click the name of the HTTP server or XDBC server to which you want to add a schema.
5. Click the Schemas icon on the left tree menu, under the specified HTTP or XDBC server.
6. Click the Add tab. The Schema Configuration page displays:

7. Enter a namespace URI and corresponding schema location.

If you are planning to store the schema in your config directory, refer to the following table for the default location of the config directory on your platform:

Platform	Schema Directory
Microsoft Windows	C:\Program Files\MarkLogic\Config
Red Hat Linux	/opt/MarkLogic/Config
Sun Solaris	/opt/MARKlogic/Config

8. Repeat step [7](#) for other schemas as needed.
9. Scroll to the top or bottom and click OK.

The schema is added to the HTTP or XDBC server.

### 20.2.3 Viewing Schema Definitions for a Group

To view a schema definition for a group, complete the following procedure:

1. Click the Groups icon on the left tree menu.
2. Click the group that contains the schema you want to view.
3. Click the Schemas icon on the left tree menu, under the specified group.

The following example shows just one schema. It specifies that the schema for namespace `http://www.w3.org/1999/xhtml` is found in the file `xhtml1.1.xsd`, which is located in the config directory of your MarkLogic Server program directory.

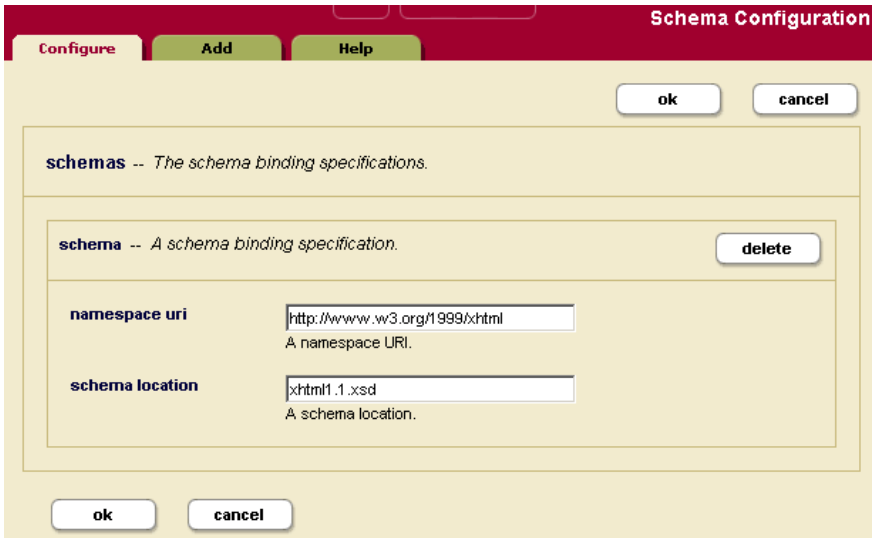
The screenshot shows a 'Schema Configuration' dialog box. At the top, there are three tabs: 'Configure', 'Add', and 'Help'. Below the tabs are 'ok' and 'cancel' buttons. The main content area is titled 'schemas -- The schema binding specifications.' and contains a list of one schema item: 'schema -- A schema binding specification.' with a 'delete' button. The schema details are: 'namespace uri' with the value 'http://www.w3.org/1999/xhtml' and a description 'A namespace URI.', and 'schema location' with the value 'xhtml1.1.xsd' and a description 'A schema location.' At the bottom are 'ok' and 'cancel' buttons.

#### 20.2.4 Viewing Schema Definitions for an HTTP or XDBC Server

To view a schema definition for an HTTP or XDBC Server, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click on the name of the group which contains the HTTP or XDBC server with the schema you want to view.
3. Click the App Servers icon.
4. Click the name of the HTTP server or XDBC server with the schema you want to view.
5. Click the Schemas icon on the left tree menu, under the specified HTTP or XDBC server.

The following example shows just one schema. It specifies that the schema for namespace `http://www.w3.org/1999/xhtml` is found in the file `xhtml1.1.xsd`, which is located in the config directory of your MarkLogic Server program directory.



### 20.2.5 Deleting a Schema Definition for a Group

To delete a schema definition for a group, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click the group from which you want to delete the schema.
3. Click the Schemas icon on the left tree menu, under the specified group.
4. Locate the schema definition to be deleted from the system and click Delete.
5. A confirmation message displays. Confirm the delete and click OK.

The schema is dropped from the group.

### 20.2.6 Deleting a Schema Definition for an HTTP or XDBC Server

To delete a schema definition for an HTTP or XDBC server, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click the name of the group which contains the HTTP or XDBC server with the schema you want to delete.
3. Click the App Servers icon.

4. Click the name of the HTTP server or XDBC server with the schema you want to delete.
5. Click the Schemas icon on the left tree menu, under the specified HTTP or XDBC server.
6. Click the Schemas icon on the left tree menu, under the specified HTTP or XDBC server.
7. Locate the schema definition to be deleted from the system and click Delete.
8. A confirmation message displays. Confirm the delete and click OK.

The schema is dropped from the HTTP or XDBC server.

## 21.0 Log Files

This chapter describes the log files and includes the following sections:

- [Understanding the Log Levels](#)
- [Configuring Log Files](#)
- [Viewing the System Log](#)
- [Viewing the File Log](#)

### 21.1 Understanding the Log Levels

MarkLogic Server sends log messages to both the operating system log and the MarkLogic Server file log. Depending on how you configure your logging functions, both logs may or may not receive the equivalent number of messages. To enhance performance, the system log should receive fewer messages than the MarkLogic Server file log.

MarkLogic Server uses the following log settings, where Finest is the most verbose while Emergency is the least verbose:

Log Level	Description
Finest	Extremely detailed debug level messages.
Finer	Very detailed debug level messages.
Fine	Detailed debug level messages.
Debug	Debug level messages.
Config	Configuration messages.
Info	Informational messages. This is the default setting.
Notice	Normal but significant conditions.
Warning	Warning conditions.
Error	Error conditions.
Critical	Critical conditions.
Alert	Immediate action required.
Emergency	System is unusable.

Log file settings are applied on a per-group basis.

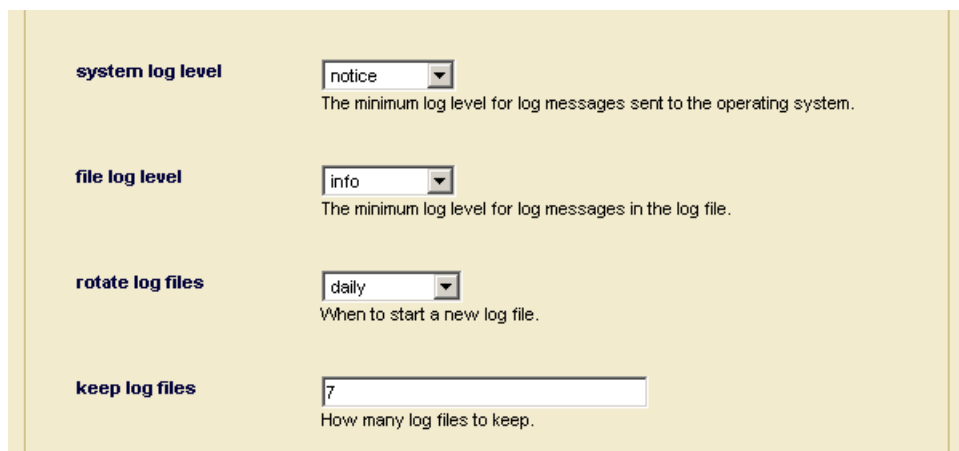
By default, the system log for a group is set to Notice while the file log is set to Info. As such, the system log receives fewer log messages than the file log. You may change these settings to suit your needs. For example, if you are debugging a system problem, you may want to set the level to Debug to get more information. Keep in mind that log levels Debug and above degrade system performance significantly, so these log levels should not normally be used.

## 21.2 Configuring Log Files

To configure how log information is generated, perform the following steps:

1. Click the Groups icon on the left tree menu.
2. Click the group for which you want to configure the log file settings.
3. Scroll down to the log settings, towards the bottom of the page.

The following example shows the default log settings:



The screenshot displays a configuration panel with a light beige background. It contains four settings, each with a label, a control element, and a descriptive text:

- system log level**: A dropdown menu set to "notice". Below it, the text reads: "The minimum log level for log messages sent to the operating system."
- file log level**: A dropdown menu set to "info". Below it, the text reads: "The minimum log level for log messages in the log file."
- rotate log files**: A dropdown menu set to "daily". Below it, the text reads: "When to start a new log file."
- keep log files**: A text input field containing the number "7". Below it, the text reads: "How many log files to keep."

