
MarkLogic Server

Release Notes

MarkLogic 5
October, 2011

Last Revised: 5.0-3, March, 2012

Table of Contents

Release Notes

1.0	Introduction	4
2.0	Installation and Upgrade	5
2.1	Supported Platforms	6
2.2	Upgrade Support	7
3.0	New Features in MarkLogic 5	8
3.1	Multi-Statement Transactions and Rollback	8
3.2	XA Distributed Transaction Support	9
3.3	Database Replication for Disaster Recovery	9
3.4	Archiving Backup Journals for Point-In-Time Recovery	9
3.5	Improved Binary Support	9
3.6	Monitoring API	9
3.7	Monitoring Dashboard	10
3.8	Query Console	10
3.9	Search API Enhancements	10
3.10	Custom Dictionaries for Tokenizing and Stemming	10
3.11	Integration with Hadoop	10
3.12	cts:field-value-query, Field Range Indexes, and Field Lexicon Functions	11
3.13	Support for Tiered Forest Storage	11
3.14	Read-Only Configuration Information	11
3.15	Exporting and Importing Configuration Information	11
3.16	Specify Output Options on the App Server Configuration	11
3.17	Additional XQuery Functions to Support SHA Hashes	12
3.18	xdmp:email Now Allows Attachments	12
3.19	New Features Added in 5.0-3	12
3.19.1	Search Equivalence Terms	12
3.19.2	Export Button on Monitoring Dashboard	12
3.19.3	xdmp:gzip and xdmp:gunzip Built-In XQuery Functions	13
4.0	Known Incompatibilities with Previous Releases	14
4.1	Quote Marks Required for Direct Attribute Constructors	14
4.2	repair-none Option of xdmp:unquote Throws Exception With Multiple Nodes ..	15
4.3	cts:entity-highlight Now Requires a Separate Download	15
4.4	XCC/J AutoCommit Methods No Longer Available	15
4.5	Post-Commit Triggers Now Execute on the D-Node	15
4.6	XCC Queries That Fetch XML Documents Return Document Nodes	16

4.7	Flexible Replication Database Need an Update to their Configuration	17
4.8	Upgrading 4.2 Custom Plugins to MarkLogic 5	17
4.9	infodev:filesystem-walk \$function Parameter Requires a Different Signature	18
4.10	Some Samples Have Moved to developer.marklogic.com	18
4.11	Docs App Server is Now Application Services App Server	18
4.12	Third Parameter of search:remove-constraint is Now Required	19
4.13	mapreduce.marklogic.output.batchsize No Longer Controls Transaction Size	19
4.14	Incompatibilities Between 4.1 and 4.2	19
4.14.1	Search Incompatibilities	19
4.14.1.1	Scores of Searches With cts:properties-query Matches are Higher	20
4.14.1.2	cts:contains Queries with a cts:properties-query and a Range Query Now Require the Needed Indexes	20
4.14.1.3	Default cts:near-query Distance Now 10	20
4.14.1.4	cts:train Uses Different Default Terms	20
4.14.2	Unicode Table Updated	21
4.14.3	XCC Jar Filed Renamed With Version Number	21
4.14.4	New mimetypes Should be Merged if You Have Customized Mimetypes ..	21
4.14.5	Default Serialization Changes for Indenting of Typed and Untyped XML (Changed in 4.2-1 to 4.2-3, the same as 4.1 in 4.2-4 and Later)	22
4.14.5.1	4.1 Behavior	22
4.14.5.2	4.2-1 to 4.2-3 Behavior	23
4.14.5.3	4.2-4 and Later Behavior	24
4.14.6	Signature Change for admin:database-geospatial-element-index and admin:database-geospatial-element-child-index	25
4.14.7	Set Journaling on Shared-Disk Failover Databases to Strict	25
4.14.8	Linux Requirements Change	25
4.14.9	Upgrade and Reindex Required	25
4.14.10	XCC AutoCommit Methods Deprecated and Cause Compile Warnings	26
5.0	Other Notes	27
5.1	Memory and Disk Space Requirements	27
5.2	Compatibility with XQuery Specifications	28
5.3	XQuery Extensions	28
5.4	Documentation	28
5.5	Browser Requirements	30
5.6	Support	31
	Product Notices	32
	COPYRIGHT	32
	TRADEMARK NOTICE	39

1.0 Introduction

MarkLogic 5 is a major release of MarkLogic Server that includes many new features. The new features are described in “New Features in MarkLogic 5” on page 8. The following lists the major features with links to where they are described:

- [Multi-Statement Transactions and Rollback](#)
- [XA Distributed Transaction Support](#)
- [Database Replication for Disaster Recovery](#)
- [Archiving Backup Journals for Point-In-Time Recovery](#)
- [Improved Binary Support](#)
- [Monitoring API](#)
- [Monitoring Dashboard](#)
- [Query Console](#)
- [Search API Enhancements](#)
- [Custom Dictionaries for Tokenizing and Stemming](#)
- [Integration with Hadoop](#)
- [cts:field-value-query, Field Range Indexes, and Field Lexicon Functions](#)
- [Support for Tiered Forest Storage](#)
- [Read-Only Configuration Information](#)
- [Exporting and Importing Configuration Information](#)
- [Specify Output Options on the App Server Configuration](#)
- [Additional XQuery Functions to Support SHA Hashes](#)
- [xdmp:email Now Allows Attachments](#)

If you are upgrading from 4.2, some applications will require minor changes to run correctly on MarkLogic 5. For details, see “Known Incompatibilities with Previous Releases” on page 14.

For a list of bugs fixed in the latest maintenance release and a list of known bugs, see the MarkLogic Technical Support website at <http://support.marklogic.com> (supported customers only).

2.0 Installation and Upgrade

This chapter describes the supported platforms and upgrade paths for MarkLogic Server, and has the following sections:

- [Supported Platforms](#)
- [Upgrade Support](#)

2.1 Supported Platforms

This release of MarkLogic Server is supported on the following platforms:

- Microsoft Windows Server 2008 (x86), Microsoft Windows 2003 Server (x86), Microsoft Windows XP SP2, Microsoft Windows Vista 32-bit Edition (x86)*
- Microsoft Windows Server 2008 (x64), Windows 2003 Server 64-bit Edition (x64), Windows Vista 64-bit Edition (x64), Windows 7 64-bit (x64)*
- Sun Solaris 10 (64-bit SPARC)
- Sun Solaris 10 (x64)
- Red Hat Enterprise Linux 5 (x86)** ***
- Red Hat Enterprise Linux 5 (x64)** *** ****
- Red Hat Enterprise Linux 6 (x64)** *** **** *****
- CentOS 5 (x64)** ***
- SUSE Linux Enterprise Server 11 (x64)** ***
- Mac OS X 10.6*****

* Microsoft Windows Vista and Windows 7 are supported for development only.

If MarkLogic Server fails to start up on Windows with the error “the application failed to initialize properly (0xc0150002)”, then a dependency is missing from your environment and you need to download and install one of the following DLLs:

32-bit versions of Windows require the DLL at the following link:

<http://www.microsoft.com/downloads/details.aspx?familyid=200B2FD9-AE1A-4A14-984D-389C36F85647&displaylang=en>.

