MarkLogic Server

Release Notes

MarkLogic 6
September, 2012

Last Revised: 6.0-6, June, 2015

Copyright © 2015 MarkLogic Corporation. All rights reserved.

Table of Contents

Release Notes

1.0 Introduction 5

2.0 Installation and Upgrade 6
- 2.1 Supported Platforms 6
- 2.2 Certified Filesystems 7
- 2.3 Upgrade Support 7

3.0 New Features in MarkLogic 6 8
- 3.1 Business Intelligence Tool Integration 8
- 3.2 REST API 8
- 3.3 MarkLogic Java API 8
- 3.4 MarkLogic Content Pump (mlcp) 9
- 3.5 Search API Enhancements 9
- 3.6 Visualization Widgets 9
- 3.7 User-Defined Aggregate Functions 9
- 3.8 Enhanced JSON Support 9
- 3.9 In-Database Analytic Functions 10
- 3.10 Path Range Indexes 10
- 3.11 FIPS 140-2 Cryptographic Compliance 10
- 3.12 Synonym Search 10
- 3.13 Hadoop Connector Enhancements 10
- 3.14 New Features Added in MarkLogic Server 6.0-3 11
  - 3.14.1 REST API Support for Dynamic Query Options 11
  - 3.14.2 Java API Support for Dynamic Query Options 11
  - 3.14.3 Java API Support for Raw Structured Query Construction 12
  - 3.14.4 mlcp Ingestion Fault Tolerance 12
  - 3.14.5 Hadoop Connector Ingestion Fault Tolerance 12

4.0 Known Incompatibilities with Previous Releases 13
- 4.1 32-bit and Solaris SPARC Platforms are No Longer Supported 13
- 4.2 Default Database Settings Changes 13
- 4.3 Serialization Changes for xdmp:from-json and xdmp:to-json 14
- 4.4 fn:format-dateTime Functions Now Round to 7 Places Instead of 3 14
- 4.5 xdmp:sha1, xdmp:sha256, xdmp:sha512 Not Supported in FIPS Mode 14
- 4.6 Values of Type xs:date Sometimes Serialize With Timezone 15
- 4.7 DocumentInputFormat Default Input Value Class is MarkLogicDocument 15
- 4.8 Application Builder Generated App Differences 16

4.9 app-user Role Now Has the rest-reader Role 16
4.10 fn:resolve-uri in XSLT Now Requires a Scheme When Using the Second Parmeter 17
4.11 MarkLogic Content Pump 1.0-2 Requires MarkLogic Server 6.0-2 17
4.12 Starting With 6.0-2, XCC Has a New insertContentCollectErrors() Method on the Java Session Interface 17
4.13 New Aggregate Input Default URI in mlcp 1.0-3 18
4.14 As of 6.0-5, HDP is no longer a supported Hadoop platform 18
4.15 Incompatibilities Between 4.2 and 5.0 18
4.15.1 Quote Marks Required for Direct Attribute Constructors 19
4.15.2 repair-none Option of xdmp:unquote Throws Exception With Multiple Nodes 19
4.15.3 cts:entity-highlight Now Requires a Separate Download 19
4.15.4 XCC/J AutoCommit Methods No Longer Available 19
4.15.5 Post-Commit Triggers Now Execute on the D-Node 20
4.15.6 XCC Queries That Fetch XML Documents Return Document Nodes 20
4.15.7 Flexible Replication Database Need an Update to their Configuration 21
4.15.8 Upgrading 4.2 Custom Plugins to MarkLogic 6 21
4.15.9 infodev:filesystem-walk $function Parameter Requires a Different Signature 22
4.15.10 Some Samples Have Moved to developer.marklogic.com 22
4.15.11 Docs App Server is Now Application Services App Server 22
4.15.12 Third Parameter of search:remove-constraint is Now Required 23
4.15.13 mapreduce.marklogic.output.batchsize No Longer Controls Transaction Size 23
4.16 Incompatibilities Between 4.1 and 4.2 23
4.16.1 Search Incompatibilities 24
4.16.1.1 Scores of Searches With cts:properties-query Matches are Higher 24
4.16.1.2 cts:contains Queries with a cts:properties-query and a Range Query Now Require the Needed Indexes 24
4.16.1.3 Default cts:near-query Distance Now 10 24
4.16.1.4 cts:train Uses Different Default Terms 25
4.16.2 Unicode Table Updated 25
4.16.3 XCC Jar Filed Renamed With Version Number 25
4.16.4 New mimetypes Should be Merged if You Have Customized Mimetypes .. 25
4.16.5 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) 26
4.16.5.1 4.1 Behavior 27
4.16.5.2 4.2-1 to 4.2-3 Behavior 27
4.16.5.3 4.2-4 and Later Behavior 28
4.16.6 Signature Change for admin:database-geospatial-element-index and admin:database-geospatial-element-child-index 29
4.16.7 Set Journaling on Shared-Disk Failover Databases to Strict 29
4.16.8 Linux Requirements Change 29