4. Go to System Log Level and change the level if needed.
5. Go to File Log Level and change the logging level of the MarkLogic Server private log file (`ErrorLog.txt`) if needed.

6. Go to Rotate Log Files and select when MarkLogic Server should start a new private log file for this group.

The following table describes each time frame:

Time Frame	Description
Never	The log file grows without bound.
Daily	A new log file is started every day at 12:00 A.M.
Sunday	A new log file is started every week on Sunday at 12:00 A.M.
Saturday	A new log file is started every week on Saturday at 12:00 A.M.
Friday	A new log file is started every week on Friday at 12:00 A.M.
Thursday	A new log file is started every week on Thursday at 12:00 A.M.
Wednesday	A new log file is started every week on Wednesday at 12:00 A.M.
Tuesday	A new log file is started every week on Tuesday at 12:00 A.M.
Monday	A new log file is started every week on Monday at 12:00 A.M.
Monthly	A new log file is started at 12:00 AM on the first day of each month.

7. Go to Keep Log Files and enter the number of private log files to keep.

The private log files are kept in an aging archive. After the number of log files grows to the value specified in the Keep Log File setting, when a new log file is started, the oldest log file archive is automatically deleted.

8. Scroll to the top or bottom and click OK.

### 21.3 Viewing the System Log

The system log messages that MarkLogic Server generates are viewable using the standard system log viewing tools available for your platform. On Windows platforms, the seven levels of logging messages are collapsed into three broad categories.

## 21.4 Viewing the File Log

The private file log is maintained as a simple text file. You may view the current or any archived file log at any time using standard text file viewing tools.

The file is stored in the MarkLogic Server data directory for your platform. You may have overridden the default location for this directory at installation time. The following table lists the default location of the file logs on your platform:

Platform	Private Log File
Microsoft Windows	C:\Program Files\MarkLogic\Data\Logs\ErrorLog.txt
Red Hat Linux	/var/opt/MarkLogic/Logs/ErrorLog.txt
Sun Solaris	/var/opt/MARKlogic/Logs/ErrorLog.txt

This file contains a set of log messages ordered chronologically. The number of messages depends on the system activity and on the log level that you set. For example, a file log set to Debug would contain many lines of messages whereas a file log set to Emergency would contain the minimum set of messages.

Any trace events are also written to the MarkLogic Server private log file. Trace events are used to debug applications. You can enable and set trace events in the Admin Interface, on the Diagnostics page for a group. You can also generate your own trace events with the `xdmp:trace` function.

**Note:** There must be sufficient disk space on the filesystem in which the log files reside. If there is no space left on the log file device, MarkLogic Server will abort. Additionally, if there is no disk space available for the log files, MarkLogic Server will fail to start.

## 22.0 Appendix A: 'Hot' versus 'Cold' Admin Tasks

“Hot” admin tasks are defined as tasks that take effect immediately and do not require the server to restart. “Cold” admin tasks are defined as tasks that require one or more instances of the server to restart to reflect the changes.

In an Enterprise Edition clustered deployment, “cold” tasks will require one or more hosts in the cluster to restart their instance of MarkLogic Server in order to reflect the changes. In an Enterprise Edition single-server deployment and in all Standard Edition deployments, “cold” tasks will cause MarkLogic Server to restart in order to reflect the changes.

The tables below show the “hot” or “cold” status for adding objects, changing configuration parameters, and dropping objects for the following object types:

- [Groups](#)
- [HTTP, XDBC, and WebDAV Servers](#)
- [Databases](#)
- [Hosts](#)
- [Forests](#)
- [Mimetypes](#)
- [Security](#)

## 22.1 Groups

Add Object	Change Configuration Parameters	Delete Object
Hot	<p>The following group parameters are hot:</p> <ul style="list-style-type: none"> <li>&gt; group name</li> <li>&gt; system log level</li> <li>&gt; file log level</li> <li>&gt; rotate log files</li> <li>&gt; keep log files</li> <li>&gt; namespaces</li> <li>&gt; schemas</li> </ul> <p>The following group parameters are cold for the hosts in the group:</p> <ul style="list-style-type: none"> <li>&gt; list cache size</li> <li>&gt; compressed tree cache size</li> <li>&gt; expanded tree cache size</li> </ul> <p>Adding and dropping hosts from groups is cold for that host.</p>	Hot

## 22.2 HTTP, XDBC, and WebDAV Servers

Add Object	Change Configuration Parameters	Delete Object
Hot	<p>The following App Server parameters are hot:</p> <ul style="list-style-type: none"> <li>&gt; server name</li> <li>&gt; root</li> <li>&gt; database</li> <li>&gt; request timeout</li> <li>&gt; keep alive timeout</li> <li>&gt; session timeout</li> <li>&gt; time limit</li> <li>&gt; realm</li> <li>&gt; security mode</li> <li>&gt; namespaces</li> <li>&gt; schemas</li> </ul> <p>The following App Server parameters are cold for all hosts in the group defining the HTTP, XDBC, or WebDAV Server:</p> <ul style="list-style-type: none"> <li>&gt; port</li> <li>&gt; address</li> <li>&gt; backlog</li> <li>&gt; threads</li> </ul>	Cold

## 22.3 Databases

Add Object	Change Configuration Parameters	Delete Object
Hot	Parameters changes are hot	Hot

## 22.4 Hosts

Add Object	Change Configuration Parameters	Delete Object
Only the added host needs to restart	Only the host whose parameters change requires a restart. The rest of the hosts remain hot.	Hot for the remaining hosts;

## 22.5 Forests

Add Object	Change Configuration Parameters	Delete Object
Hot	Parameter changes are hot. Backup is hot. Restore, clear and drop are hot	Hot

## 22.6 Mimetypes

Add Object	Change Configuration Parameters	Delete Object
Cold	Parameter changes are cold.	Cold

## 22.7 Security

Add Object	Change Configuration Parameters	Delete Object
Hot	Parameter changes are hot.	Hot

## 23.0 Appendix B: Pre-defined Execute Privileges

The pre-defined execute privileges listed below are included with every installation of MarkLogic Server.

Name	Action URI	Description	Protects Function
admin-ut	<a href="http://marklogic.com/xdmp/privileges/admin-ui">http://marklogic.com/xdmp/privileges/admin-ui</a>	privilege to use the Admin Interface	admin built-ins
amp-add-roles	<a href="http://marklogic.com/xdmp/privileges/amp-add-roles">http://marklogic.com/xdmp/privileges/amp-add-roles</a>	privilege to assign additional roles to the amp	sec:amp-add-roles()
amp-get-roles	<a href="http://marklogic.com/xdmp/privileges/amp-get-roles">http://marklogic.com/xdmp/privileges/amp-get-roles</a>	privilege to get the roles associated with the amp	sec:amp-get-roles()
amp-remove-roles	<a href="http://marklogic.com/xdmp/privileges/amp-remove-roles">http://marklogic.com/xdmp/privileges/amp-remove-roles</a>	privilege to remove roles assigned to the amp	sec:amp-remove-roles()
amp-set-roles	<a href="http://marklogic.com/xdmp/privileges/amp-set-roles">http://marklogic.com/xdmp/privileges/amp-set-roles</a>	privilege to set the roles associated with the amp	sec:amp-set-roles()
any-collection	<a href="http://marklogic.com/xdmp/privileges/any-collection">http://marklogic.com/xdmp/privileges/any-collection</a>	privilege to add to or remove from any collection, regardless of whether it is protected	xdmp:document-add-collections(), xdmp:document-remove-collections(), xdmp:document-set-collections()
any-uri	<a href="http://marklogic.com/xdmp/privileges/any-uri">http://marklogic.com/xdmp/privileges/any-uri</a>	privilege to create a document with any uri, regardless of whether the uri is protected	xdmp:document-insert(), xdmp:document-load, xdmp:load()
cancel-any-requests	<a href="http://marklogic.com/xdmp/privileges/cancel-any-requests">http://marklogic.com/xdmp/privileges/cancel-any-requests</a>	privilege to cancel requests issued by any user attempting to cancel a request	admin built-ins
cancel-my-requests	<a href="http://marklogic.com/xdmp/privileges/cancel-my-requests">http://marklogic.com/xdmp/privileges/cancel-my-requests</a>	privilege to cancel requests issued by the user attempting to cancel a request	admin built-ins
collection-add-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-add-permissions">http://marklogic.com/xdmp/privileges/collection-add-permissions</a>	privilege to add permissions to a collection	sec:get-collections(), sec:collection-add-permissions
collection-get-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-get-permissions">http://marklogic.com/xdmp/privileges/collection-get-permissions</a>	privilege to get permissions on a collection	sec:collection-get-permissions()
collection-remove-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-remove-permissions">http://marklogic.com/xdmp/privileges/collection-remove-permissions</a>	privilege to remove permissions from a collection	sec:get-collections(), sec:collection-remove-permissions()