64-bit versions of Windows require the following DLL:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=eb4ebe2d-33c0-4a47-9dd4-b9a6d7bd44da&DisplayLang=en>.

** The deadline I/O scheduler is required on Red Hat Linux platforms. The deadline scheduler is optimized to ensure efficient disk I/O for multi-threaded processes, and MarkLogic Server can have many simultaneous threads. For information on the deadline scheduler, see the Red Hat documentation (for example, <http://www.redhat.com/magazine/008jun05/features/schedulers/>).

***The `redhat-lsb`, `glibc`, and `gdb` packages are required on Red Hat Linux. Additionally, on 64-bit Red Hat Linux, both the 32-bit and the 64-bit `glibc` packages are required.

****Red Hat Linux 5 and 6 (x64) are also supported in a VMWare ESX 3.0.2, ESX 3.5.3, or ESX 4.0 (installed on bare metal) environment.

*****Mac OS X is supported for development only. Conversion (Office and PDF) and entity enrichment are not available on Mac OS X. Mac OS X 10.6 (Snow Leopard) on a 64-bit capable processor is required (<http://support.apple.com/kb/HT3696>).

*****Red Hat Enterprise Linux 6 (x64) is supported in releases 5.0-2 and later.

2.2 Upgrade Support

This section describes upgrade support to MarkLogic 5. For details on installing MarkLogic Server and for the upgrade procedure, see the *Installation Guide*.

Upgrading is supported from 4.1-1 or later. If you are running a release prior to 4.1, you must first upgrade to 4.1 or 4.2 before upgrading to MarkLogic 5. If you are upgrading an Enterprise Edition cluster, you must first upgrade the node in which the Security database forest is located before you upgrade other nodes in the cluster. MarkLogic Corporation strongly recommends performing a backup of your databases before upgrading to MarkLogic 5. Additionally, MarkLogic Corporation recommends that you first upgrade to the latest maintenance release of 4.1 or 4.2 before upgrading to MarkLogic 5.

An upgrade from 4.1 or 4.2 to MarkLogic 5 will reindex any databases that have `reindex enable` set to `true`. If you choose not to reindex your databases, they will run in either 3.0, 3.1, 3.2, 4.0, 4.1, or 4.2 compatibility mode, depending on the version of MarkLogic Server in which they were last loaded or reindexed. Running in compatibility mode will disable certain MarkLogic 5 features and may treat all content in the database as English language content. For details on database compatibility, see the *Installation Guide*.

There are some known incompatibilities between MarkLogic 5 and 4.2. You might need to make some minor code changes to your 4.2 applications before they can run correctly in MarkLogic 5. For details on the incompatibilities, see “Known Incompatibilities with Previous Releases” on page 14. For instructions on upgrading to MarkLogic 5, including information about database compatibility between MarkLogic 5 and 4.2, see the *Installation Guide*.

3.0 New Features in MarkLogic 5

This chapter describes the new features in MarkLogic 5. The feature descriptions are as follows:

- [Multi-Statement Transactions and Rollback](#)
- [XA Distributed Transaction Support](#)
- [Database Replication for Disaster Recovery](#)
- [Archiving Backup Journals for Point-In-Time Recovery](#)
- [Improved Binary Support](#)
- [Monitoring API](#)
- [Monitoring Dashboard](#)
- [Query Console](#)
- [Search API Enhancements](#)
- [Custom Dictionaries for Tokenizing and Stemming](#)
- [Integration with Hadoop](#)
- [cts:field-value-query, Field Range Indexes, and Field Lexicon Functions](#)
- [Support for Tiered Forest Storage](#)
- [Read-Only Configuration Information](#)
- [Exporting and Importing Configuration Information](#)
- [Specify Output Options on the App Server Configuration](#)
- [Additional XQuery Functions to Support SHA Hashes](#)
- [xdmp:email Now Allows Attachments](#)
- [New Features Added in 5.0-3](#)

3.1 Multi-Statement Transactions and Rollback

MarkLogic 5 includes multi-statement transactions, which allows updates made by one statement in a transaction to be visible in subsequent statements in the same transaction, before committing to the database. As part of multi-statement transactions, you can also abort a transaction and roll back the previous statements, resulting in none of the statements ever being committed to the database. In previous versions of MarkLogic, each statement is an auto-committed transaction (which is still the default mode in MarkLogic 5).

For details on transactions, including a description of how multi-statement transactions work, see the [Understanding Transactions in MarkLogic Server](#) chapter in the *Application Developer's Guide* and the Transaction `xdmp` built-in functions in the *MarkLogic XQuery and XSLT Function Reference*.

3.2 XA Distributed Transaction Support

MarkLogic 5 includes support for participating in distributed transactions by acting as a Resource Manager in an XA/JTA transaction. XA is a standard for distributed transaction processing defined by The Open Group. You can write a Java program to participate in XA transactions with MarkLogic Server and any other XA-compliant database (it could be another MarkLogic database or an XA-compliant third-party database). For details, see [Participating in XA Transactions](#) in the *XCC Developer's Guide*.

3.3 Database Replication for Disaster Recovery

MarkLogic 5 provides a facility to replicate a database on another cluster. The replication is transaction-aware and high-performance. Database Replication is a good choice for a disaster recovery plan, as you can replicate into a completely separate MarkLogic Server cluster. For details, see the [Database Replication Guide](#).

3.4 Archiving Backup Journals for Point-In-Time Recovery

MarkLogic 5 provides the ability to back up each transaction in a database after a full backup, enabling you to “roll forward” to a point in time after your full backup. Each transaction is stored as part of a journal archive that is associated with the backup, and that archive is used to replay transactions up to a point in time. For details, see [Backing Up and Restoring a Database](#) in the *Administrator's Guide*.

3.5 Improved Binary Support

MarkLogic 5 includes improved support for binary documents. There are several parts to this support:

- Support for large binaries. You can now store binaries of any size in MarkLogic.
- Support for external binaries. You can now reference a binary file outside of a MarkLogic database and access that file from within the database.
- The ability to extract text and metadata from many binary formats using the `xdmp:document-filter` built-in function.

For more details about binary document support in MarkLogic Server, see [Working With Binary Documents](#) in the *Application Developer's Guide*. For details about extracting text and metadata from binary files, see [Extracting Metadata and Text From Binary Documents](#) in the *Search Developer's Guide*.

3.6 Monitoring API

MarkLogic 5 includes a RESTful API to get monitoring information from MarkLogic Server. You can use this API with system monitoring tools such as Nagios and HP Operations Manager, both of which have deep integrations with MarkLogic Server. For more details, see the [Monitoring MarkLogic Guide](#).

3.7 Monitoring Dashboard

Using the Monitoring API, MarkLogic Server includes a monitoring dashboard, which monitors activity and resources on a MarkLogic Server cluster. The monitoring dashboard is supplemental to any monitoring you do a system level using tools such as Nagios or HP Open-View. The monitoring dashboard is a web application that runs on port 8002 of MarkLogic Server. For details on the monitoring dashboard, see [Using the MarkLogic Server Monitoring Dashboard](#) in the *Monitoring MarkLogic Guide*.

