
MarkLogic Server

XCC Developer's Guide

Release 3.2
May, 2007

Copyright

© Copyright 2002-2007 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 Inxight Software, Inc. Copyright © 1996-2006. All rights reserved. www.inxight.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

XCC Developer's Guide

Copyright	2
1.0 Introduction to XCC	5
1.1 Overview of XCC	5
1.1.1 XCC Client Libraries Communicate With an XDBC Server	5
1.1.2 Client-Server Architecture	6
1.1.3 Automatically Pools Connections	6
1.2 API and Other Documentation	7
1.3 XCC Requirements	7
1.3.1 XCC MarkLogic Server Requirements	7
1.3.2 XML Contentbase Connector for Java (XCC/J) Requirements	7
1.3.3 XML Contentbase Connector for .NET (XCC/.NET) Requirements	7
2.0 Programming in XCC	8
2.1 Configuring an XDBC Server	8
2.2 XCC Sessions	8
2.3 Point-In-Time Queries	9
2.4 Invoking Modules from MarkLogic Server 3.0 With XCC/J	9
2.5 Automatically Retries Exceptions	9
2.6 Coding Basics	10
3.0 Downloading and Using the XCC API	11
3.1 XCC/J Java Packages	11
3.2 XCC/.NET C# Packages	12
4.0 Using the Sample Applications	13
4.1 Setting Up Your Environment	13
4.1.1 Setting Up Your MarkLogic Server Environment	13
4.1.2 Setting Up Your Java Environment	13
4.1.3 Setting Up Your .NET Environment	14
4.2 Sample Applications	14
4.2.1 ContentFetcher	14
4.2.2 ContentLoader	15
4.2.3 HelloWorld	15
4.2.4 ModuleRunner	15
4.2.5 SimpleQueryRunner	16

Technical Support 17

1.0 Introduction to XCC

The XML Contentbase Connector (XCC) is an interface to communicate with MarkLogic Server from a Java or .NET middleware application layer. This chapter provides background on XCC and includes the following sections:

- [Overview of XCC](#)
- [API and Other Documentation](#)
- [XCC Requirements](#)

1.1 Overview of XCC

The XML Contentbase Connector (XCC) is used to communicate between a Java or .NET application layer and MarkLogic Server. Starting with MarkLogic Server 3.1, XCC replaces the XDBC libraries, and the XDBC libraries are now deprecated. You can still use XDBC 3.0 to communicate with MarkLogic Server 3.1, but if you want to use any of the new 3.1 functionality (for example, point-in-time queries), you must use XCC; XDBC will no longer have any new features added (bug fixes only), so you should use XCC.

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

- [XCC Client Libraries Communicate With an XDBC Server](#)
- [Client-Server Architecture](#)
- [Automatically Pools Connections](#)

1.1.1 XCC Client Libraries Communicate With an XDBC Server

XCC has a set of client libraries that you use to build applications that communicate with MarkLogic Server. There are Java and .NET versions of the client libraries. XCC requires that an XDBC server is configured in MarkLogic Server.

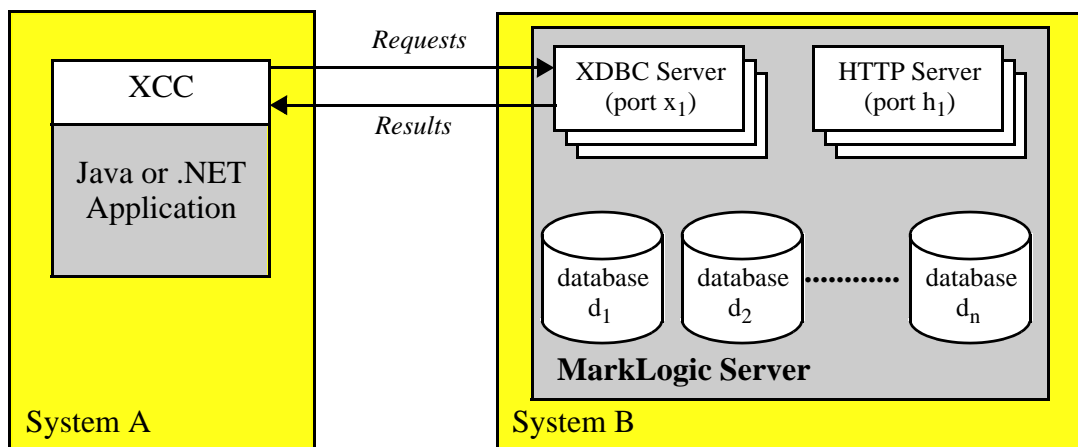
An XDBC server responds to XDBC and XCC requests. XDBC and XCC use the same wire protocol to communicate with MarkLogic Server. You can write applications either as standalone applications or ones that run in an application server environment. Your XCC-enabled application connects to a specified port on a system that is running MarkLogic Server, and communicates with MarkLogic Server by submitting requests (for example, XQuery statements) and processing the results returned by those programs. These XQuery programs can incorporate calls to XQuery functions stored and accessible by MarkLogic Server, and accessible from any XDBC-enabled application. The XQuery programs can perform the full suite of XQuery functionality, including loading, querying, updating and deleting content.