Name	Action URI	Description	Protects Function
collection-set-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-set-permissions">http://marklogic.com/xdmp/privileges/collection-set-permissions</a>	privilege to set permissions on a collection	sec:get-collections() sec:collection-set-permissions
count-builtins	<a href="http://marklogic.com/xdmp/privileges/counts">http://marklogic.com/xdmp/privileges/counts</a>	privilege to run xdmp:forest-counts	xdmp:forst-counts
create-amp	<a href="http://marklogic.com/xdmp/privileges/create-amp">http://marklogic.com/xdmp/privileges/create-amp</a>	privilege to create an amp	sec:create-amp()
create-domain	<a href="http://marklogic.com/xdmp/privileges/create-domain">http://marklogic.com/xdmp/privileges/create-domain</a>	privilege to create an domain	dom:create()
create-pipeline	<a href="http://marklogic.com/xdmp/privileges/create-pipeline">http://marklogic.com/xdmp/privileges/create-pipeline</a>	privilege to create a pipeline	p:insert(), p:create()
create-privilege	<a href="http://marklogic.com/xdmp/privileges/create-privilege">http://marklogic.com/xdmp/privileges/create-privilege</a>	privilege to create a privilege	sec:create-role()
create-role	<a href="http://marklogic.com/xdmp/privileges/create-role">http://marklogic.com/xdmp/privileges/create-role</a>	privilege to create a role	sec:create-role()
create-trigger	<a href="http://marklogic.com/xdmp/privileges/create-trigger">http://marklogic.com/xdmp/privileges/create-trigger</a>	privilege to create a trigger	trgr:create-trigger()
create-user	<a href="http://marklogic.com/xdmp/privileges/create-user">http://marklogic.com/xdmp/privileges/create-user</a>	privilege to create a user	sec:create-user()
get-amp	<a href="http://marklogic.com/xdmp/privileges/get-amp">http://marklogic.com/xdmp/privileges/get-amp</a>	privilege to get an amp	sec:get-amp()
get-privilege	<a href="http://marklogic.com/xdmp/privileges/get-privilege">http://marklogic.com/xdmp/privileges/get-privilege</a>	privilege to get a privilege from action uri and type	sec:get-privilege()
get-role-ids	<a href="http://marklogic.com/xdmp/privileges/get-role-ids">http://marklogic.com/xdmp/privileges/get-role-ids</a>	privilege to get role ids	internal functions
get-role-names	<a href="http://marklogic.com/xdmp/privileges/get-role-names">http://marklogic.com/xdmp/privileges/get-role-names</a>	privilege to get role names	internal functions

Name	Action URI	Description	Protects Function
grant-all-roles	<a href="http://marklogic.com/xdmp/privileges/grant-all-roles">http://marklogic.com/xdmp/privileges/grant-all-roles</a>	privilege to grant a user all roles. Either grant-all-roles or grant-my-roles would be needed by functions that assign roles.	sec:create-user(), sec:user-set-roles(), sec:user-add-roles(), sec:user-remove-roles(), sec:create-role(), sec:role-set-roles(), sec:role-add-roles(), sec:role-remove-roles(), sec:remove-role-from-roles(), sec:remove-role-from-privileges(), sec:remove-role-from-amps(), sec:create-role(), sec:privilege-set-roles(), sec:privilege-add-roles(), sec:privilege-remove-roles(), sec:create-amp(), sec:amp-set-roles(), sec:amp-add-roles(), sec:amp-remove-roles()

Name	Action URI	Description	Protects Function
grant-my-roles	<a href="http://marklogic.com/xdmp/privileges/grant-my-roles">http://marklogic.com/xdmp/privileges/grant-my-roles</a>	privilege to grant a user my roles. Either grant-all-roles or grant-my-roles would be needed by functions that assign roles.	sec:create-user(), sec:user-set-roles(), sec:user-add-roles(), sec:user-remove-roles(), sec:create-role(), sec:role-set-roles(), sec:role-add-roles(), sec:role-remove-roles(), sec:remove-role-from-roles(), sec:remove-role-from-privileges(), sec:remove-role-from-amps(), sec:create-role(), sec:privilege-set-roles(), sec:privilege-add-roles(), sec:privilege-remove-roles(), sec:create-amp(), sec:amp-set-roles(), sec:amp-add-roles(), sec:amp-remove-roles()
privilege-add-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-add-roles">http://marklogic.com/xdmp/privileges/privilege-add-roles</a>	privilege to assign the privilege to additional roles	sec:privilege-add-roles()
privilege-get-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-get-roles">http://marklogic.com/xdmp/privileges/privilege-get-roles</a>	privilege to get all roles associated with a privilege	sec:privilege-get-roles()
privilege-remove-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-remove-roles">http://marklogic.com/xdmp/privileges/privilege-remove-roles</a>	privilege to remove privilege from roles to which it is assigned	sec:privilege-remove-roles()
privilege-set-name	<a href="http://marklogic.com/xdmp/privileges/privilege-set-name">http://marklogic.com/xdmp/privileges/privilege-set-name</a>	privilege to set a privilege's name	sec:privilege-set-name()
privilege-set-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-set-roles">http://marklogic.com/xdmp/privileges/privilege-set-roles</a>	privilege to set roles associated with a privilege	sec:privilege-set-roles()
profile-any-requests	<a href="http://marklogic.com/xdmp/privileges/profile-any-requests">http://marklogic.com/xdmp/privileges/profile-any-requests</a>	privilege to profile requests initiated by any user	prof:enable and other profile APIs

Name	Action URI	Description	Protects Function
profile-my-requests	<a href="http://marklogic.com/xdmp/privileges/profile-my-requests">http://marklogic.com/xdmp/privileges/profile-my-requests</a>	privilege to profile requests initiated by the user running the request from which profiling is called	prof.enable and other profile APIs
protect-collection	<a href="http://marklogic.com/xdmp/privileges/protect-collection">http://marklogic.com/xdmp/privileges/protect-collection</a>	privilege to make a new or existing collection protected	sec:protect-collection()
remove-amp	<a href="http://marklogic.com/xdmp/privileges/remove-amp">http://marklogic.com/xdmp/privileges/remove-amp</a>	privilege to remove an amp from the security database	sec:remove-amp()
remove-privilege	<a href="http://marklogic.com/xdmp/privileges/remove-privilege">http://marklogic.com/xdmp/privileges/remove-privilege</a>	privilege to remove a privilege from the security database	sec:remove-privilege()
remove-role	<a href="http://marklogic.com/xdmp/privileges/remove-role">http://marklogic.com/xdmp/privileges/remove-role</a>	privilege to remove a role from the security database	sec:remove-role()
remove-role-from-amps	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-amps">http://marklogic.com/xdmp/privileges/remove-role-from-amps</a>	privilege to remove a role from all amps in the security database	sec:remove-role-from-amps()
remove-role-from-privileges	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-privileges">http://marklogic.com/xdmp/privileges/remove-role-from-privileges</a>	privilege to remove a role from all privileges in the security database	sec:remove-role-from-privileges()
remove-role-from-roles	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-roles">http://marklogic.com/xdmp/privileges/remove-role-from-roles</a>	privilege to remove a role from all roles in the security database	sec:remove-role-from-roles()
remove-role-from-users	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-users">http://marklogic.com/xdmp/privileges/remove-role-from-users</a>	privilege to remove a role from all users in the security database	sec:remove-role-from-users()
remove-user	<a href="http://marklogic.com/xdmp/privileges/remove-user">http://marklogic.com/xdmp/privileges/remove-user</a>	privilege to remove a user from the security database	sec:remove-user()
role-add-roles	<a href="http://marklogic.com/xdmp/privileges/role-add-roles">http://marklogic.com/xdmp/privileges/role-add-roles</a>	privilege to add roles to the roles of a specified role	sec:role-add-roles()
role-get-default-collections	<a href="http://marklogic.com/xdmp/privileges/role-get-default-collections">http://marklogic.com/xdmp/privileges/role-get-default-collections</a>	privilege to get a role's default collections	sec:role-get-default-collections()
role-get-default-permissions	<a href="http://marklogic.com/xdmp/privileges/role-get-default-permissions">http://marklogic.com/xdmp/privileges/role-get-default-permissions</a>	privilege to get a role's default permissions	sec:role-get-default-permissions()
role-get-description	<a href="http://marklogic.com/xdmp/privileges/role-get-description">http://marklogic.com/xdmp/privileges/role-get-description</a>	privilege to get a role's description	sec:role-get-description()
role-get-roles	<a href="http://marklogic.com/xdmp/privileges/role-get-roles">http://marklogic.com/xdmp/privileges/role-get-roles</a>	privilege to get all the roles included in the specified role	sec:role-get-roles()