3.8 Query Console

MarkLogic 5 includes Query Console, which allows you to evaluate queries against MarkLogic Server databases. Query Console is part of Application Services, and runs on port 8000 by default along with Application Builder and Information Studio. Query Console has much of the functionality of CQ, which is a popular open-source project to evaluate queries against MarkLogic Server databases, but in a supported and integrated package with many usability enhancements. For an overview of Query Console, see [Introduction to Query Console](#) in the *Query Console User's Guide*.

3.9 Search API Enhancements

MarkLogic 5 includes enhancements to make it easy to create a facet based on a value stored in a property fragment. Additionally, you can now create value and range constraints on fields. The Search API makes it easy to create search and facet applications. For more details, see the `search:search` API documentation in *MarkLogic XQuery and XSLT Function Reference* and see [Search API: Understanding and Using](#) in the *Search Developer's Guide*.

3.10 Custom Dictionaries for Tokenizing and Stemming

MarkLogic 5 includes the ability to create language-specific custom dictionaries to modify the stemming and tokenization behavior for that language. Custom dictionaries are useful to improve search recall in cases where your content includes words are not in the stemming dictionary. For details, see [Custom Dictionaries for Tokenizing and Stemming](#) in the *Search Developer's Guide*.

3.11 Integration with Hadoop

MarkLogic 5 includes a Java API which serves as a connector to feed map-reduce jobs with data from a MarkLogic cluster into hadoop and to output the results back to MarkLogic Server.

For details, see the Hadoop JavaDoc and the *MarkLogic Connector for Hadoop Developer's Guide*.

3.12 `cts:field-value-query`, Field Range Indexes, and Field Lexicon Functions

To complement the existing `cts:field-word-query`, MarkLogic 5 includes `cts:field-value-query`, which allows you to search for entire values on fields. Additionally, you can create a range index on a field and perform lexicon functions and range queries that use that index. This functionality enables building facets on field values. For details on these functions, see the *MarkLogic XQuery and XSLT Function Reference* and [Field Value Lexicons](#) in the *Search Developer's Guide*.

3.13 Support for Tiered Forest Storage

In MarkLogic Server, you can specify three different types of storage for forests. By default, the forest data is stored in the *data directory*. Additionally, you can specify a *fast data directory*, which is typically fast solid state device (SSD) storage (for example, a PCI card that uses I/O accelerator technology), and you can also specify a large data directory, which is where large binary documents are stored. This flexibility allows you to get much of the benefit of fast expensive filesystems without the need to store all of your data on that expensive filesystem. For more details, see [Disk Storage Considerations](#) in the *Query Performance and Tuning Guide*.

3.14 Read-Only Configuration Information

MarkLogic 5 allows you to set up the system such that users without the `admin` role can view the configuration information for a cluster. The Configuration Manager allows read-only access to configuration information for databases, App Servers, and other objects in your MarkLogic Server cluster. For details, see [Using the Configuration Manager](#) in the *Administrator's Guide*.

3.15 Exporting and Importing Configuration Information

MarkLogic 5 allows you to export a configuration package containing the configuration for one cluster and then import that package into another cluster, making it easier to re-create a configuration on another cluster. You can perform the exports and imports using the Configuration Manger, which comes installed as part of Application Services on port 8000. For details, see [Packaging MarkLogic Server Configurations](#) in the *Administrator's Guide*.

3.16 Specify Output Options on the App Server Configuration

MarkLogic 5 makes it easier to customize the default output behavior of queries with a set of output options to set on HTTP and XDBC App Servers. The rich set of output options that were previously available at the individual query level are now available to set as defaults. These options correspond to the `xsl:output` settings that you can set in XSLT and in XQuery via the `xdmp:output` option. For details, see [Setting Output Options for an HTTP Server](#) in the *Administrator's Guide*.

3.17 Additional XQuery Functions to Support SHA Hashes

MarkLogic 5 includes the following built-in XQuery functions to calculate SHA hashes and hash-based message authentication SHA hashes:

- `xdmp:hmac-sha1`
- `xdmp:hmac-sha256`
- `xdmp:hmac-sha512`
- `xdmp:sha1`
- `xdmp:sha256`
- `xdmp:sha512`

For details on these functions, see the *MarkLogic XQuery and XSLT Function Reference*.

3.18 `xdmp:email` Now Allows Attachments

You can now send attachments in emails sent with `xdmp:email`. For details, see `xdmp:email` in the *MarkLogic XQuery and XSLT Function Reference*.

3.19 New Features Added in 5.0-3

MarkLogic 5 maintenance release 5.0-3 and later include the following features:

- [Search Equivalence Terms](#)
- [Export Button on Monitoring Dashboard](#)
- [xdmp:gzip and xdmp:gunzip Built-In XQuery Functions](#)

3.19.1 Search Equivalence Terms

Beginning in 5.0-3, leaf-level `cts:query` constructors take an additional option ("`synonym`") which treats all terms passed into that query constructor as the same term for scoring purposes. For example, given the following query:

```
cts:word-query(("cat", "kitty"), "synonym")
```

documents with both the words `cat` and `kitty` will be treated, for scoring purposes, as if they only had the word `cat`. Without the `synonym` option, there would be one contribution to the score from `cat` and one from `kitty`.

3.19.2 Export Button on Monitoring Dashboard

Beginning in 5.0-3, you can export data saved using the monitoring dashboard. For details, see [Exporting Monitoring Data](#) in the *Monitoring MarkLogic Guide*.

3.19.3 **xdmp:gzip and xdmp:gunzip Built-In XQuery Functions**

Beginning in 5.0-3, two new XQuery built-in functions are available:

- `xdmp:gzip`
- `xdmp:gunzip`

These functions allow you to compress and uncompress nodes using gzip compression.

4.0 Known Incompatibilities with Previous Releases

The vast majority of applications implemented on MarkLogic Server 4.2-* will run either without modifications or with very minor modifications in MarkLogic 5. There are, however, a number of changes that will cause compatibility issues with 4.2 applications. This section describes those incompatibilities and includes the following topics:

- [Quote Marks Required for Direct Attribute Constructors](#)
- [repair-none Option of xdmp:unquote Throws Exception With Multiple Nodes](#)
- [cts:entity-highlight Now Requires a Separate Download](#)
- [XCC/J AutoCommit Methods No Longer Available](#)
- [Post-Commit Triggers Now Execute on the D-Node](#)
- [XCC Queries That Fetch XML Documents Return Document Nodes](#)
- [Flexible Replication Database Need an Update to their Configuration](#)
- [Upgrading 4.2 Custom Plugins to MarkLogic 5](#)
- [infodev:filesystem-walk \\$function Parameter Requires a Different Signature](#)
- [Some Samples Have Moved to developer.marklogic.com](#)
- [Docs App Server is Now Application Services App Server](#)
- [Third Parameter of search:remove-constraint is Now Required](#)
- [mapreduce.marklogic.output.batchsize No Longer Controls Transaction Size](#)
- [Incompatibilities Between 4.1 and 4.2](#)