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).

1.1.2 Client-Server Architecture

XCC communicates with MarkLogic Server with a client-server architecture, where the XCC application is the client and MarkLogic Server is the server. The following figure illustrates the high-level architecture:



As shown in the diagram above, the XCC-enabled application can run on the same system as an instance of MarkLogic Server (a host), or it can run on a completely different system, as long as the two systems are networked together.

In the diagram, the XCC application running on System A has opened an XDBC connection to port x_1 on System B. On System B, MarkLogic Server is configured with an XDBC Server listening to port x_1 , and that XDBC Server connects to database d_1 . Consequently, the configuration shown in the diagram above allows the XCC application on System A to submit XQuery requests (including query, load, update and delete) for evaluation against database d_1 .

1.1.3 Automatically Pools Connections

XCC automatically does connection pooling, so you do not need to write any connection pooling logic in your application. The XCC `session` object automatically obtains and releases connections for XCC applications as needed.

1.2 API and Other Documentation

This document provides an introduction to the XCC developer libraries. For detailed API documentation for XCC and for MarkLogic Server, or to learn how to configure XDBC servers in MarkLogic Server, see the appropriate documents:

- Java API documentation (Javadoc available on developer.marklogic.com)
- .NET API documentation (available on developer.marklogic.com)
- MarkLogic Server *Developer's Guide*
- MarkLogic Server *Administrator's Guide*
- *Mark Logic Built-In and Module Functions Reference*

1.3 XCC Requirements

This section lists the requirements for XCC and has the following parts:

- [XCC MarkLogic Server Requirements](#)
- [XML Contentbase Connector for Java \(XCC/J\) Requirements](#)
- [XML Contentbase Connector for .NET \(XCC/.NET\) Requirements](#)

1.3.1 XCC MarkLogic Server Requirements

XCC requires MarkLogic Server 3.0 or MarkLogic Server 3.1. You must connect to MarkLogic Server 3.1 or later to use any of the 3.1-specific features (for example, point-in-time queries).

Note: When using XCC against MarkLogic Server 3.0, module invokes are a special case. See “Invoking Modules from MarkLogic Server 3.0 With XCC/J” on page 9.

1.3.2 XML Contentbase Connector for Java (XCC/J) Requirements

XCC/J has the following requirements:

- Java 1.4 or 1.5
- MarkLogic Server 3.0 or later (on any platform)

1.3.3 XML Contentbase Connector for .NET (XCC/.NET) Requirements

XCC/.NET has the following requirements:

- .NET runtime version 1.1
- MarkLogic Server 3.0 or later (on any platform)

2.0 Programming in XCC

XCC allows you to create multi-tier applications that communicate with MarkLogic Server as the underlying content repository. This chapter describes some of the basic programming concepts used in XCC. It includes the following sections:

- [Configuring an XDBC Server](#)
- [XCC Sessions](#)
- [Point-In-Time Queries](#)
- [Invoking Modules from MarkLogic Server 3.0 With XCC/J](#)
- [Automatically Retries Exceptions](#)
- [Coding Basics](#)

2.1 Configuring an XDBC Server

Use the Admin console to set up an XDBC server, specifying a name, port, a database to access, and other configuration parameters. For detailed instructions how to configure an XDBC Server, see the *Administrator's Guide*. You need an XDBC Server for an XCC program to communicate with MarkLogic Server.

2.2 XCC Sessions

XCC programs use the `Session` interface to set up and control communication with MarkLogic Server. XCC automatically creates and releases connections to MarkLogic Server as needed, and automatically pools the connections so that multiple requests are handled efficiently.

A `Session` handles authentication with MarkLogic Server and holds a dynamic state, but it is a lightweight object. It is OK to create and release `Session` objects as needed and as makes logical sense for your program. Do not expend effort to pool and reuse them, however, because they are not expensive to create. For example, if your program is doing multiple requests one after another, create a `Session` object at the beginning and close it when the last request is complete.

You set up the connection details with the `ContentSource` object. You can submit the connection details when you invoke the XCC program with a URL that has the following form:

```
xcc://username:password@host:port/database
```

Also, there are discrete arguments to the constructors in the API to set up any or all portions of the connection details.

2.3 Point-In-Time Queries

Point-in-time queries allow you to query older versions of content in a database. In an XCC application, you set up the options for any requests submitted to MarkLogic Server with the `RequestOptions` class. One of the options you can set is the effective point-in-time option. Therefore, to set up a query to run at a different point in time, you just set that option (the `setEffectivePointInTime` method in Java) on the `RequestOptions`. The query will then run at the specified point in time.