Name	Action URI	Description	Protects Function
role-privileges	<a href="http://marklogic.com/xdmp/privileges/role-privileges">http://marklogic.com/xdmp/privileges/role-privileges</a>	privilege to get all the privileges for a given role	sec:role-privileges()
role-remove-roles	<a href="http://marklogic.com/xdmp/privileges/role-remove-roles">http://marklogic.com/xdmp/privileges/role-remove-roles</a>	privilege to remove roles from the roles of a specified role	sec:role-remove-roles()
role-set-default-collections	<a href="http://marklogic.com/xdmp/privileges/role-set-default-collections">http://marklogic.com/xdmp/privileges/role-set-default-collections</a>	privilege to set a role's default collections	sec:role-set-default-collections()
role-set-default-permissions	<a href="http://marklogic.com/xdmp/privileges/role-set-default-permissions">http://marklogic.com/xdmp/privileges/role-set-default-permissions</a>	privilege to set a role's default permissions	sec:role-set-default-permissions()
role-set-description	<a href="http://marklogic.com/xdmp/privileges/role-set-description">http://marklogic.com/xdmp/privileges/role-set-description</a>	privilege to set a role's name	sec:role-set-description()
role-set-name	<a href="http://marklogic.com/xdmp/privileges/role-set-name">http://marklogic.com/xdmp/privileges/role-set-name</a>	privilege to change a role's name	sec:role-set-name()
role-set-roles	<a href="http://marklogic.com/xdmp/privileges/role-set-roles">http://marklogic.com/xdmp/privileges/role-set-roles</a>	privilege to change all the roles in the specified role	sec:role-set-roles()
status-builtins	<a href="http://marklogic.com/xdmp/privileges/status-builtins">http://marklogic.com/xdmp/privileges/status-builtins</a>	privilege to access the status built-ins	status built-ins
unprotect-collection	<a href="http://marklogic.com/xdmp/privileges/unprotect-collection">http://marklogic.com/xdmp/privileges/unprotect-collection</a>	privilege to change roles for a collection	xdmp:document-add-collections(), xdmp:document-remove-collections(), xdmp:document-set-collections()
unprotected-collections	<a href="http://marklogic.com/xdmp/privileges/unprotected-collections">http://marklogic.com/xdmp/privileges/unprotected-collections</a>	privilege to add to or remove from collections that are unprotected	xdmp:document-add-collections(), xdmp:document-remove-collections(), xdmp:document-set-collections()
unprotected-uri	<a href="http://marklogic.com/xdmp/privileges/unprotected-uri">http://marklogic.com/xdmp/privileges/unprotected-uri</a>	privilege to create document with uri's that are unprotected	xdmp:document-insert(), xdmp:load()
user-add-roles	<a href="http://marklogic.com/xdmp/privileges/user-add-roles">http://marklogic.com/xdmp/privileges/user-add-roles</a>	privilege to add roles to a user	sec:user-add-roles()
user-get-default-collections	<a href="http://marklogic.com/xdmp/privileges/user-gt-default-collections">http://marklogic.com/xdmp/privileges/user-gt-default-collections</a>	privilege to get a user's default collections	sec:user-get-default-collections()
user-get-default-permissions	<a href="http://marklogic.com/xdmp/privileges/user-get-default-permissions">http://marklogic.com/xdmp/privileges/user-get-default-permissions</a>	privilege to get user's default permissions	sec:user-get-default-permissions()
user-get-description	<a href="http://marklogic.com/xdmp/privileges/user-get-description">http://marklogic.com/xdmp/privileges/user-get-description</a>	privilege to get user's description	sec:user-get-description (if not logged in as user)
user-get-roles	<a href="http://marklogic.com/xdmp/privileges/user-get-roles">http://marklogic.com/xdmp/privileges/user-get-roles</a>	privilege to get user's roles	sec:user-get-roles() (if not logged in as user)

Name	Action URI	Description	Protects Function
user-privileges	<a href="http://marklogic.com/xdmp/privileges/user-privileges">http://marklogic.com/xdmp/privileges/user-privileges</a>	privilege to get a user's complete privileges	sec:user-privileges() (if not logged in as user)
user-remove-roles	<a href="http://marklogic.com/xdmp/privileges/user-remove-roles">http://marklogic.com/xdmp/privileges/user-remove-roles</a>	privilege to remove roles from a user	sec:user-remove-roles()
user-set-default-collections	<a href="http://marklogic.com/xdmp/privileges/user-set-default-collections">http://marklogic.com/xdmp/privileges/user-set-default-collections</a>	privilege to set a user's default collections	sec:user-set-default-collections()
user-set-default-permissions	<a href="http://marklogic.com/xdmp/privileges/user-set-default-permissions">http://marklogic.com/xdmp/privileges/user-set-default-permissions</a>	privilege to set a user's default permissions	sec:user-set-default-permissions()
user-set-description	<a href="http://marklogic.com/xdmp/privileges/user-set-description">http://marklogic.com/xdmp/privileges/user-set-description</a>	privilege to set a user's description	sec:user-set-description (if not logged in as user)
user-set-name	<a href="http://marklogic.com/xdmp/privileges/user-set-name">http://marklogic.com/xdmp/privileges/user-set-name</a>	privilege to set a user's name	sec:user-set-name() (if not logged in as user)
user-set-password	<a href="http://marklogic.com/xdmp/privileges/user-set-password">http://marklogic.com/xdmp/privileges/user-set-password</a>	privilege to set user's password	sec:user-set-password() (if not logged in as user)
user-set-roles	<a href="http://marklogic.com/xdmp/privileges/user-set-roles">http://marklogic.com/xdmp/privileges/user-set-roles</a>	privilege to set a user's role	sec:user-set-roles()
xdbc:eval	<a href="http://marklogic.com/xdmp/privileges/xdbc-eval">http://marklogic.com/xdmp/privileges/xdbc-eval</a>	privilege to execute eval statements from xcc or xdbc	xdmp:eval()
xdbc:eval-in	<a href="http://marklogic.com/xdmp/privileges/xdbc-eval-in">http://marklogic.com/xdmp/privileges/xdbc-eval-in</a>	privilege to execute eval-in statements from xcc or xdbc	xdmp:eval-in()
xdbc:insert	<a href="http://marklogic.com/xdmp/privileges/xdbc-insert">http://marklogic.com/xdmp/privileges/xdbc-insert</a>	privilege to execute insert statements from xcc or xdbc	xcc or xdbc inserts
xdbc:insert-in	<a href="http://marklogic.com/xdmp/privileges/xdbc-insert-in">http://marklogic.com/xdmp/privileges/xdbc-insert-in</a>	privilege to execute insert statements from xcc or xdbc	xdbc or xcc inserts into another database
xdbc:invoke	<a href="http://marklogic.com/xdmp/privileges/xdbc-invoke">http://marklogic.com/xdmp/privileges/xdbc-invoke</a>	privilege to execute invoke statements from xcc or xdbc	xdbc or xcc invokes
xdbc:invoke-in	<a href="http://marklogic.com/xdmp/privileges/xdbc-invoke-in">http://marklogic.com/xdmp/privileges/xdbc-invoke-in</a>	privilege to execute invoke statements from xcc or xdbc	xdbc or xcc invokes into another database
xdbc:spawn	<a href="http://marklogic.com/xdmp/privileges/xdbc-spawn">http://marklogic.com/xdmp/privileges/xdbc-spawn</a>	privilege to execute spawn statements from xcc or xdbc	xdbc or xcc spawns