4.16.9 Upgrade and Reindex Required 30
4.16.10XCC AutoCommit Methods Deprecated and Cause Compile Warnings .30

5.0 Other Notes 31
5.1 Memory and Disk Space Requirements 31
5.2 Compatibility with XQuery Specifications 32
5.3 XQuery Extensions 32
5.4 Documentation 32
5.5 Browser Requirements 35
5.6 Support 35

Combined Product Notices 36
COPYRIGHT 36
TRADEMARK NOTICE 48

1.0 Introduction

MarkLogic 6 is a major release of MarkLogic Server that includes many new features. The new features are described in “New Features in MarkLogic 6” on page 8. The following lists the major features with links to where they are described:

- Business Intelligence Tool Integration
- REST API
- MarkLogic Java API
- MarkLogic Content Pump (mlcp)
- Search API Enhancements
- Visualization Widgets
- User-Defined Aggregate Functions
- Enhanced JSON Support
- In-Database Analytic Functions
- Path Range Indexes
- FIPS 140-2 Cryptographic Compliance
- Synonym Search

If you are upgrading from 5.0, some applications will require minor changes to run correctly on MarkLogic 6. For details, see “Known Incompatibilities with Previous Releases” on page 13.

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
- Certified Filesystems
- Upgrade Support

2.1 Supported Platforms

MarkLogic Server is supported on the following platforms:

- Microsoft Windows Server 2008 (x64), Windows Vista 64-bit Edition (x64), Windows 7 64-bit (x64)*
- Sun Solaris 10 (x64)
- Red Hat Enterprise Linux 5 (x64)** *** ****
- Red Hat Enterprise Linux 6 (x64)** *** ****
- SUSE Linux Enterprise Server 11 (x64)** ***
- CentOS 5 (x64)** ***
- CentOS 6 (x64)** ***
- Mac OS X 10.6 or 10.7*****

* 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 the following DLL for 64-bit versions of Windows: 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, ESX 4.0, or ESXi 5.0 (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 or 10.7 (Snow Leopard or Lion) on a 64-bit capable processor is required (http://support.apple.com/kb/HT3696).

2.2 Certified Filesystems

MarkLogic relies on the operating system for filesystem operations. While any filesystem that works properly (including under heavy load) should work, the following table lists the operating systems along with the filesystems under which they are certified. Other filesystems may work but have not been thoroughly tested by MarkLogic.

Operating System	Certified Filesystems
Linux (all varieties)	EXT3, EXT4, as well as the clustered filesystems for shared-disk failover mentioned in Requirements for Shared-Disk Failover in the *Scalability, Availability, and Failover Guide*.
Sun Solaris	UFS, as well as the clustered filesystems for shared-disk failover mentioned in Requirements for Shared-Disk Failover in the *Scalability, Availability, and Failover Guide*.
Windows	NTFS
Mac OS	HFS+

2.3 Upgrade Support

This section describes upgrade support to MarkLogic 6. 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 6. 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 6. Additionally, MarkLogic Corporation recommends that you first upgrade to the latest maintenance release of 4.1 or 4.2 before upgrading to MarkLogic 6.

An upgrade from 4.1 or 4.2 to MarkLogic 6 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 6 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 6 and 4.2. You might need to make some minor code changes to your 4.2 applications before they can run correctly in MarkLogic 6. For details on the incompatibilities, see “Known Incompatibilities with Previous Releases” on page 13. For instructions on upgrading to MarkLogic 6, including information about database compatibility between MarkLogic 6 and 4.2, see the *Installation Guide*.

3.0 New Features in MarkLogic 6

This chapter describes the new features in MarkLogic 6. The feature descriptions are as follows:

- Business Intelligence Tool Integration
- REST API
- MarkLogic Java API
- MarkLogic Content Pump (mlcp)
- Search API Enhancements
- Visualization Widgets
- User-Defined Aggregate Functions
- Enhanced JSON Support
- In-Database Analytic Functions
- Path Range Indexes
- FIPS 140-2 Cryptographic Compliance
- Synonym Search
- Hadoop Connector Enhancements
- New Features Added in MarkLogic Server 6.0-3

3.1 Business Intelligence Tool Integration

MarkLogic 6 allows integration via ODBC with the Business Intelligence (BI) tools IBM Cognos and Tableau. The ODBC connections are made to a new kind of App Server in MarkLogic Server, an ODBC App Server. You can create SQL views on data stored in range indexes, and then use the BI tools to create reports and visualizations on the data. For details, see the *SQL Data Modeling Guide*.