4.1 Quote Marks Required for Direct Attribute Constructors

In MarkLogic 5, quote marks (single or double) are required around the attribute value in a direct attribute constructor. On previous releases, the quote marks are inserted if they are not supplied. The XQuery specification requires them and, in MarkLogic 5, MarkLogic Server also requires them. The following works in 4.2, but throws a syntax error in MarkLogic 5:

```
<a href={ "link" }/>
```

If you have code that uses this pattern, you will need to rewrite to use quote marks. You can make those changes in 4.2 if you like, as they work in 4.2 as well. For example, the following works in both 4.2 and MarkLogic 5, and produces the same result in MarkLogic 5 as the previous code does in 4.2.

```
<a href="{ "link" }"/>
```

4.2 repair-none Option of xdmp:unquote Throws Exception With Multiple Nodes

The `repair-none` option (which is the default for the 1.0 and 1.0-m1 XQuery dialects) throws an exception if the string to parse is not valid XML. In MarkLogic 5, it is a little more strict than in 4.2, and if the string represents multiple root nodes, then it throws an exception. The following throws an exception in MarkLogic 5, but returns two document nodes in 4.2:

```
xdmp:unquote(' <a/><b/>')
```

If you have code that relies on the 4.2 behavior, you will need to rewrite it (for example, by breaking it into multiple calls to `xdmp:unquote` or by using the `repair-full` option).

4.3 cts:entity-highlight Now Requires a Separate Download

The `cts:entity-highlight` built-in function, which performs entity enrichment, requires that you install the entity enrichment package after installing MarkLogic 5. For details, see [Installing the Entity Enrichment Libraries](#) in the *Search Developer's Guide*.

4.4 XCC/J AutoCommit Methods No Longer Available

The previously deprecated XCC `Session.setAutoCommit` and `Session.getAutoCommit` methods are no longer available in MarkLogic 5. If you have code that relies on these methods, you must rewrite it to use the new `Session.setTransactionMode` and `Session.getTransactionMode` methods.

4.5 Post-Commit Triggers Now Execute on the D-Node

In 4.2, the module spawned from a post-commit trigger is evaluated on the node that initiated the query that caused the trigger (that is, on the evaluator node, also known as the e-node). In MarkLogic 5, the module spawned from a post-commit trigger is evaluated on the data node (d-node) that stores the forest data for the triggering document. This ensures that post-commit triggers for documents participating in distributed transactions (XA transactions) will fire properly after the distributed transaction is committed.

In most cases, this change will not cause any incompatibilities, but there is potential for incompatibilities in the following cases:

- If the d-node is in a different group than the e-node, and if the d-node expanded tree cache is set to very low values, then it is possible for the post-commit trigger code to fail with an expanded tree cache full exception. If this happens, the expanded tree cache setting might need to be increased for the d-node group.
- If your trigger action code is stored in the filesystem under the `Modules` directory, you must ensure that the code is there for all of the nodes in the cluster.

4.6 XCC Queries That Fetch XML Documents Return Document Nodes

In MarkLogic 5, XML documents returned by XCC queries return the document node in the `ResultItem`. Previously, such queries returned the root element node in the `ResultItem`. For example, in MarkLogic 5 if a `ResultItem` represents an XML document then `ResultItem.getItem()` in Java or `ResultItem.Item()` in C# returns an `XdmDocument` object. In previous releases, an `XdmElement` object is returned.

For example, in previous releases, the following code prints the name of the root element of an XML document:

```
...
import com.marklogic.xcc.types.XdmNode;
import org.w3c.dom.Node;

public static void main(String[] args) throws Exception {
    ...
    Request request =
        session.newAdhocQuery("fn:doc('somedoc.xml')");
    ResultSequence rs = session.submitRequest(request);
    ResultItem item = rs.next();

    Node node = ((XdmNode)item.getItem()).asW3cNode();
    System.out.println(node.getNodeName());
    ...
}
```

To achieve the same effect in MarkLogic 5, use `org.w3c.dom.Document.getDocumentElement`:

```
...
import com.marklogic.xcc.types.XdmDocument;
import org.w3c.dom.Document;

public static void main(String[] args) throws Exception {
    ...
    Request request =
        session.newAdhocQuery("fn:doc('somedoc.xml')");
    ResultSequence rs = session.submitRequest(request);
    ResultItem item = rs.next();

    Document doc = ((XdmDocument)item.getItem()).asW3cDocument();
    System.out.println(doc.getDocumentElement().getNodeName());
    ...
}
```

4.7 Flexible Replication Database Need an Update to their Configuration

If you are using Flexible Replication in 4.2, any database that is being replicated needs a minor configuration update. This configuration update adds a new range index and will therefore trigger a reindex on the database, so it is preferable to do this before the MarkLogic 5 reindexing occurs. To update the Flexible Replication setup for a database, run the following for each database in which you have Flexible Replication configured.

```
xquery version "1.0-ml";

import module namespace flexrep =
  "http://marklogic.com/xdmp/flexible-replication"
  at "/MarkLogic/flexrep.xqy";
import module namespace admin = "http://marklogic.com/xdmp/admin"
  at "MarkLogic/admin.xqy";

(:
  This script is for the "Documents" database: substitute the name of
  your database in which Flexible Replication is configured.
: )
admin:save-configuration(
  flexrep:configure-database(
    admin:get-configuration(),
    xdm:database("Documents"))) )
```

4.8 Upgrading 4.2 Custom Plugins to MarkLogic 5

The directory where Information Studio looks for its plugins has changed. In 4.2, the plugins were in the system plugin directory. In MarkLogic 5, they are under the directory show in [Plugin Directory](#) in the *Information Studio Developer's Guide*.

If you have plugins created in MarkLogic Server 4.2, you must move them from the old directory:

```
<marklogic-dir>/Plugins
```

to the new directory:

```
<marklogic-dir>/Assets/plugins/marklogic/appservices
```

For details on plugins in Information Studio, see [Information Studio Plugin Framework](#) in the *Information Studio Developer's Guide*.

4.9 infodev:filesystem-walk \$function Parameter Requires a Different Signature

The signature for the function that you pass into the `infodev:filesystem-walk` function in MarkLogic 5 has a different signature from 4.2. In MarkLogic 5, the first argument of the process function signature must be a `node()?` instead of a `node()`. The signature for the `$function` parameter must now be as follows:

```
declare function my:function(  
  $document as node()?,  
  $source-location as xs:string,  
  $ticket-id as xs:string,  
  $policy-deltas as element(info:options)?,  
  $context as item()?)  
  as empty-sequence()
```

If you have code that uses `infodev:filesystem-walk`, you must modify the function that is referenced in the `$function` argument to match this signature; that is, you must change the first argument of the process function to be of type `node()?` instead of `node()`.

4.10 Some Samples Have Moved to developer.marklogic.com

In MarkLogic 5, the `<marklogic-dir>/Samples` directory no longer contains the alerting sample application, the cq application (although MarkLogic 5 includes Query Console, which is installed on port 8000 as part of the installation or upgrade process), the heatmaps sample application, and the pairs sample application. Those applications are now available on developer.marklogic.com.