Name	Action URI	Description	Protects Function
xdbc:spawn-in	<a href="http://marklogic.com/xdmp/privileges/xdbc-spawn-in">http://marklogic.com/xdmp/privileges/xdbc-spawn-in</a>	privilege to execute spawn statements from xcc or xdbc	xdbc or xcc spawns into another database
xdmp:address-bindable	<a href="http://marklogic.com/xdmp/privileges/xdmp-address-bindable">http://marklogic.com/xdmp/privileges/xdmp-address-bindable</a>	privilege to perform admin functions.	admin built-ins
xdmp:amp-roles	<a href="http://marklogic.com/xdmp/privileges/xdmp-amp-roles">http://marklogic.com/xdmp/privileges/xdmp-amp-roles</a>	privilege to get an amp's roles	xdmp:amp-roles()
xdmp:castable-as	<a href="http://marklogic.com/xdmp/privileges/xdmp-castable-as">http://marklogic.com/xdmp/privileges/xdmp-castable-as</a>	privilege to perform admin functions	admin built-ins
xdmp:compressed-tree-cache-size	<a href="http://marklogic.com/xdmp/privileges/xdmp-compressed-tree-cache-size">http://marklogic.com/xdmp/privileges/xdmp-compressed-tree-cache-size</a>	privilege to perform admin functions	admin built-ins
xdmp:compressed-tree-cache-partitions	<a href="http://marklogic.com/xdmp/privileges/xdmp-compressed-tree-cache-partitions">http://marklogic.com/xdmp/privileges/xdmp-compressed-tree-cache-partitions</a>	privilege to perform admin functions	admin built-ins
xdmp:data-directory	<a href="http://marklogic.com/xdmp/privileges/xdmp-data-directory">http://marklogic.com/xdmp/privileges/xdmp-data-directory</a>	privilege to access the data directory	admin built-ins
xdmp:database-backup	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-backupt">http://marklogic.com/xdmp/privileges/xdmp-database-backupt</a>	privilege to perform a database backup	admin built-ins
xdmp:database-backup-cancel	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-backupt-cancel">http://marklogic.com/xdmp/privileges/xdmp-database-backupt-cancel</a>	privilege to cancel a database backup	admin built-ins
xdmp:database-backup-status	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-backupt-status">http://marklogic.com/xdmp/privileges/xdmp-database-backupt-status</a>	privilege to get status for a database backup	admin built-ins
xdmp:database-backup-validate	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-backupt-validate">http://marklogic.com/xdmp/privileges/xdmp-database-backupt-validate</a>	privilege to validate a database backup	admin built-ins
xdmp:database-restore	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-restore">http://marklogic.com/xdmp/privileges/xdmp-database-restore</a>	privilege to perform a database restore	admin built-ins
xdmp:database-restore-cancel	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-backupt">http://marklogic.com/xdmp/privileges/xdmp-database-backupt</a>	privilege to cancel a database restore	admin built-ins
xdmp:database-restore-status	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-restore-status">http://marklogic.com/xdmp/privileges/xdmp-database-restore-status</a>	privilege to get status for a database restore	admin built-ins
xdmp:database-restore-validate	<a href="http://marklogic.com/xdmp/privileges/xdmp-database-restore-validate">http://marklogic.com/xdmp/privileges/xdmp-database-restore-validate</a>	privilege to validate a database restore	admin built-ins
xdmp:default-in-memory-limit	<a href="http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-limit">http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-limit</a>	privilege to perform admin functions.	admin built-ins
xdmp:default-in-memory-list-size	<a href="http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-list-size">http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-list-size</a>	privilege to perform admin functions.	admin built-ins
xdmp:default-in-memory-range-in-dex-size	<a href="http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-range-index-size">http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-range-index-size</a>	privilege to perform admin functions	admin built-ins
xdmp:default-in-memory-tree-size	<a href="http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-tree-size">http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-tree-size</a>	privilege to perform admin functions	admin built-ins
xdmp:default-journal-count	<a href="http://marklogic.com/xdmp/privileges/xdmp-default-journal-count">http://marklogic.com/xdmp/privileges/xdmp-default-journal-count</a>	privilege to perform admin functions.	admin built-ins

Name	Action URI	Description	Protects Function
xdmp:default-journal-size	http://marklogic.com/xdmp/privileges/xdmp-default-journal-size	privilege to perform admin functions.	admin built-ins
xdmp:default-preallocate-journals	http://marklogic.com/xdmp/privileges/xdmp-default-preallocate-journals	privilege to perform admin functions.	admin built-ins
xdmp:delete-cluster-config-file	http://marklogic.com/xdmp/privileges/xdmp-delete-cluster-config-file	privilege to perform admin functions	admin built-ins
xdmp:delete-host-config-file	http://marklogic.com/xdmp/privileges/xdmp-delete-host-config-file	privilege to perform admin functions	admin built-ins
xdmp:disable-event	http://marklogic.com/xdmp/privileges/xdmp-disable-event	privilege to perform admin functions	admin built-ins
xdmp:document-get	http://marklogic.com/xdmp/privileges/xdmp-document-get	privilege to execute function	xdmp:document-get()
xdmp:document-load	http://marklogic.com/xdmp/privileges/xdmp-document-load	privilege to execute function	xdmp:document-load()
xdmp:email	http://marklogic.com/xdmp/privileges/xdmp-email	privilege to email	xdmp:email()
xdmp:email-address	http://marklogic.com/xdmp/privileges/xdmp-email-address	privilege to perform admin functions	admin built-ins
xdmp:enable-event	http://marklogic.com/xdmp/privileges/xdmp-enable-event	privilege to perform admin functions	admin built-ins
xdmp:eval	http://marklogic.com/xdmp/privileges/xdmp-eval	privilege to perform eval functions	xdmp:eval
xdmp:eval-in	http://marklogic.com/xdmp/privileges/xdmp-eval-in	privilege to perform eval-in functions	xdmp:eval-in
xdmp:expanded-tree-cache-size	http://marklogic.com/xdmp/privileges/xdmp-expanded-tree-cache-size	privilege to perform admin functions	admin built-ins
xdmp:expanded-tree-cache-partitions	http://marklogic.com/xdmp/privileges/xdmp-expanded-tree-cache-partitions	privilege to perform admin functions	admin built-ins
xdmp:filesystem-directory	http://marklogic.com/xdmp/privileges/xdmp-filesystem-directory	not supported	xdmp:filesystem-directory()
xdmp:forest-backup	http://marklogic.com/xdmp/privileges/xdmp-forest-backup	privilege to perform admin functions	admin built-ins
xdmp:forest-clear	http://marklogic.com/xdmp/privileges/xdmp-forest-clear	privilege to perform admin functions	admin built-ins
xdmp:forest-delete	http://marklogic.com/xdmp/privileges/xdmp-forest-delete	privilege to perform admin functions	admin built-ins
xdmp:forest-restore	http://marklogic.com/xdmp/privileges/xdmp-forest-restore	privilege to perform admin functions	admin built-ins
xdmp:forest-status	http://marklogic.com/xdmp/privileges/xdmp-forest-status	privilege to perform admin functions	admin built-ins
xdmp:get	http://marklogic.com/xdmp/privileges/xdmp-get	privilege to get a document into memory	xdmp:get()

Name	Action URI	Description	Protects Function
xdmp:get-forest-keys	http://marklogic.com/xdmp/privileges/xdmp-get-forest-keys	privilege to perform admin functions	admin built-ins
xdmp:get-hot-updates	http://marklogic.com/xdmp/privileges/xdmp-get-hot-updates	privilege to perform admin functions	admin built-ins
xdmp:host-cpus	http://marklogic.com/xdmp/privileges/xdmp-host-cpus	privilege to perform admin functions	admin built-ins
xdmp:host-size	http://marklogic.com/xdmp/privileges/xdmp-host-size	privilege to perform admin functions	admin built-ins
xdmp:hostname	http://marklogic.com/xdmp/privileges/xdmp-hostname	privilege to perform admin functions	admin built-ins
xdmp:install-directory	http://marklogic.com/xdmp/privileges/xdmp-install-directory	privilege to access the installation directory	admin built-ins
xdmp:invoke	http://marklogic.com/xdmp/privileges/xdmp-invoke	privilege to perform invoke functions	xdmp:invoke
xdmp:invoke-in	http://marklogic.com/xdmp/privileges/xdmp-invoke-in	privilege to perform invoke-in functions	xdmp:invoke-in
xdmp:license-accepted	http://marklogic.com/xdmp/privileges/xdmp-license-accepted	privilege to perform admin functions	admin built-ins
xdmp:license-fee	http://marklogic.com/xdmp/privileges/xdmp-license-fee	privilege to perform admin functions	admin built-ins
xdmp:license-key	http://marklogic.com/xdmp/privileges/xdmp-license-key	privilege to perform admin functions	admin built-ins
xdmp:license-key-agreement	http://marklogic.com/xdmp/privileges/xdmp-license-key-agreement	privilege to perform admin functions	admin built-ins
xdmp:license-key-cpus	http://marklogic.com/xdmp/privileges/xdmp-license-key-cpus	privilege to perform admin functions	admin built-ins
xdmp:license-key-decode	http://marklogic.com/xdmp/privileges/xdmp-license-key-decode	privilege to perform admin functions	admin built-ins
xdmp:license-key-encode	http://marklogic.com/xdmp/privileges/xdmp-license-key-encode	privilege to perform admin functions	admin built-ins
xdmp:license-key-expires	http://marklogic.com/xdmp/privileges/xdmp-key-expires	privilege to perform admin functions	admin built-ins
xdmp:license-key-options	http://marklogic.com/xdmp/privileges/xdmp-license-key-options	privilege to perform admin functions	admin built-ins
xdmp:license-key-size	http://marklogic.com/xdmp/privileges/xdmp-license-key-size	privilege to perform admin functions	admin built-ins
xdmp:license-key-valid	http://marklogic.com/xdmp/privileges/xdmp-license-key-valid	privilege to perform admin functions	admin built-ins
xdmp:licensee	http://marklogic.com/xdmp/privileges/xdmp-licensee	privilege to perform admin functions	admin built-ins
xdmp:list-cache-size	http://marklogic.com/xdmp/privileges/xdmp-list-cache-size	privilege to perform admin functions	admin built-ins