3.2 REST API

MarkLogic 6 includes a REST API that allows you to perform searches, create documents, read documents, update documents, and delete documents. The REST API allows you to build fully functional MarkLogic applications without writing any XQuery code. It also allows you to directly load JSON documents. For details, see the *REST Application Developer's Guide*.

3.3 MarkLogic Java API

MarkLogic 6 includes a new Java API which allows you full-featured access to MarkLogic functionality without writing any XQuery code. The MarkLogic Java API is written on top of the REST API, and has all of its functionality such as paginated search with facets and snippets, full document CRUD operations, and so on. For details, see the *Java Application Developer's Guide*.

3.4 MarkLogic Content Pump (mlcp)

MarkLogic 6 include MarkLogic Content Pump (mlcp). mlcp is a command-line tool for loading content into MarkLogic Server and for migrating content from one instance of MarkLogic to another, even if they are on different platforms. If you have a Hadoop cluster, mlcp takes advantage of Hadoop to parallelize the loading. mlcp takes much of the functionality of the open source projects Record Loader and xqsync and bundles them in a single package, and allows them to take advantage of Hadoop if it is available; Hadoop is not required to use mlcp, but is used if it is available. For information on mlcp, see Using MarkLogic Content Pump in the *Loading Content Into MarkLogic Server Guide*.

3.5 Search API Enhancements

MarkLogic 6 includes several enhancement to the search API, including:

- Structured Search
- Extracting metadata at search time
- Modify the unconstrained term behavior using `<search:term>`
- Range constraints for path range indexes
- Ability to return values from range indexes with `search:values`
- JSON key support

For details on the Search API, see Search API: Understanding and Using in the *Search Developer's Guide*.

3.6 Visualization Widgets

MarkLogic 6 includes visualization widgets that are driven off of data from MarkLogic. You can access these widgets with Application Builder.

3.7 User-Defined Aggregate Functions

MarkLogic 6 includes a new facility to run user-defined aggregate functions (UDF) with custom C or C++ code that is dynamically linked into MarkLogic Server. The custom code has access to range index data on each data node, and you can use this to perform parallel in database map-reduce jobs. For details, see User-Defined Functions in the *Application Developer's Guide*.

3.8 Enhanced JSON Support

MarkLogic 6 includes a new JSON library that makes it easy to store JSON documents as key-value stores, and to convert them back and forth between JSON and XML. The REST Client API and the MarkLogic Client API for Java make use of this functionality to make it easier to load and work with JSON documents. This includes the `json:transform-from-json` and `json:transform:to-json` XQuery APIs as well as low-level XQuery support. For details, see Working With JSON in the *Application Developer's Guide*.

3.9 In-Database Analytic Functions

MarkLogic 6 includes several built-in XQuery functions to perform analytic and statistical functions. These functions are in the `cts` and `math` namespace. For details, see the Math Lexicon and Math function references.

3.10 Path Range Indexes

MarkLogic 6 includes support for range index specified by a path. You can specify a subset of XPath as the definition of what goes into an index. For example, consider the following path:

```
/foo/bar
```

If you specify this path in a Path Range Index configuration, the resulting range index includes `bar` values that have a `foo` parent, where `foo` is the root element of the document. This allows you to create much finer grained range indexes than an element or an element-attribute range index, and still have the advantages of being able to use the lexicon functions and the range query constructors. The search API can take advantage of path range indexes to create range constraints on them. Path range indexes are also useful when setting up SQL views on data stored in a MarkLogic database.

3.11 FIPS 140-2 Cryptographic Compliance

MarkLogic 6 includes the OpenSSL Federal Information Processing Standard (FIPS) Object Module, which was evaluated by the National Institute of Standards (NIST) for FIPS 140-2 compliance. For details on the OpenSSL Object Module and on the FIPS 140-2 compliance, see http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140sp/140sp1051.pdf.

MarkLogic Server operates in FIPS mode by default, and FIPS mode ensures that the server only uses secure cryptography. The cryptography is used with SSL/TLS connections to App Servers. For details on configuring FIPS mode, see Clusters in the *Administrator's Guide*.

3.12 Synonym Search

MarkLogic 6 includes synonym search, where 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`.