There are several things you must set up on MarkLogic Server in order to perform point-in-time queries. For details, see the “Point-In-Time Queries” chapter of the *Developer’s Guide*.

2.4 Invoking Modules from MarkLogic Server 3.0 With XCC/J

Although XCC was released with MarkLogic Server 3.1, it is backwards compatible with MarkLogic Server 3.0. If you are running an XCC Java application against a 3.0 version of MarkLogic Server and you are invoking modules (using the `ModuleInvoke` interface), you must set a system property on the JVM in which the application runs. You can either do that in the startup script for the JVM or in as a method call from your application. The following system property must be set to `true`:

```
xcc.module.invoke.oldstyle
```

To set this property in the JVM startup script, add the following options to the startup script:

```
-Dxcc.module.invoke.oldstyle=true
```

To set this property from your java program, use the following method:

```
System.setProperty("xcc.module.invoke.oldstyle", "true")
```

See the Java system documentation for the syntax of the `System.setProperty` method.

2.5 Automatically Retries Exceptions

Certain exceptions that MarkLogic Server throws are *retryable*; that is, the exception is thrown because of a condition that is transitory, and applications can try the request again after getting the exception. XCC will automatically retry any retryable exceptions. You can control the maximum number of retryable exceptions with the `RequestOptions` interface.

2.6 Coding Basics

To use XCC, there are several basic things you need to do in your Java or .NET code:

- Import the needed libraries.
- Set up the `ContentSource` object to authenticate against MarkLogic Server.
- Create a new `Session` object.
- Add a request to the session object.
- Get back a `ResultSequence` object from MarkLogic Server.
- Do something with the results (print them out, for example).
- Close the session.

The following are Java code samples that illustrate these basic design patterns:

```
package com.marklogic.xcc.examples;

import com.marklogic.xcc.ContentSource;
import com.marklogic.xcc.ContentSourceFactory;
import com.marklogic.xcc.Session;
import com.marklogic.xcc.Request;
import com.marklogic.xcc.ResultSequence;

URI uri = new URI("xcc://user:pass@localhost:8000/mycontent");
ContentSource contentSource =
    ContentSourceFactory.newContentSource (uri);

Session session = contentSource.newSession();

Request request = session.newAdhocQuery ("\"Hello World\"");

ResultSequence rs = session.submitRequest (request);

System.out.println (rs.asString());

session.close();
```

3.0 Downloading and Using the XCC API

The XCC API is available by downloading the XCC packages from developer.marklogic.com.

This chapter describes the basics of setting up your XCC environment, and includes the following sections:

- [XCC/J Java Packages](#)
- [XCC/.NET C# Packages](#)

For a description of the sample applications included with XCC, see “Using the Sample Applications” on page 13.

3.1 XCC/J Java Packages

The Java distribution of XCC has the following directory structure:

Document or Directory	Description
docs/	Includes the Javadoc for XCC in both expanded HTML and compressed zip format.
lib/	Contains the <code>xcc.jar</code> file, which is the XCC libraries, and the <code>xccexamples.jar</code> file, which has the compiled versions of the sample applications.
src/	Includes the source code for the sample applications.
Readme.txt	Includes the version number and any last-minute updates not included in the documentation.

3.2 XCC/.NET C# Packages

The ,NET distribution of XCC has the following directory structure:

Document or Directory	Description
<code>docs/</code>	Includes the .NET API documentation for XCC in both expanded HTML and compressed zip format.
<code>lib/</code>	Contains the <code>xcc.dll</code> file, which is the XCC libraries, and the <code>xccexamples.dll</code> file, which has the compiled versions of the sample applications.
<code>src/</code>	Includes the source code for the sample applications.
<code>Readme.txt</code>	Includes the version number and any last-minute updates not included in the documentation.

4.0 Using the Sample Applications

The XCC packages contain a number of sample applications. Each sample application is provided along with its source code, giving you a starting point for creating your own applications. This chapter describes the sample applications and contains the following sections:

- [Setting Up Your Environment](#)
- [Sample Applications](#)

4.1 Setting Up Your Environment

Before running the sample applications, be sure to set up the necessary environment to run the application. This section has the following parts:

- [Setting Up Your MarkLogic Server Environment](#)
- [Setting Up Your Java Environment](#)
- [Setting Up Your .NET Environment](#)

4.1.1 Setting Up Your MarkLogic Server Environment

Before you run the sample applications, complete the following steps:

1. Install MarkLogic Server, or have a MarkLogic Server installation to which you can connect. For details on installing MarkLogic Server, see the *Enterprise Edition Installation Guide*.
2. Create and configure an XDBC Server using the Admin interface. See the *MarkLogic Server Administrator's Guide* for details on how to create and configure an XDBC Server.
3. Configure a user for the XDBC Server you created. For example, add a user to the security database with the username as `user` and the password as `pass`. See the *MarkLogic Server Administrator's Guide* for details on adding a user to the security database.