4.11 Docs App Server is Now Application Services App Server

MarkLogic 5 includes extended Application Services features. To accommodate those features, the Docs App Server that existed in previous release on port 8000 is renamed to App Services. Additionally, what was previously the App Services App Server on port 8002 is now the Manage App Server, and services the Management API, a REST API that drives several management and monitoring features in MarkLogic 5.

The MarkLogic 5 upgrade attempts to make these changes to your configuration. In most cases, there are only the MarkLogic-installed application running on these ports, and the upgrade will add all of the new MarkLogic 5 Application Services functionality. If you do have other applications running on these ports, however, depending on how you have things configured, the App Server configuration for those applications might change, and you would therefore need to restore that functionality by created another App Server on a different port. For details on the upgrade and on the MarkLogic 5 Application Services App Server settings, see [Prerequisites for Application Services Portion of the Upgrade](#) in the *Installation Guide*.

4.12 Third Parameter of `search:remove-constraint` is Now Required

The `search:remove-constraint` function takes an options node as a third parameter. In MarkLogic 5, the options node parameter is required. In 4.2, it was optional and it defaulted to the default options node. The change is because the default options do not include any constraints, and you can therefore not perform any useful actions with this function using the default options node. If you have code that uses the 2-parameter version of the function, you must rewrite that code to use the 3-parameter version and pass in an options node.

4.13 `mapreduce.marklogic.output.batchsize` No Longer Controls Transaction Size

As of MarkLogic Server version 5.0-3, the meaning of the MarkLogic Connector for Hadoop configuration property `mapreduce.marklogic.output.batchsize` changed from controlling transaction size to controlling the number records per request. For details, see [Controlling Transaction Boundaries](#) in the *MarkLogic Connector for Hadoop Developer's Guide*.

4.14 Incompatibilities Between 4.1 and 4.2

MarkLogic 5 allows you to upgrade either from 4.2 or 4.1. If you are upgrading from 4.0, there are some known incompatibilities between 4.0 and 4.1 that are documented in the *4.1 Release Notes*. If you are upgrading from 4.2, you can skip this section. For convenience, the incompatibilities between 4.1 and 4.2 are repeated here, and are as follows:

- [Search Incompatibilities](#)
- [Unicode Table Updated](#)
- [XCC Jar Filed Renamed With Version Number](#)
- [New mimetypes Should be Merged if You Have Customized Mimetypes](#)
- [Default Serialization Changes for Indenting of Typed and Untyped XML \(Changed in 4.2-1 to 4.2-3, the same as 4.1 in 4.2-4 and Later\)](#)
- [Signature Change for `admin:database-geospatial-element-index` and `admin:database-geospatial-element-child-index`](#)
- [Set Journaling on Shared-Disk Failover Databases to Strict](#)
- [Linux Requirements Change](#)
- [Upgrade and Reindex Required](#)
- [XCC AutoCommit Methods Deprecated and Cause Compile Warnings](#)

4.14.1 Search Incompatibilities

Most search code will function identically in 4.2 as it did in 4.1, but there are a few cases where the behavior might be slightly different. This section describes the incompatibilities in search features between 4.1 and 4.2, and includes the following incompatibilities:

- [Scores of Searches With `cts:properties-query` Matches are Higher](#)
- [`cts:contains` Queries with a `cts:properties-query` and a Range Query Now Require the Needed Indexes](#)
- [Default `cts:near-query` Distance Now 10](#)
- [`cts:train` Uses Different Default Terms](#)

4.14.1.1 Scores of Searches With `cts:properties-query` Matches are Higher

Searches that match a `cts:properties-query` in a `cts:and-query` will have higher scores than in 4.1. In 4.2, the score contributions of the `cts:properties-query` and the other `cts:query` constructors that operate on the document are added together, whereas in 4.1 the score was the maximum of the property and document scores.

4.14.1.2 `cts:contains` Queries with a `cts:properties-query` and a Range Query Now Require the Needed Indexes

If you have an expression that uses `cts:contains` with a `cts:properties-query` (or a `cts:document-fragment-query`), those expressions will now use the indexes to perform index resolution. Therefore, any `cts:query` constructors used with `cts:properties-query` OR `cts:document-fragment-query` will throw an exception if any needed indexes are not available. For example, the following query will now throw an exception if there is no range index on the element `my-element`:

```
cts:contains(//a, cts:properties-query(
  cts:element-range-query(xs:QName("my-element"), "=", "my-value")))
```

In 4.1, this query would complete even if there was no range index, but it would have to gather all of the needed fragments in order to perform the `cts:properties-query`.

4.14.1.3 Default `cts:near-query` Distance Now 10

The default distance for `cts:near-query` is now 10. Previously, the default was 100. If your applications use `cts:near-query` without specifying a `$distance` parameter, they will only match words within a 10 word distance by default. If you want them to match according to the old 100 word value, you must add a value of 100 to the `$distance` parameter.

4.14.1.4 `cts:train` Uses Different Default Terms

The default database terms that `cts:train` uses have changed in 4.2. In 4.2, the defaults are `db:stemmed-searches`, `db:fast-phrase-searches`, `db:fast-element-word-searches`, and `db:fast-element-phrase-searches`. In 4.1, the default was only `db:word-searches`. If you want the same behavior as 4.1, you can modify the options to `cts:train` to use `db:word-searches`. For most data, however, the new defaults should give better classifications.

4.14.2 Unicode Table Updated

The Unicode table has been updated to the following Unicode character mappings:

<http://unicode.org/versions/Unicode5.2.0/>

These are minor changes, but they might cause some minor incompatibilities in applications where a character was mapped to one set before and is mapped to another set now. For example, it can cause some regular expressions that matched before to not match now, or vice-versa.

4.14.3 XCC Jar Files Renamed With Version Number

The XCC jar files, used for Java client programs to connect to MarkLogic Server, now include the version number as part of the file name. The names are of the form `marklogic-xcc-4.2.1.jar`. Previously, the jar files were all named `xcc.jar`, regardless of the version. When you use the new jar file in your existing java programs, you will have to either rename the jar file before copying it to your application or modify the Java startup scripts to use the new name.

4.14.4 New mimetypes Should be Merged if You Have Customized Mimetypes

MarkLogic Server 4.2 includes the following new mimetype values:

Name	Extensions	Format	Description
<code>application/atom+xml</code>	N/A	<code>xml</code>	Mimetype for Atom Feeds.
<code>text/cache-manifest</code>	<code>manifest</code>	<code>text</code>	Mimetype for cached manifest files for offline web applications.
<code>text/csv</code>	<code>csv</code>	<code>text</code>	Mimetype for comma separated value files.
<code>text/html-sandboxed</code>	<code>sandboxed</code>	<code>text</code>	Mimetype for sandboxed html files.
<code>text/ping</code>	N/A	<code>text</code>	Mimetype for ping feeds.

If you upgraded to 4.2 from a previous version of MarkLogic Server and if you have ever customized your mimetypes (for example, using the MIME Types Configuration page of the Admin Interface), the upgrade will not automatically add the new mimetypes to your configuration. If you have not added any mimetypes, then the new mimetypes will be automatically added during the upgrade. You can check if you have these mimetypes configured by going to the Mimetype page of the Admin Interface and checking if the above mimetypes exist. If they do, then there is nothing you need to do.