3.13 Hadoop Connector Enhancements

The MarkLogic Connector for Hadoop v1.1 has the following new features:

- New configuration property: `mapreduce.marklogic.input.indented`. When set to true, XML documents is pretty-printed as it is extracted from MarkLogic Server.
- New configuration property `mapreduce.marklogic.output.content.namespace`. Specifies a namespace declaration to be added to XML documents inserted into MarkLogic Server.
- New classes are added to support querying lexicon and range index values. See the following classes in the package `com.marklogic.mapreduce.functions`: `PathReference`, `ValueCooccurrences`, `ValueMatch`, and `Values`.
- Beginning with MarkLogic Server 6.0-2, the MarkLogic Connector for Hadoop supports the Hortonworks Data Platform, installed using the MarkLogic Connector for Hortonworks Data Platform. For details, see Using the MarkLogic Connector for the Hortonworks Data Platform in *MarkLogic Connector for Hadoop Developer's Guide*.

3.14 New Features Added in MarkLogic Server 6.0-3

The following new features are introduced in MarkLogic Server 6.0-3:

- REST API Support for Dynamic Query Options
- Java API Support for Dynamic Query Options
- Java API Support for Raw Structured Query Construction
- mlcp Ingestion Fault Tolerance
- Hadoop Connector Ingestion Fault Tolerance

3.14.1 REST API Support for Dynamic Query Options

The REST API now supports *combined query*, a new query format that enables you to customize a query using query options, without pre-configuring persistent query options. Combined queries are supported on the following methods:

- `POST /v1/search`
- `POST /v1/values/{name}`

The new POST method for `/v1/values/{name}` also accepts a structured query in the POST body.

For details, see Specifying Dynamic Query Options with Combined Query in *REST Application Developer's Guide*.

3.14.2 Java API Support for Dynamic Query Options

The REST API now supports *combined query*, a new query format that enables you to customize a query using query options, without pre-configuring persistent query options. Combined queries are supported through the class `com.marklogic.client.query.RawCombinedQueryDefinition`.

For details, see Apply Dynamic Query Options to Document Searches in *Java Application Developer's Guide*.

3.14.3 Java API Support for Raw Structured Query Construction

You can use the new class `com.marklogic.client.query.RawStructuredQueryDefinition` to construct a structured query directly in XML or JSON instead of using a query builder.

For details see Creating a Structured Query From Raw XML or JSON in *Java Application Developer's Guide*.

3.14.4 mlcp Ingestion Fault Tolerance

Marklogic Content Pump (mlcp) has a new `-tolerate_errors` option which determines whether or not ingestion continues if an error occurs on one or more documents in a batch. By default, mlcp can reject an entire batch if any document in the batch causes an error.

For details, see Import Options in *Loading Content Into MarkLogic Server Guide*.

3.14.5 Hadoop Connector Ingestion Fault Tolerance

MarkLogic Connector for Hadoop has a new output configuration property that determines whether or not ingestion into the database continues if an error occurs on one or more documents in a batch. By default, an entire batch can be rejected if any document in the batch causes an error.

For details, see `com.marklogic.output.content.tolerateerrors` in Output Configuration Properties in *MarkLogic Connector for Hadoop Developer's Guide*.

4.0 Known Incompatibilities with Previous Releases

The vast majority of applications implemented on MarkLogic Server 5.0-* will run either without modifications or with very minor modifications in MarkLogic 6. There are, however, a number of changes that will cause compatibility issues with 5.0 applications. This section describes those incompatibilities and includes the following topics:

- 32-bit and Solaris SPARC Platforms are No Longer Supported
- Default Database Settings Changes
- Serialization Changes for xdmp:from-json and xdmp:to-json
- fn:format-dateTime Functions Now Round to 7 Places Instead of 3
- xdmp:sha1, xdmp:sha256, xdmp:sha512 Not Supported in FIPS Mode
- Values of Type xs:date Sometimes Serialize With Timezone
- DocumentInputFormat Default Input Value Class is MarkLogicDocument
- Application Builder Generated App Differences
- app-user Role Now Has the rest-reader Role
- fn:resolve-uri in XSLT Now Requires a Scheme When Using the Second Parmeter
- MarkLogic Content Pump 1.0-2 Requires MarkLogic Server 6.0-2
- Starting With 6.0-2, XCC Has a New insertContentCollectErrors() Method on the Java Session Interface
- As of 6.0-5, HDP is no longer a supported Hadoop platform
- New Aggregate Input Default URI in mlcp 1.0-3
- Incompatibilities Between 4.2 and 5.0
- Incompatibilities Between 4.1 and 4.2

4.1 32-bit and Solaris SPARC Platforms are No Longer Supported

MarkLogic 6 discontinues support for all 32-bit platforms (32-bit Windows and 32-bit Linux) and also discontinues support for Sun Solaris SPARC. 64-bit operating systems are required. Sun Solaris is supported on x64 platforms only. If you are using one of the discontinued platforms, you will need to migrate your environment to a supported platform. For details on supported platforms, see “Supported Platforms” on page 6.