Name	Action URI	Description	Protects Function
xdmp:list-tree-cache-partitions	http://marklogic.com/xdmp/privileges/xdmp-list-tree-cache-partitions	privilege to perform admin functions	admin built-ins
xdmp:load	http://marklogic.com/xdmp/privileges/xdmp-load	privilege needed to load a document from the file system	xdmp:load()
xdmp:login	http://marklogic.com/xdmp/privileges/xdmp-login	privilege to log in a user without the corresponding password	xdmp-login()
xdmp:merge	http://marklogic.com/xdmp/privileges/xdmp-merge	privilege to start merging the forests	xdmp-merge()
xdmp:merging	http://marklogic.com/xdmp/privileges/xdmp-merging	privilege to get forest ids of forests currently merging	xdmp:merging()
xdmp:pre-release-expires	http://marklogic.com/xdmp/privileges/xdmp-pre-release-expires	privilege to perform admin functions	admin built-ins
xdmp:privilege-roles	http://marklogic.com/xdmp/privileges/xdmp-privilege-roles	privilege needed to get a role's privileges	xdmp:privilege-roles()
xdmp:read-cluster-config-file	http://marklogic.com/xdmp/privileges/xdmp-read-cluster-config-file	privilege to perform admin functions	admin built-ins
xdmp:read-host-config-file	http://marklogic.com/xdmp/privileges/xdmp-read-host-config-file	privilege to perform admin functions	admin built-ins
xdmp:restart	http://marklogic.com/xdmp/privileges/xdmp-restart	privilege to perform admin functions	admin built-ins
xdmp:role-roles	http://marklogic.com/xdmp/privileges/xdmp-role-roles	privilege to get a role's roles	xdmp:role-roles()
xdmp:save	http://marklogic.com/xdmp/privileges/xdmp-save	privilege needed to save a document to the file system	xdmp:save()
xdmp:server-backup	http://marklogic.com/xdmp/privileges/xdmp-server-backup	privilege to perform admin functions	admin built-ins
xdmp:server-import-qualities	http://marklogic.com/xdmp/privileges/xdmp-server-import-qualities	privilege to perform admin functions	admin built-ins
xdmp:server-restore	http://marklogic.com/xdmp/privileges/xdmp-server-restore	privilege to perform admin functions	admin built-ins
xdmp:set-hot-updates	http://marklogic.com/xdmp/privileges/xdmp-set-hot-updates	privilege to perform admin functions	admin built-ins
xdmp:shutdown	http://marklogic.com/xdmp/privileges/xdmp-shutdown	privilege to perform admin functions	admin built-ins
xdmp:smtp-relay	http://marklogic.com/xdmp/privileges/xdmp-smtp-relay	privilege to perform admin functions	admin built-ins
xdmp:timestamp	http://marklogic.com/xdmp/privileges/xdmp-timestamp	privilege to perform point-in-time queries	xdmp:eval, xdmp:invoke (timestamp option)

Name	Action URI	Description	Protects Function
xdmp:user-roles	<a href="http://marklogic.com/xdmp/privileges/xdmp-user-roles">http://marklogic.com/xdmp/privileges/xdmp-user-roles</a>	privilege to get a user's roles	xdmp:user-roles()
xdmp:username	<a href="http://marklogic.com/xdmp/privileges/xdmp-username">http://marklogic.com/xdmp/privileges/xdmp-username</a>	privilege to perform admin functions	admin built-ins
xdmp:write-cluster-config-file	<a href="http://marklogic.com/xdmp/privileges/xdmp-write-cluster-config-file">http://marklogic.com/xdmp/privileges/xdmp-write-cluster-config-file</a>	privilege to perform admin functions	admin built-ins
xdmp:write-host-config-file	<a href="http://marklogic.com/xdmp/privileges/xdmp-write-host-config-file">http://marklogic.com/xdmp/privileges/xdmp-write-host-config-file</a>	privilege to perform admin functions	admin built-ins

## 24.0 Appendix C: Pre-defined Roles

The following roles are pre-defined in every installation of MarkLogic Server. To give a user execute privileges listed for each pre-defined role, you may add the execute privileges individually to an existing role for the user, or add the pre-defined role to the user's set of roles.

The following are the pre-built roles in MarkLogic Server:

- [admin](#)
- [admin-builtins](#)
- [domain-management](#)
- [filesystem-access](#)
- [merge](#)
- [pipeline-management](#)
- [security](#)
- [trigger-management](#)

### 24.1 admin

The admin role is given all privileges and permissions to perform any action in the system. There are no default permissions associated with the admin role.

### 24.2 admin-builtins

The admin-builtins role has the execute privileges to call the admin built-in functions. The execute privileges given to the admin-builtins role are:

Name	Action URI
cancel-any-request	http://marklogic.com/xdmp/privileges/cancel-any-request
cancel-my-request	http://marklogic.com/xdmp/privileges/cancel-my-request
count-builtins	http://marklogic.com/xdmp/privileges/counts
xdmp:address-bindable	http://marklogic.com/xdmp/privileges/xdmp-address-bindable
xdmp:amp-roles	http://marklogic.com/xdmp/privileges/xdmp-amp-roles
xdmp:castable-as	http://marklogic.com/xdmp/privileges/xdmp-castable-as
xdmp:compressed-tree-cache-size	http://marklogic.com/xdmp/privileges/xdmp-compressed-tree-cache-size
xdmp:compressed-tree-cache-partitions	http://marklogic.com/xdmp/privileges/xdmp-compressed-tree-cache-partitions
xdmp:default-in-memory-limit	http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-limit
xdmp:default-in-memory-list-size	http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-list-size
xdmp:default-in-memory-range-index-size	http://marklogic.com/xdmp/privileges/xdmp-default-in-memory-range-index-size
xdmp:in-memory-tree-size	http://marklogic.com/xdmp/privileges/xdmp-in-memory-tree-size