If you do not have the above mimetypes after upgrading to 4.2, you can manually add the mimetypes to your configuration using the Admin Interface. To manually add the configuration, perform the following:

1. Open the Admin Interface in a browser (for example, open <http://localhost:8001>).
2. Navigate to the Mimetypes page, near the bottom of the tree menu.
3. Click the Create tab.
4. Enter the name, the extension, and the format for the mimetype (see the table above).
5. Click OK.
6. Repeat the preceding steps for each mimetype in the above table.

4.14.5 Default Serialization Changes for Indenting of Typed and Untyped XML (Changed in 4.2-1 to 4.2-3, the same as 4.1 in 4.2-4 and Later)

The default serialization of untyped XML (that is, XML for which there is no in-scope schema) has changed in 4.2. Additionally, the behavior changed again in 4.2-4. This section describes these changes as follows:

- [4.1 Behavior](#)
- [4.2-1 to 4.2-3 Behavior](#)
- [4.2-4 and Later Behavior](#)

The examples below show how you can use the new serialization capabilities in 4.2 to control the XML serialization. Note that you can add any of the 4.2 `xdrm:output` prolog statements to your 4.1 code, and they will simply be ignored in 4.1 (because `declare option` is designed to ignore unknown options, and the `xdrm:output` option does not exist in 4.1, but does in 4.2).

Note: If you are using CQ, the `xdrm:output` options shown in the examples below will be ignored and instead use the defaults (because it performs an `xdrm:eval` of your query); to see the results of the `xdrm:output` options, evaluate the query directly against an App Server.

4.14.5.1 4.1 Behavior

In 4.1, untyped XML is not indented when serialized.

```
(: in 4.1 :)
xquery version "1.0-ml" ;

<test>
  <indented>hello</indented><not/></test>

=>
<test>
  <indented>hello</indented><not/></test>
```

4.14.5.2 4.2-1 to 4.2-3 Behavior

In 4.2-3 and earlier 4.2 releases, untyped XML is indented by default. Additionally, 4.2 included the `xdmp:output indent` option which allows you to control the indentation of XML. In 4.2-3 and earlier, the `indent` option controls both typed and untyped XML (XML with a schema and without).

```
(: in 4.2-1 to 4.2-3 :)
xquery version "1.0-ml" ;

<test>
  <indented>hello</indented><not/></test>

=>
<test>
  <indented>hello</indented>
  <not/>
</test>

(: in 4.2-1 to 4.2-3 (emulates the 4.1 behavior) :)
xquery version "1.0-ml" ;
declare option xdmp:output "indent=no" ;

<test>
  <indented>hello</indented><not/></test>

=>
<test>
  <indented>hello</indented><not/></test>

(: in 4.2-1 to 4.2-3 :)
xquery version "1.0-ml" ;
declare option xdmp:output "indent=yes" ;

<test>
  <indented>hello</indented><not/></test>

=>
<test>
  <indented>hello</indented>
  <not/>
</test>
```

4.14.5.3 4.2-4 and Later Behavior

In 4.2-4 and later releases, the `indent` option no longer applies to untyped XML (XML with no in-scope schema); it only applies to typed XML (XML with an in-scope schema). An additional `xdrm:output` option—"indent-untyped=yes" or "indent-untyped=no"—controls the serialization of untyped XML. Additionally, you can enable trace events and add the trace event `OUTPUT INDENT UNTYPED` to the diagnostics page in the Admin Interface to make the default for untyped XML serialiazation be indented (like it was in 4.2-1 to 4.2-3).

```
(: in 4.2-4 and later :)
xquery version "1.0-ml" ;

<test>
  <indented>hello</indented><not/></test>

=>
<test>
  <indented>hello</indented><not/></test>

(: in 4.2-4 and later :)
xquery version "1.0-ml" ;
declare option xdrm:output "indent-untyped=no" ;

<test>
  <indented>hello</indented><not/></test>
=>
<test>
  <indented>hello</indented><not/></test>

(: in 4.2-4 and later :)
xquery version "1.0-ml" ;
declare option xdrm:output "indent-untyped=yes" ;

<test>
  <indented>hello</indented><not/></test>
=>
<test>
  <indented>hello</indented>
  <not/>
</test>

(: in 4.2-4 and later, OUTPUT INDENT UNTYED trace event enabled :)
xquery version "1.0-ml" ;

<test>
  <indented>hello</indented><not/></test>
=>
<test>
  <indented>hello</indented>
  <not/>
</test>
```

4.14.6 Signature Change for `admin:database-geospatial-element-index` and `admin:database-geospatial-element-child-index`

The signature has changed for the following functions in the Admin API:

- `admin:database-geospatial-element-index`
- `admin:database-geospatial-element-child-index`

The order of the last 2 arguments, and the `$point-format` argument is now last and is optional. If you have code that uses either of these functions, then see the function signatures in *MarkLogic XQuery and XSLT Function Reference* and determine if you need to change your code. If you use the `$point-format` argument, you will need to change the order of the last 2 arguments.

4.14.7 Set Journaling on Shared-Disk Failover Databases to Strict

In 4.2, there are two types of failover: shared-disk failover and local-disk failover. Shared-disk failover is the same as the failover available in 4.1. In 4.1, failover forests always use the `journaling` database option of `strict`, but you cannot manually set the journaling option in 4.1. In 4.2, MarkLogic Server allows you to manually set the `journaling` option at the database level, but the default is `fast`. If you are upgrading from 4.1 to 4.2 and are using failover (shared-disk failover, which was available in 4.1), to have the same behavior as 4.1, you should change the database option for `journaling` on any databases in which you have failover forests configured to `strict`. If you leave the option at the default of `fast`, your system will be protected against the MarkLogic Server process going down during transactions, but not protect against the computer going down.

For details on the journaling settings, see [Understanding the Locking and Journaling Database Settings for Bulk Loads](#) in the *Application Developer's Guide*. For details on failover, see *Scalability, Availability, and Failover Guide*.

4.14.8 Linux Requirements Change

To run 4.2 on Linux distributions, Red Hat 5 is required. MarkLogic Server will no longer start up on Red Hat 4. If you are using Red Hat 4, you must upgrade your operating system before upgrading MarkLogic Server. For details on system requirements, see “Supported Platforms” on page 6.

4.14.9 Upgrade and Reindex Required

When you log into the Admin Interface after installing 4.2, you will be prompted to upgrade the security database and the configuration files. The 4.2 security database and configuration files are not backwards compatible with 4.1. Make sure to do a full backup of your databases and data directory before upgrading. If you do not want to reindex after upgrading, turn off reindexing for each database in 4.1 before installing 4.2. For details and index compatibility, see [Upgrades and Database Compatibility](#) in the *Installation Guide*.