4.2 Default Database Settings Changes

There are several changes to the default settings when you create a new database. The changes are to turn off some features that you might not need by default and to enable the URI lexicon by default. If you need the feature, you can enable it when or after you create the database configuration. The following are the new default values for the settings that changed:

- `maintain last modified` is set to `false`
- `directory creation` is set to `manual`
- `inherit permissions` is set to `false`
- `inherit collections` is set to `false`
- `inherit quality` is set to `false`
- `URI lexicon` is set to `true`

If you plan to use WebDAV on a database, then set `directory creation` to `automatic`. If you want any of the other options to behave as they did with the previous defaults, then modify them when you create the database or after you create the database. If you have scripts that relied on the old default values, you will need to update the scripts to modify the values.

4.3 Serialization Changes for xdmp:from-json and xdmp:to-json

The `xdmp:from-json` built-in function serializes as xml that is a serialization of a `json:array` in MarkLogic 6, with key/value pairs expressed as `json:objects`. Previously, they were serialized as a `map:map` and as sequences. Also, null values behave differently and are now included in the JSON output from `xdmp:to-json`. JSON array objects are now returned as `json:arrays` instead of sequences. Because of these changes, any code you have that uses `xdmp:to-json` and `xdmp:from-json` will not work the same way as it did in MarkLogic 5, and will need rewriting.

If you have code that uses these functions, consider rewriting it to use the new json:transform-from-json and json:transform-to-json XQuery library modules. For details, see Working With JSON in the *Application Developer's Guide*.

4.4 fn:format-dateTime Functions Now Round to 7 Places Instead of 3

In MarkLogic 6, the `fn:format-dateTime` function, as well as all of the other format functions that show fractional seconds, rounds the seconds to 7 fractional digits. Previously, it rounded to 3 fractional digits, so it is now more precise than it used to be. If you have code that relies on the old rounding behavior, you might see different (and more accurate) answers in MarkLogic 6.

4.5 xdmp:sha1, xdmp:sha256, xdmp:sha512 Not Supported in FIPS Mode

MarkLogic 6 runs by default in FIPS mode, which require ciphers and other security requirements that are secure according to FIPS 140-2. The `xdmp:sha1`, `xdmp:sha256`, and `xdmp:sha512` functions are not FIPS 140-2 compliant, and are therefore not allowed in the default FIPS mode; in FIPS mode, these functions throw an exception. If you want to use these functions, you will need to turn FIPS mode off. Alternately, you can change your code to use the more secure `xdmp:hmac-sha1`, `xdmp:hmac-sha256`, or `xdmp:hmac-sha512` functions, which are supported in FIPS mode.

4.6 Values of Type xs:date Sometimes Serialize With Timezone

In MarkLogic 6, some range index values that have timezones associated with them include the timezone when retrieved through lexicon functions (`cts:element-values`, `cts:element-attribute-values`, and so on). This difference only occurs for date values that have a timezone different from the current timezone, and only when stored in a range index. Previously, those date values did not include the timezone. For example,

```
xquery version "1.0-ml";

(:
   assumes an element range index is
   configured for the element "hello"
:)
xdmp:document-insert("/date.xml",
<hello>{fn:adjust-date-to-timezone(xs:date("2012-08-10"),
        xs:dayTimeDuration("-PT5H"))}</hello>);

cts:element-values(xs:QName("hello"), (), "type=date")

(:
   returns 2012-08-10-05:00 in MarkLogic 6
   returns 2012-08-10 in MarkLogic 5
:)
```

If you have date range index values stored in range indexes and if your code expects those values to not have timezones, then you will need to modify the code to handle timezones when they occur.

4.7 DocumentInputFormat Default Input Value Class is MarkLogicDocument

In MarkLogic Connector for Hadoop 1.1 the default input value class for `com.marklogic.mapreduce.MarkLogicDocument` is `com.marklogic.mapreduce.DocumentInputFormat`. In 1.0, the default input value class was `com.marklogic.mapreduce.MarkLogicNode`. To restore the 1.0 behavior, add the following to your MapReduce configuration file:

```
<property>
  <name>mapreduce.marklogic.input.valueclass</name>
  <value>com.marklogic.mapreduce.MarkLogicNode</value>
</property>
```

4.8 Application Builder Generated App Differences

The applications generated from Application Builder are considerably different in MarkLogic 6 than they were in MarkLogic 5. In MarkLogic 6, the applications use the REST API and do not have any XQuery code. The following are some of the differences:

- The generated applications use the REST API and are hosted on a REST API instance.
- The generated applications are now 100% html and javascript; there is no XQuery code directly in the application.
- The default quality weight is now 0. Previously, it was 1. The new default value of 0 is the default for the REST API (which applications generated with Application Builder use) makes search performance faster.
- Any customizations built in MarkLogic 5 will be lost if you re-generate the application in MarkLogic 6.