Name	Action URI
xmdp:delete-cluster-config-file	<a href="http://marklogic.com/xmdp/privileges/xmdp-delete-cluster-config-file">http://marklogic.com/xmdp/privileges/xmdp-delete-cluster-config-file</a>
xmdp:delete-host-config-file	<a href="http://marklogic.com/xmdp/privileges/xmdp-delete-host-config-file">http://marklogic.com/xmdp/privileges/xmdp-delete-host-config-file</a>
xmdp:directory	<a href="http://marklogic.com/xmdp/privileges/xmdp-directory">http://marklogic.com/xmdp/privileges/xmdp-directory</a>
xmdp:disable-event	<a href="http://marklogic.com/xmdp/privileges/xmdp-disable-event">http://marklogic.com/xmdp/privileges/xmdp-disable-event</a>
xmdp:email	<a href="http://marklogic.com/xmdp/privileges/xmdp-email">http://marklogic.com/xmdp/privileges/xmdp-email</a>
xmdp:email-address	<a href="http://marklogic.com/xmdp/privileges/xmdp-email-address">http://marklogic.com/xmdp/privileges/xmdp-email-address</a>
xmdp:enable-event	<a href="http://marklogic.com/xmdp/privileges/xmdp-enable-event">http://marklogic.com/xmdp/privileges/xmdp-enable-event</a>
xmdp:expanded-tree-cache-size	<a href="http://marklogic.com/xmdp/privileges/xmdp-expanded-tree-cache-size">http://marklogic.com/xmdp/privileges/xmdp-expanded-tree-cache-size</a>
xmdp:expanded-tree-cache-partitions	<a href="http://marklogic.com/xmdp/privileges/xmdp-expanded-tree-cache-partitions">http://marklogic.com/xmdp/privileges/xmdp-expanded-tree-cache-partitions</a>
xmdp:forest-backup	<a href="http://marklogic.com/xmdp/privileges/xmdp-forest-backup">http://marklogic.com/xmdp/privileges/xmdp-forest-backup</a>
xmdp:forest-clear	<a href="http://marklogic.com/xmdp/privileges/xmdp-forest-clear">http://marklogic.com/xmdp/privileges/xmdp-forest-clear</a>
xmdp:forest-delete	<a href="http://marklogic.com/xmdp/privileges/xmdp-forest-delete">http://marklogic.com/xmdp/privileges/xmdp-forest-delete</a>
xmdp:forest-restore	<a href="http://marklogic.com/xmdp/privileges/xmdp-forest-restore">http://marklogic.com/xmdp/privileges/xmdp-forest-restore</a>
xmdp:forest-status	<a href="http://marklogic.com/xmdp/privileges/xmdp-forest-status">http://marklogic.com/xmdp/privileges/xmdp-forest-status</a>
xmdp:forest-keys	<a href="http://marklogic.com/xmdp/privileges/xmdp-forest-keys">http://marklogic.com/xmdp/privileges/xmdp-forest-keys</a>
xmdp:get-hot-updates	<a href="http://marklogic.com/xmdp/privileges/xmdp-get-hot-updates">http://marklogic.com/xmdp/privileges/xmdp-get-hot-updates</a>
xmdp:hostname	<a href="http://marklogic.com/xmdp/privileges/xmdp-hostname">http://marklogic.com/xmdp/privileges/xmdp-hostname</a>
xmdp:license-accepted	<a href="http://marklogic.com/xmdp/privileges/xmdp-license-accepted">http://marklogic.com/xmdp/privileges/xmdp-license-accepted</a>
xmdp:list-cache-size	<a href="http://marklogic.com/xmdp/privileges/xmdp-list-cache-size">http://marklogic.com/xmdp/privileges/xmdp-list-cache-size</a>
xmdp:list-cache-partitions	<a href="http://marklogic.com/xmdp/privileges/xmdp-list-cache-partitions">http://marklogic.com/xmdp/privileges/xmdp-list-cache-partitions</a>
xmdp:pre-release-expires	<a href="http://marklogic.com/xmdp/privileges/xmdp-pre-release-expires">http://marklogic.com/xmdp/privileges/xmdp-pre-release-expires</a>
xmdp:read-cluster-config-file	<a href="http://marklogic.com/xmdp/privileges/xmdp-read-cluster-config-file">http://marklogic.com/xmdp/privileges/xmdp-read-cluster-config-file</a>
xmdp:read-host-config-file	<a href="http://marklogic.com/xmdp/privileges/xmdp-read-host-config-file">http://marklogic.com/xmdp/privileges/xmdp-read-host-config-file</a>
xmdp:restart	<a href="http://marklogic.com/xmdp/privileges/xmdp-restart">http://marklogic.com/xmdp/privileges/xmdp-restart</a>
xmdp:server-backup	<a href="http://marklogic.com/xmdp/privileges/xmdp-server-backup">http://marklogic.com/xmdp/privileges/xmdp-server-backup</a>
xmdp:server-import-qualities	<a href="http://marklogic.com/xmdp/privileges/xmdp-server-import-qualities">http://marklogic.com/xmdp/privileges/xmdp-server-import-qualities</a>
xmdp:server-restore	<a href="http://marklogic.com/xmdp/privileges/xmdp-server-restore">http://marklogic.com/xmdp/privileges/xmdp-server-restore</a>
xmdp:set-hot-updates	<a href="http://marklogic.com/xmdp/privileges/xmdp-set-hot-updates">http://marklogic.com/xmdp/privileges/xmdp-set-hot-updates</a>
xmdp:shutdown	<a href="http://marklogic.com/xmdp/privileges/xmdp-shutdown">http://marklogic.com/xmdp/privileges/xmdp-shutdown</a>
xmdp:smtp-relay	<a href="http://marklogic.com/xmdp/privileges/xmdp-smtp-relay">http://marklogic.com/xmdp/privileges/xmdp-smtp-relay</a>
xmdp:username	<a href="http://marklogic.com/xmdp/privileges/xmdp-username">http://marklogic.com/xmdp/privileges/xmdp-username</a>
xmdp:write-cluster-config-file	<a href="http://marklogic.com/xmdp/privileges/xmdp-write-cluster-config-file">http://marklogic.com/xmdp/privileges/xmdp-write-cluster-config-file</a>
xmdp:write-host-config-file	<a href="http://marklogic.com/xmdp/privileges/xmdp-write-host-config-file">http://marklogic.com/xmdp/privileges/xmdp-write-host-config-file</a>

There are no default permissions associated with the admin-builtins role.

### 24.3 domain-management

The domain-management role has the privileges to create and modify content processing domains. The domain-management role has no execute privileges associated with it, but it has the following default permissions:

Role	Capability
domain-management	Read
domain-management	Update

### 24.4 filesystem-access

The filesystem-access role has the privileges to access the file system. The execute privileges given to the filesystem-access role are:

Name	Action URI
xdmp:document-get	<a href="http://marklogic.com/xdmp/privileges/xdmp-document-get">http://marklogic.com/xdmp/privileges/xdmp-document-get</a>
xdmp:document-load	<a href="http://marklogic.com/xdmp/privileges/xdmp-document-load">http://marklogic.com/xdmp/privileges/xdmp-document-load</a>
xdmp:get	<a href="http://marklogic.com/xdmp/privileges/xdmp-get">http://marklogic.com/xdmp/privileges/xdmp-get</a>
xdmp:load	<a href="http://marklogic.com/xdmp/privileges/xdmp-load">http://marklogic.com/xdmp/privileges/xdmp-load</a>
xdmp:save	<a href="http://marklogic.com/xdmp/privileges/xdmp-save">http://marklogic.com/xdmp/privileges/xdmp-save</a>

There are no default permissions associated with the filesystem-access role.

### 24.5 merge

The merge role has the privileges related to forest merging. The execute privileges given to the merge role are:

Name	Action URI
xdmp:merge	<a href="http://marklogic.com/xdmp/privileges/xdmp-merge">http://marklogic.com/xdmp/privileges/xdmp-merge</a>
xdmp:merging	<a href="http://marklogic.com/xdmp/privileges/xdmp-merging">http://marklogic.com/xdmp/privileges/xdmp-merging</a>

There are no default permissions associated with the admin-builtins role.

## 24.6 pipeline-management

The pipeline-management role has the privileges to create and modify content processing pipelines. The pipeline-management role has no execute privileges associated with it, but it has the following default permissions:

Role	Capability
pipeline-management	Read
pipeline-management	Update

## 24.7 security

The security role has the privileges needed to perform security functions. The execute privileges given to the security role are:

Name	Action URI
amp-add-roles	<a href="http://marklogic.com/xdmp/privileges/amp-add-roles">http://marklogic.com/xdmp/privileges/amp-add-roles</a>
amp-get-roles	<a href="http://marklogic.com/xdmp/privileges/amp-get-roles">http://marklogic.com/xdmp/privileges/amp-get-roles</a>
amp-remove-roles	<a href="http://marklogic.com/xdmp/privileges/amp-remove-roles">http://marklogic.com/xdmp/privileges/amp-remove-roles</a>
amp-set-roles	<a href="http://marklogic.com/xdmp/privileges/amp-set-roles">http://marklogic.com/xdmp/privileges/amp-set-roles</a>
any-collection	<a href="http://marklogic.com/xdmp/privileges/any-collection">http://marklogic.com/xdmp/privileges/any-collection</a>
any-uri	<a href="http://marklogic.com/xdmp/privileges/any-uri">http://marklogic.com/xdmp/privileges/any-uri</a>
collection-add-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-add-permissions">http://marklogic.com/xdmp/privileges/collection-add-permissions</a>
collection-get-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-get-permissions">http://marklogic.com/xdmp/privileges/collection-get-permissions</a>
collection-remove-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-remove-permissions">http://marklogic.com/xdmp/privileges/collection-remove-permissions</a>
collection-set-permissions	<a href="http://marklogic.com/xdmp/privileges/collection-set-permissions">http://marklogic.com/xdmp/privileges/collection-set-permissions</a>
create-amp	<a href="http://marklogic.com/xdmp/privileges/create-amp">http://marklogic.com/xdmp/privileges/create-amp</a>
create-privilege	<a href="http://marklogic.com/xdmp/privileges/create-privilege">http://marklogic.com/xdmp/privileges/create-privilege</a>
create-role	<a href="http://marklogic.com/xdmp/privileges/create-role">http://marklogic.com/xdmp/privileges/create-role</a>
create-user	<a href="http://marklogic.com/xdmp/privileges/create-user">http://marklogic.com/xdmp/privileges/create-user</a>
get-amp	<a href="http://marklogic.com/xdmp/privileges/get-amp">http://marklogic.com/xdmp/privileges/get-amp</a>
get-privilege	<a href="http://marklogic.com/xdmp/privileges/get-privilege">http://marklogic.com/xdmp/privileges/get-privilege</a>
get-role-ids	<a href="http://marklogic.com/xdmp/privileges/get-role-ids">http://marklogic.com/xdmp/privileges/get-role-ids</a>
grant-all-roles	<a href="http://marklogic.com/xdmp/privileges/grant-all-roles">http://marklogic.com/xdmp/privileges/grant-all-roles</a>
grant-my-roles	<a href="http://marklogic.com/xdmp/privileges/grant-my-roles">http://marklogic.com/xdmp/privileges/grant-my-roles</a>
permission	<a href="http://marklogic.com/xdmp/privileges/permission">http://marklogic.com/xdmp/privileges/permission</a>
privilege-add-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-add-roles">http://marklogic.com/xdmp/privileges/privilege-add-roles</a>
privilege-get-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-get-roles">http://marklogic.com/xdmp/privileges/privilege-get-roles</a>
privilege-remove-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-remove-roles">http://marklogic.com/xdmp/privileges/privilege-remove-roles</a>
privilege-set-name	<a href="http://marklogic.com/xdmp/privileges/privilege-set-name">http://marklogic.com/xdmp/privileges/privilege-set-name</a>