4.14.10 XCC AutoCommit Methods Deprecated and Cause Compile Warnings

The XCC `Session.setAutoCommit` and `Session.getAutoCommit` methods are now deprecated, and beginning with 4.2-4, they will cause compilation warnings if used in XCC code. These methods did not provide any functionality, and future releases will not include these methods. If you have any code that uses these methods, you should remove them from your code. You do not have to remove the methods immediately, however, as they will continue to exist in 4.1 and 4.2 versions of XCC. However, these methods may not exist in future versions of XCC.

5.0 Other Notes

This section provides the following information about MarkLogic Server:

- [Memory and Disk Space Requirements](#)
- [Compatibility with XQuery Specifications](#)
- [XQuery Extensions](#)
- [Documentation](#)
- [Browser Requirements](#)
- [Support](#)

5.1 Memory and Disk Space Requirements

MarkLogic Server requires at least 512 MB of system memory.

The first time it runs, MarkLogic Server automatically configures itself to the amount of memory on the system, reserving as much as it can for its own use. If you need to change the default configuration, you can manually override these defaults at a later time using the Admin Interface.

MarkLogic recommends the following two guidelines for server sizing:

- Configure your server with 1 GB of physical memory for every 16 GB of source content you expect to manage.
- Configure your server with at least one CPU (or core) per 100 GB of source content.

Pragmatically, we recommend running most configurations with a minimum of two CPUs (or two cores).

MarkLogic Server requires 3 times the disk space of the total forest size. Specifically, each forest on a filesystem requires its filesystem to have at least 3 times the forest size in disk space. This translates to three times the disk space of the source content after it is loaded. For example, if you plan on loading content that will result in a 10 GB database, reserve at least 30 GB of disk space. The disk space reserve is required for merges.

On UNIX systems, MarkLogic Server requires swap space at least equal to the amount of physical memory on the machine. Swap space equal to twice the amount of physical memory is highly recommended. For example, if you have a UNIX machine with 32 GB of memory, you should ideally configure the swap space to be 64 GB (and at least 32 GB). This is true on Windows systems also, but the system is normally set up to grow the swap (page) file as needed.

For more details about memory, disk, and swap requirements, see [Memory, Disk Space, and Swap Space Requirements](#) in the *Installation Guide*.

5.2 Compatibility with XQuery Specifications

This release implements the XQuery language, functions and operators specified in the W3C XQuery 1.0 Recommendations:

- <http://www.w3.org/TR/xquery/>
- <http://http://www.w3.org/TR/xquery-operators/>

Additionally, there is backwards compatibility with the May 2003 version of the XQuery 1.0 Draft specification used in MarkLogic Server 3.2 and previous versions. For details on the XQuery implementation in MarkLogic Server 4.1, including the three different dialects supported, see the *XQuery and XSLT Reference Guide*.

5.3 XQuery Extensions

Working within the W3C XQuery 1.0 Recommendation, MarkLogic has created a number of language extensions enabling key functionality not supported in the current release of the language specification. These extensions provide transactional update capabilities, assorted search and retrieval features, various data manipulation functions, and administrative tools.

The extensions, as well as the XQuery standard functions, are documented at <http://developer.marklogic.com>.

5.4 Documentation

MarkLogic Server includes the following documentation, available through the support web site and through <http://developer.marklogic.com/>:

Documentation	Description
<i>Installation Guide</i>	Provides procedures for installing MarkLogic Server.
<i>Administrator's Guide</i>	Provides procedures for administrative tasks such as creating servers, creating databases, backing up databases, creating users, setting up your security policy, and so on.
<i>Scripting Administrative Tasks Guide</i>	Provides information on writing code to script various administrative tasks such as creating and modifying databases, App Servers, and so on.
<i>Database Replication Guide</i>	Provides information on database replication, useful for disaster recovery scenarios.
<i>Flexible Replication Guide</i>	Provides information on Flexible Replication, useful for information sharing from one database to another.

Documentation	Description
<i>Application Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about general MarkLogic Server application development tasks.
<i>Search Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about search-based application development tasks.
<i>Application Builder Developer's Guide</i>	Provides step-by-step information on using Application Builder to build a search application, and also provides information on customizing applications built with Application Builder.
<i>Information Studio Developer's Guide</i>	Provides step-by-step information on using Information Studio to load and transform content.
<i>Monitoring MarkLogic Guide</i>	Provides information on monitoring MarkLogic Server, including using the built-in monitoring tools and integrating with external tools such as Nagios and HP Operations Manager.
<i>MarkLogic Connector for Hadoop Developer's Guide</i>	Provides information on the MarkLogic Connector for Hadoop, a Java library to help you build applications that combine MarkLogic Server and Hadoop map-reduce jobs.
<i>Content Processing Framework Guide</i>	Provides an introduction to the Content Processing Framework and procedures for installing the default content processing framework.
<i>Understanding and Using Security Guide</i>	Provides information on the role-based security model in MarkLogic Server.
<i>Query Performance and Tuning Guide</i>	Provides performance-related information that is useful to application developers and administrators.
<i>Scalability, Availability, and Failover Guide</i>	Provides information on large-scale system architecture, clustering, availability, and details on setting up shared-disk and local-disk failover.
<i>Query Console User's Guide</i>	Provides step-by-step information on using Query Console, a tool to create and run arbitrary XQuery code.
<i>MarkLogic XQuery and XSLT Function Reference</i>	API documentation for the MarkLogic built-in and module extensions to the XQuery standard functions, as well as API documentation for the W3C functions implemented in MarkLogic Server.

Documentation	Description
<i>MarkLogic 5 Release Notes</i>	Contains a summary of new features, upgrade compatible information, and known issues.
XCC Javadoc and .Net C# API Documentation	API documentation for the MarkLogic XML Contentbase Connector for Java API (XCC/J) and .Net XCC C# API documentation.
<i>XCC Developer's Guide</i>	An overview of the what you can do with the XCC libraries, examples of how to use XCC, and an overview of the sample applications included with XCC.
<i>XQuery and XSLT Reference Guide</i>	A condensed overview of the XQuery language, including information on the three XQuery dialects in MarkLogic Server. This book does include some syntax information, although it is primarily intended as in introduction and quick-reference to the language, not as a comprehensive reference.
<i>Getting Started with MarkLogic Server</i>	A quick, step-by-step overview of how to get up and running with MarkLogic Server.

XQuery language documentation is provided through the W3C working group drafts specified in “Compatibility with XQuery Specifications” on page 28. Sample code is provided through the demo server at <http://localhost:8000/>, which is automatically installed as part of the MarkLogic Server installation process. Additionally, there are many samples available on the MarkLogic developer site (<http://developer.marklogic.com>).

XQuery language extensions specific to MarkLogic Server are documented online in the *MarkLogic XQuery and XSLT Function Reference*. Example code snippets are provided as part of that documentation. The Admin Interface provides a large-scale example of complex XQuery programming, using many of the MarkLogic Server XQuery language extensions.

The Admin Interface includes built-in help screens that explain the purpose of the various controls and parameters in the Admin Interface.

Known bugs are documented online as we find them or as they are reported to us. See <http://support.marklogic.com> (supported customers only) for more details.

5.5 Browser Requirements

The MarkLogic Server Admin Interface is supported on Internet Explorer 6 or newer and Firefox browsers.