If you have customizations of your generated applications or if you rely on any of the old application features, then you will need to rework those in MarkLogic 6. For details on Application Builder and the applications it generates, see the *Application Builder Developer's Guide*.

4.9 app-user Role Now Has the rest-reader Role

In MarkLogic 6, the `app-user` role now inherits the `rest-reader` role. Previously it did not have this role. The REST API uses the `rest-reader` role by default to put permissions on documents, so when you create documents without explicitly specifying permissions using the REST API or the Java API, users with the `rest-reader` role (and therefore uses with the `app-user` role because it inherits `rest-reader`) can read the documents. This should not cause an incompatibility for existing applications, as the `rest-reader` role did not exist in previous versions of MarkLogic, but those users might be able to read new documents created in MarkLogic 6 using the REST API or the Java API. If you have users who have the `app-user` role, be aware that they will have this additional role.

4.10 fn:resolve-uri in XSLT Now Requires a Scheme When Using the Second Parmeter

If you use `fn:resolve-uri` in an XSLT stylesheet, it now runs in strict mode and therefore, if you sepcify the second parameter, the second parameter must be a valid base URI and include a scheme; that is, it cannot start with a slash. The following query demonstrates this change:

```
xquery version "1.0-ml";

xdmp:xslt-eval(
<xsl:stylesheet version="2.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:template match="/">
  <xsl:value-of select="resolve-uri('../foo','/bar/boo/bat.xml')"/>
</xsl:template>
</xsl:stylesheet>, document{<_/>})
(:
  returns /bar/foo in 5.0, throws
  XDMP-RELURIBASEARG  in 6.0
:)
```

If you have XSLT code that relies on the old behavior, you will have to modify your code so it does not throw `XDMP-RELURIBASEARG`. You can put a scheme (such as `http://`) in the second parameter, call the function in an XQuery `1.0-ml` module, or use another means besides `fn:resolve-uri`.

4.11 MarkLogic Content Pump 1.0-2 Requires MarkLogic Server 6.0-2

MarkLogic Content Pump (mlcp) 1.0 should not be used with MarkLogic Server 6.0-2 or later. Use mlcp 1.0-2 or later with MarkLogic Server 6.0-2 or later.

4.12 Starting With 6.0-2, XCC Has a New insertContentCollectErrors() Method on the Java Session Interface

Starting with 6.0-2, XCC has a new method on the `Session` interface. If you recompile a program with the 6.0-2 or later jar file, and if that program implements the `Session` interface, then you will need to add the `insertContentCollectErrors()` method to your source code so your interface implements this function. For example, the open source XQSync program implements the `Session` interface, so if you have old source code for it, when you recompile it you will need to implement the new method in its `Session` interface. You should only need to do this if you are recompiling; if you are not recompiling the program, then it should continue to work with the new XCC jar.

4.13 New Aggregate Input Default URI in mlcp 1.0-3

The following change might affect you if you use mlcp to load documents from aggregate input (`-input_file_type aggregates`) and rely on the default URI generation. If you use `-aggregate_uri_id` when loading from aggregate files, you should not be affected by this change.

Documents created with earlier versions of mlcp used the following default URI template:

```
hash_or_taskid-seqnum
```

Starting with mlcp version 1.0-3, the default URI template for documents created from aggregates is:

```
/path/filename-split_start-seqnum
```

For details, see Default Document URI Construction in *Loading Content Into MarkLogic Server Guide*.

4.14 As of 6.0-5, HDP is no longer a supported Hadoop platform

Starting with 6.0-5, Hortonworks Data Platform (HDP) is no longer a supported Hadoop distribution for use with the MarkLogic Connector for Hadoop or the distributed mode of MarkLogic Content Pump (mlcp).

4.15 Incompatibilities Between 4.2 and 5.0

MarkLogic 6 allows you to upgrade either from 5.0, 4.2 or 4.1. If you are upgrading from 4.0, you must first upgrade to at least 4.1, and 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 5.0, you can skip this section. For convenience, the incompatibilities between 5.0 and 4.2 are repeated here, and are as follows:

- 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 6
- 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

4.15.1 Quote Marks Required for Direct Attribute Constructors

In MarkLogic 5 and later, 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 and later, MarkLogic Server also requires them. The following works in 4.2, but throws a syntax error in MarkLogic 5 and later:

```
<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 later, and produces the same result in MarkLogic 5 and later as the previous code does in 4.2.

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

4.15.2 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-ml` XQuery dialects) throws an exception if the string to parse is not valid XML. In MarkLogic 5 and later, 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 and later 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).

4.15.3 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 and later. For details, see Installing the Entity Enrichment Libraries in the *Search Developer's Guide*.

4.15.4 XCC/J AutoCommit Methods No Longer Available

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

4.15.5 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 and later, 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.

4.15.6 XCC Queries That Fetch XML Documents Return Document Nodes

In MarkLogic 5 and later, 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 and later 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 and later, 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());
    ...
}
```