4.1.2 Setting Up Your Java Environment

If you are using XCC/J, you must have Java installed on your client machine. Additionally, you will need the following set up to run the sample applications:

- Set your `JAVA_HOME` environment variable, if it is not already set. For example, if you are running a Windows machine, set `JAVA_HOME` in a command window as in the following example:

```
set JAVA_HOME=c:\jdk142
```

Substitute the directory in which Java is installed in your environment.

- Set your `CLASSPATH` environment variable correctly, or use the `-classpath` option to pass the appropriate classpath on the command line.

4.1.3 Setting Up Your .NET Environment

If you are using XCC/.NET, you must have the .NET runtime environment installed on your client machine.

4.2 Sample Applications

The source code and API documentation for the sample applications are included in the XCC packages. A `xccexamples.jar` file is included with the Java distribution of XCC. The sample applications are as follows:

Sample	Description
ContentFetcher	This class fetches documents from the contentbase and writes their serialized contents to a provided <code>OutputStream</code> .
ContentLoader	This program accepts a server URI (in the format expected by <code>ContentSourceFactory.newContentSource(java.net.URI)</code>) and one or more file pathnames of documents to load.
HelloWorld	This very simple class prints out the string "Hello World".
ModuleRunner	This is a very simple class that will invoke a named XQuery module on the server and return the result.
SimpleQueryRunner	This is a very simple class that will submit an XQuery string to the server and return the result.

4.2.1 ContentFetcher

This program fetches a document from MarkLogic Server and serializes its contents. You can serialize the contents to the standard output (display it on the screen) or to a file using the `-o` option. The following is a sample command to run the `ModuleRunner` class:

```
java -classpath "c:/xcc/xcc.jar;c:/xcc/xccexamples.jar"
    com.marklogic.xcc.examples.ContentFetcher
    xcc://username:password@localhost:8021
    /mydocs/hello.xml -o myHelloFile.xml
```

This sends the contents of the document at `/mydocs/hello.xml` to the file `myHelloFile.xml` (in the same directory in which the command is run). It connects to the default database of the XDBC Server listening on port 8021 of the local machine, using the credentials `username` and `password` to authenticate the connection.

4.2.2 ContentLoader

This program loads the specified document in the database. It loads the file with a URI equal to the fully-qualified pathname of the file. The following is a sample command to run the `ContentLoader` class:

```
java -classpath "c:/xcc/xcc.jar;c:/xcc/xccexamples.jar"  
com.marklogic.xcc.examples.ContentLoader  
xcc://username:password@localhost:8021 hello.xml
```

This loads the file at `hello.xml` to a document with the fully-qualified pathname of `hello.xml` (for example, `c:\xcc\examples\hello.xml`). It loads it into the default database of the XDBC Server listening on port 8021 of the local machine, using the credentials `username` and `password` to authenticate the connection.

4.2.3 HelloWorld

This program runs a query on MarkLogic Server that returns the string "Hello World". The following is a sample command to run the `HelloWorld` class:

```
java -classpath "c:/xcc/xcc.jar;c:/xcc/xccexamples.jar"  
com.marklogic.xcc.examples.SimpleQueryRunner  
xcc://username:password@localhost:8021 hello.xqy
```

4.2.4 ModuleRunner

This program allows you to invoke a module on the server. The module must exist under the XDBC server root, either in the database (when a modules database is configured) or on the filesystem (when the filesystem is configured for modules). The following is a sample command to run the `ModuleRunner` class:

```
java -classpath "c:/xcc/xcc.jar;c:/xcc/xccexamples.jar"  
com.marklogic.xcc.examples.ModuleRunner  
xcc://username:password@localhost:8021 hello.xqy
```

This invokes the module named `hello.xqy`. The request is submitted to the XDBC Server running on the local machine at port 8021, using the credentials `username` and `password` to authenticate the connection. The module path is resolved relative to the XDBC Server root.

4.2.5 SimpleQueryRunner

This program allows you to store XQuery in a file and then submit the XQuery to MarkLogic Server. The following is a sample command to run the `SimpleQueryRunner` class:

```
java -classpath "c:/xcc/xcc.jar;c:/xcc/xccexamples.jar"  
com.marklogic.xcc.examples.SimpleQueryRunner  
xcc://username:password@localhost:8021 hello.xqy
```

This submits the contents of the `hello.xqy` file to a MarkLogic Server XDBC Server running on the local machine at port 8021, using the credentials `username` and `password` to authenticate the connection.

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

If reporting a query evaluation problem, please be sure to include the sample XQuery code.