Application Builder and the applications generated with Application Builder are supported on Internet Explorer 6 or 7 on Windows, Firefox 3.0 or newer on Windows, and Safari 3 or newer on Mac OS 10.5. Other browser/platform combinations may work but are not as thoroughly tested.

Authentication functions in MarkLogic Server are supported for Internet Explorer 6 or newer and Firefox.

5.6 Support

MarkLogic provides technical support according to the terms detailed in your Software License Agreement. For evaluation licenses, MarkLogic 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 MarkLogic 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 MarkLogic 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.

Product Notices

COPYRIGHT

Copyright © 2012 MarkLogic Corporation. All rights reserved.

The MarkLogic software is protected by United States and international copyright laws, and incorporates certain third party libraries and components which are subject to the attributions, terms, conditions and disclaimers set forth below.

1. Contains SAP BusinessObjects Text Analysis XI from SAP AG. Copyright © 1996-2011. All rights reserved.
2. Highslide Software from Highslide Software Torstein Honsi. All Highslide Software is protected by local and international copyright laws. All rights reserved.
3. Icons developed by Yusuke Kamiyamane. Copyright © 2011 Yusuke Kamiyamane. All rights reserved. Icons are licensed subject to <http://creativecommons.org/licenses/by/3.0/legalcode>.
4. Antenna House OfficeHTML Copyright © 2000-2008 Antenna House, Inc. All rights reserved.
5. Argus Copyright ©1999-2008 Icenit Technology Ltd. All rights reserved.
6. Rosette Linguistics Platform 6.5.2/6.5.3 from Basis Technology Corporation, Copyright © 2004-2008 Basis Technology Corporation. All rights reserved.
7. ISYS Search. Copyright © 2011 ISYS™ Search Software, Inc. All rights reserved.
8. Software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>) Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved. Copyright © 1998-2011 The OpenSSL Project. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)" 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org. 5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project. 6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)" THIS SOFTWARE

IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com). Additional terms may apply to the foregoing software as further set forth at <http://www.openssl.org/source/license.html>

9. ICU v. 4.2.1 available from <http://site.icu-project.org/> Copyright © 1995-2011 International Business Machines Corporation and others. All rights reserved. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE. Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

10. Tidy available at <http://www.w3.org/People/Raggett/tidy/>. Copyright © 1998-2008 World Wide Web Consortium (Massachusetts Institute of Technology, European Research Consortium for Informatics and Mathematics, Keio University). All Rights Reserved.

11. 'zlib' general purpose compression library v. 1.2.3. Copyright © 1995-2010 Jean-loup Gailly and Mark Adler. http://www.zlib.net/zlib_license.html

12. RSA Data Security, Inc. MD5 Message-Digest Algorithm, copyright © 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved. License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function. License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work. RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind. These notices must be retained in any copies of any part of this documentation and/or software.

13. TRE software, available at <http://laurikari.net/tre/>. Copyright © 2001-2009 Ville Laurikari <vl@iki.fi>. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. (2) Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

14. Eigen Library, version 2.0.10, available at http://eigen.tuxfamily.org/index.php?title=Main_Page. Eigen Library is free software: you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version. Eigen Library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details. You should have received a copy of the GNU Lesser General Public License along with Eigen Library. If not, see: <http://www.gnu.org/copyleft/lesser.html>

15. FunctX XQuery Library v.1.0 available at <http://www.xqueryfunctions.com/xq/>, Copyright © Datypic (Priscilla Walmsley). You can redistribute such library and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This library is distributed in the

hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License as available at <http://www.gnu.org/licenses/lgpl-2.1.html#SEC4> for more details.

16. DeployJava.js v. 1.8, Copyright © 2006, 2011, Oracle and/or its affiliates. All rights reserved. ORACLE PROPRIETARY/CONFIDENTIAL. Use of the foregoing library is subject to the following license terms. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Oracle nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

17. YUI libraries v.2.8 available at <http://developer.yahoo.com/yui/>. Copyright © 2011 Yahoo! Inc. All rights reserved. Use of the foregoing libraries is subject to the following license terms. Redistribution and use of this software in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Yahoo! Inc. nor the names of YUI's contributors may be used to endorse or promote products derived from this software without specific prior written permission of Yahoo! Inc. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING

NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. Additional notices related to components included in the foregoing library are also available at the following URL: <http://yuilibrary.com/license/>

18. XSLTForms libraries available at <http://sourceforge.net/projects/xsltforms/>. You can redistribute such libraries and/or modify them under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. These libraries are distributed in the hope that they will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License as available at <http://www.gnu.org/licenses/lgpl-2.1.html#SEC4> for more details.

19. jQuery v.1.7 libraries available at <http://jquery.com/> and jQuery UI v.1.8.2 libraries available at <http://jqueryui.com/>. Copyright © 2011 John Resig. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

20. JSON v. 2 libraries available at <http://www.json.org/json2.js>. Copyright © 2002 JSON.org. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. The Software shall be used for Good, not Evil. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

21. CodeMirror libraries available at <http://codemirror.net/>. Copyright © 2011 by Marijn Haverbeke marijnh@gmail.com. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. **THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.**

22. cURL v. 7.22.0 available at <http://curl.haxx.se/>. Copyright © 1996 - 2011, Daniel Stenberg, <daniel@haxx.se>. All rights reserved. Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies. **THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.** Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

23. History.js available <https://github.com/balupton/History.js/> Copyright © 2011, Benjamin Arthur Lupton. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Benjamin Arthur Lupton nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. **THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,**

PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

24. Sample content provided for use with the software is copyrighted by the respective authors as contributed by each of them to Wikipedia pursuant to the Creative Commons Attribution-ShareAlike License found at <http://creativecommons.org/licenses/by-sa/3.0/>.

25. Visual Studio Tools for Office 3.0 (Microsoft VSTO 3.0) licensed under Microsoft Public License Microsoft Public License (Ms-PL) found at <http://msdn.microsoft.com/en-us/library/ff647676.aspx>.

26. The MarkLogic Connector for Hadoop contains Apache Jakarta Commons Modeler. Copyright 2001-2007 The Apache Software Foundation. Licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

27. jemalloc v 2.2.4. <http://www.canonware.com/jemalloc/index.html>, released under the terms of the following BSD-derived licenses: Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice(s), this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice(s), this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER(S) ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER(S) BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. Portions of the software are copyright (C) 2002-2011 Jason Evans, and portions copyright (C) 2009-2011 Facebook, Inc. and copyright (C) 2007-2010 Mozilla Foundation. Neither the name of Facebook, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. All rights reserved.

28. GzipHelper; version: 1.0, Feb, 2003, authored by Gao Dasheng, is distributed subject to the following terms: This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software. Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions: (1.) the origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required; (2) altered source versions must be plainly marked as such, and must not be misrepresented as being the original software; (3) this notice may not be removed or altered from any source distribution.

TRADEMARK NOTICE

The MarkLogic name and logo are registered trademarks of MarkLogic Corporation in the United States and/or other countries. Excel and PowerPoint are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Any other names or logos included in the Software, this notice or the documentation are property of the respective trademark owners.