Name	Action URI
privilege-set-roles	<a href="http://marklogic.com/xdmp/privileges/privilege-set-roles">http://marklogic.com/xdmp/privileges/privilege-set-roles</a>
protect-collection	<a href="http://marklogic.com/xdmp/privileges/protect-collection">http://marklogic.com/xdmp/privileges/protect-collection</a>
remove-amp	<a href="http://marklogic.com/xdmp/privileges/remove-amp">http://marklogic.com/xdmp/privileges/remove-amp</a>
remove-privilege	<a href="http://marklogic.com/xdmp/privileges/remove-privilege">http://marklogic.com/xdmp/privileges/remove-privilege</a>
remove-role	<a href="http://marklogic.com/xdmp/privileges/remove-role">http://marklogic.com/xdmp/privileges/remove-role</a>
remove-role-from-amps	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-amps">http://marklogic.com/xdmp/privileges/remove-role-from-amps</a>
remove-role-from-privileges	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-privileges">http://marklogic.com/xdmp/privileges/remove-role-from-privileges</a>
remove-role-from-roles	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-roles">http://marklogic.com/xdmp/privileges/remove-role-from-roles</a>
remove-role-from-users	<a href="http://marklogic.com/xdmp/privileges/remove-role-from-users">http://marklogic.com/xdmp/privileges/remove-role-from-users</a>
remove-user	<a href="http://marklogic.com/xdmp/privileges/remove-user">http://marklogic.com/xdmp/privileges/remove-user</a>
role-add-roles	<a href="http://marklogic.com/xdmp/privileges/role-add-roles">http://marklogic.com/xdmp/privileges/role-add-roles</a>
role-get-default-collections	<a href="http://marklogic.com/xdmp/privileges/role-get-default-collections">http://marklogic.com/xdmp/privileges/role-get-default-collections</a>
role-get-default-permissions	<a href="http://marklogic.com/xdmp/privileges/role-get-default-permissions">http://marklogic.com/xdmp/privileges/role-get-default-permissions</a>
role-get-roles	<a href="http://marklogic.com/xdmp/privileges/role-get-roles">http://marklogic.com/xdmp/privileges/role-get-roles</a>
role-privileges	<a href="http://marklogic.com/xdmp/privileges/role-privileges">http://marklogic.com/xdmp/privileges/role-privileges</a>
role-remove-roles	<a href="http://marklogic.com/xdmp/privileges/role-remove-roles">http://marklogic.com/xdmp/privileges/role-remove-roles</a>
role-set-default-collections	<a href="http://marklogic.com/xdmp/privileges/role-set-default-collections">http://marklogic.com/xdmp/privileges/role-set-default-collections</a>
role-set-default-permissions	<a href="http://marklogic.com/xdmp/privileges/role-set-default-permissions">http://marklogic.com/xdmp/privileges/role-set-default-permissions</a>
role-set-description	<a href="http://marklogic.com/xdmp/privileges/role-set-description">http://marklogic.com/xdmp/privileges/role-set-description</a>
role-set-name	<a href="http://marklogic.com/xdmp/privileges/role-set-name">http://marklogic.com/xdmp/privileges/role-set-name</a>
role-set-roles	<a href="http://marklogic.com/xdmp/privileges/role-set-roles">http://marklogic.com/xdmp/privileges/role-set-roles</a>
unprotect-collection	<a href="http://marklogic.com/xdmp/privileges/unprotect-collection">http://marklogic.com/xdmp/privileges/unprotect-collection</a>
user-add-roles	<a href="http://marklogic.com/xdmp/privileges/user-add-roles">http://marklogic.com/xdmp/privileges/user-add-roles</a>
user-get-default-collections	<a href="http://marklogic.com/xdmp/privileges/user-gt-default-collections">http://marklogic.com/xdmp/privileges/user-gt-default-collections</a>
user-get-default-permissions	<a href="http://marklogic.com/xdmp/privileges/user-get-default-permissions">http://marklogic.com/xdmp/privileges/user-get-default-permissions</a>
user-get-description	<a href="http://marklogic.com/xdmp/privileges/user-get-description">http://marklogic.com/xdmp/privileges/user-get-description</a>
user-get-roles	<a href="http://marklogic.com/xdmp/privileges/user-get-roles">http://marklogic.com/xdmp/privileges/user-get-roles</a>
user-privileges	<a href="http://marklogic.com/xdmp/privileges/user-privileges">http://marklogic.com/xdmp/privileges/user-privileges</a>
user-remove-roles	<a href="http://marklogic.com/xdmp/privileges/user-remove-roles">http://marklogic.com/xdmp/privileges/user-remove-roles</a>
user-set-default-collections	<a href="http://marklogic.com/xdmp/privileges/user-set-default-collections">http://marklogic.com/xdmp/privileges/user-set-default-collections</a>
user-set-default-permissions	<a href="http://marklogic.com/xdmp/privileges/user-set-default-permissions">http://marklogic.com/xdmp/privileges/user-set-default-permissions</a>
user-set-description	<a href="http://marklogic.com/xdmp/privileges/user-set-description">http://marklogic.com/xdmp/privileges/user-set-description</a>
user-set-name	<a href="http://marklogic.com/xdmp/privileges/user-set-name">http://marklogic.com/xdmp/privileges/user-set-name</a>
user-set-password	<a href="http://marklogic.com/xdmp/privileges/user-set-password">http://marklogic.com/xdmp/privileges/user-set-password</a>
user-set-roles	<a href="http://marklogic.com/xdmp/privileges/user-set-roles">http://marklogic.com/xdmp/privileges/user-set-roles</a>
xdmp:amp-roles	<a href="http://marklogic.com/xdmp/privileges/xdmp:amp-roles">http://marklogic.com/xdmp/privileges/xdmp:amp-roles</a>
xdmp:privilege-roles	<a href="http://marklogic.com/xdmp/privileges/xdmp:privilege-roles">http://marklogic.com/xdmp/privileges/xdmp:privilege-roles</a>
xdmp:role-roles	<a href="http://marklogic.com/xdmp/privileges/xdmp:role-roles">http://marklogic.com/xdmp/privileges/xdmp:role-roles</a>
xdmp:user-roles	<a href="http://marklogic.com/xdmp/privileges/xdmp:user-roles">http://marklogic.com/xdmp/privileges/xdmp:user-roles</a>

Default permissions for the security role are:

Role	Capability
security	Read
security	Insert
security	Update

## 24.8 trigger-management

The trigger-management role has the privileges to create and modify triggers. The trigger-management role has no execute privileges associated with it, but it has the following default permissions:

Role	Capability
trigger-management	Read
trigger-management	Update

## Technical Support

Mark Logic provides technical support according to the terms detailed in your Software License Agreement. For evaluation licenses, Mark Logic may provide support on an “as possible” basis.

For registered customers, we invite you to visit our support website at <http://support.marklogic.com> to access our full suite of documentation and help materials. For all customers, including community licensed customers, visit the Mark Logic Developer’s site at <http://developer.marklogic.com>, which includes full product documentation, downloads, and developer community open-source projects.

If you have questions or comments, you may contact Mark Logic Technical Support at the following email address:

[support@marklogic.com](mailto:support@marklogic.com)

If reporting a query evaluation problem, please be sure to include the sample XQuery code.