4.15.7 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 and later 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(),
    xdmp:database("Documents")))
```

4.15.8 Upgrading 4.2 Custom Plugins to MarkLogic 6

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 and later, 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*.

4.15.9 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 and later has a different signature from 4.2. In MarkLogic 5 and later, 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()`.

4.15.10 Some Samples Have Moved to developer.marklogic.com

In MarkLogic 5 and later, the `<marklogic-dir>/Samples` directory no longer contains the alerting sample application, the cq application (although MarkLogic 5 and later 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.

4.15.11 Docs App Server is Now Application Services App Server

MarkLogic 5 and later 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 6.

The MarkLogic 5 and later 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 and later 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 and later Application Services App Server settings, see Prerequisites for Application Services Portion of the Upgrade in the *Installation Guide*.

4.15.12 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 and later, 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.

4.15.13 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*.

4.16 Incompatibilities Between 4.1 and 4.2

MarkLogic 6 allows you to upgrade from 5.0, 4.2 or 4.1. If you are upgrading from 4.0, you must first upgrade to at least 4.1, and 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 5.0, 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

4.16.1 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

4.16.1.1 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.

4.16.1.2 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`.

4.16.1.3 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.

4.16.1.4 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.

4.16.2 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.

4.16.3 XCC Jar Filed 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.

4.16.4 New mimetypes Should be Merged if You Have Customized Mimetypes

MarkLogic Server 4.2 includes the following new mimetype values:

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

4.16.5 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 `xdmp: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 `xdmp:output` option does not exist in 4.1, but does in 4.2).

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

4.16.5.1 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>
```

4.16.5.2 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>
```

4.16.5.3 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 `xdmp: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 xdmp: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 xdmp: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>
```

4.16.6 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.

4.16.7 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 *Loading Content Into MarkLogic Server Guide*. For details on failover, see *Scalability, Availability, and Failover Guide*.

4.16.8 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.

4.16.9 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*.

4.16.10 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.

It is critical for swap space to be properly configured on your system according to the recommendations for your platform, as described in Memory, Disk Space, and Swap Space Requirements in the *Installation Guide*.

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 developer web site at http://developer.marklogic.com/:

Documentation	Description
Release Notes	Contains a summary of new features, upgrade compatible information, and known issues.
Installation Guide	Provides procedures for installing MarkLogic Server.
Administrator's Guide	Provides procedures for administrative tasks such as creating servers, creating databases, backing up databases, creating users, setting up your security policy, and so on.
Scripting Administrative Tasks Guide	Provides information on writing code to script various administrative tasks such as creating and modifying databases, App Servers, and so on.
Database Replication Guide	Provides information on database replication, useful for disaster recovery scenarios.
Flexible Replication Guide	Provides information on Flexible Replication, useful for information sharing from one database to another.

Documentation	Description
Application Developer's Guide	Provides procedures, methodologies, and conceptual information about general MarkLogic Server application development tasks.
Search Developer's Guide	Provides procedures, methodologies, and conceptual information about search-based application development tasks.
Application Builder Developer's Guide	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.
Information Studio Developer's Guide	Provides step-by-step information on using Information Studio to load and transform content.
REST Application Developer's Guide	Provides information on administration of and application development using the MarkLogic REST API.
Java Application Developer's Guide	Provides procedures, methodologies, and conceptual information about developing MarkLogic Server applications using the Java API.
Monitoring MarkLogic Guide	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.
MarkLogic Connector for Hadoop Developer's Guide	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.
Content Processing Framework Guide	Provides an introduction to the Content Processing Framework and procedures for installing the default content processing framework.
Understanding and Using Security Guide	Provides information on the role-based security model in MarkLogic Server.
Query Performance and Tuning Guide	Provides performance-related information that is useful to application developers and administrators.
Scalability, Availability, and Failover Guide	Provides information on large-scale system architecture, clustering, availability, and details on setting up shared-disk and local-disk failover.

Documentation	Description
Query Console User's Guide	Provides step-by step information on using Query Console, a tool to create and run arbitrary XQuery code.
SQL Data Modeling Guide	Provides information on how to use MarkLogic's SQL interface, including the creation of relational schemas and views.
Loading Content Into MarkLogic Server Guide	Provides procedures, methodologies, and conceptual information about loading content into MarkLogic Server. Includes an overview of ingestion techniques available using XQuery, Java, REST, .NET, and the MarkLogic Content Pump (mlcp).
MarkLogic XQuery and XSLT Function Reference	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.
MarkLogic REST API Reference	API documentation for the REST API.
MarkLogic 6 Release Notes	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.
XCC Developer's Guide	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.
XQuery and XSLT Reference Guide	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.
Getting Started with MarkLogic Server	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 32. 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 Admin Interface, Application Builder, the applications generated with Application Builder, and the other GUI tools (Query Console, Information Studio, and so on) are certified on Internet Explorer 8 and 9 on Windows, Firefox 34 on Windows and Mac OS, and Chrome 42 on Windows and Mac OS. Other browser/platform combinations may work but are not as thoroughly tested.

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 Express and Academic 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.

Combined Product Notices

MarkLogic Server 6.0 and supporting products.
Last updated: August 19, 2013

COPYRIGHT

Copyright © 2013 MarkLogic Corporation. All rights reserved.
This technology is protected by U.S. Patent No. 7,127,469B2, U.S. Patent No. 7,171,404B2, U.S. Patent No. 7,756,858 B2, and U.S. Patent No 7,962,474 B2.

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 Iceni 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 OFMERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ANDNONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BELIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTIONOF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTIONWITH 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. 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.

23. 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/.

24. 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.

25. 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.

26. jemalloc 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) 2002-2013 Jason Evans, and portions copyright (C) 2009-2013 Facebook, Inc. and copyright (C) 2007-2012 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.

27. 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.

28. Some features of MarkLogic allow you to develop and use Google Maps. Review Google Map's privacy policy at http://www.google.com/policies/privacy/, legal notices at http://maps.google.com/help/legalnotices_maps.html, and terms of use at http://maps.google.com/help/terms_maps.html.

29. commons-codec, HttpComponents, commons-logging, and Google Maps Utility are used under the Licensed under the Apache License, Version 2.0 (the "License"). You may not use these files 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.

30. Mapstraction. Copyright (c) 2012, Mapstraction. 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. 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.

31. logback-classic and logback-core are provided in object code form only under the Eclipse Public License - v 1.0 available at http://www.eclipse.org/legal/epl-v10.html. Copyright (C) 1999-2012, QOS.ch. All rights reserved. Source code for the Program is available http://logback.qos.ch/.

32. Simple Logging Facade for Java (SLF4J). 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.

33. 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.

34. heatmap.js distributed subject to the "THE BEER-WARE LICENSE" (Revision 42): <phk@FreeBSD.ORG> wrote this file. As long as you retain this notice you can do whatever you want with this stuff. If we meet some day, and you think this stuff is worth it, you can buy me a beer in return Poul-Henning Kamp.

35. Except for components that are otherwise subject to other open source licenses, the MarkLogic ODBC Driver is released under the GNU Lesser General Public License as published by the Free Software Foundation, version 2.1 of the License.

36. Solr-commons-csv-3.5.0 Copyright 2005-2006 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, this file is 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.

37. Avro v. 1.7.1. Copyright © 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, this file is 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.

38. Guava v. 11.0.2. Copyright © Google. 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, this file is 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.

39. Hadoop-auth v. 2.0.0. Copyright © 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, this file is 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.

40. hadoop-hdfs v. 2.0. Copyright © 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, this file is 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.

41. commons-lang v. 2.4, 2.5. Copyright © 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, this file is 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.

42. commons-configuration v. 2.0. Copyright © 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, this file is 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.

43. jackson-core-asl v. 1.8.8. Copyright © FasterXML. 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, this file is 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.

44. jackson-mapper-asl v. 1.8.8. Copyright © FasterXML. 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, this file is 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.

45. log4j v. 1.2.16. Copyright © 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, this file is 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.

46. protobuf-java v. 2.4.0a. Copyright © Google. 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 Google 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.

47. xpp3 v. 1.1.3.3. Copyright © 2002 Extreme! Lab, Indiana University. 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. The end-user documentation included with the redistribution, if any, must include the following acknowledgment: "This product includes software developed by the Indiana University Extreme! Lab(http://www.extreme.indiana.edu/)." Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear. The names "Indiana Univeristy" and "Indiana Univeristy Extreme! Lab" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact http://www.extreme.indiana.edu/. Products derived from this software may not use "Indiana Univeristy" name nor may "Indiana Univeristy" appear in their name, without prior written permission of the Indiana University. THIS SOFTWARE IS PROVIDED "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 AUTHORS, COPYRIGHT HOLDERS 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 IFADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

48. xstream v. 1.4.2. Copyright © 2003-2006, Joe Walnes. Copyright © 2006-2009, 2011 XStream Committers. 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 XStream 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.

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.