

---

# MarkLogic Server

---

## Release Notes

MarkLogic 9  
May, 2017

Last Revised: 9.0-13, July, 2020



---



---

## Table of Contents

---

### Release Notes

1.0	Introduction .....	13
1.1	Bug Fixes .....	13
2.0	Installation and Upgrade .....	15
2.1	Supported Platforms .....	15
2.2	Supported Filesystems .....	15
2.3	Upgrade Support .....	15
3.0	New Features in MarkLogic 9 .....	17
3.1	Template Driven Extraction (TDE) .....	18
3.2	SQL Enhancements .....	18
3.3	Optic API .....	18
3.4	Enhanced Tiered Storage .....	19
3.5	Encryption at Rest .....	19
3.6	Element Level Security .....	19
3.7	Redaction .....	20
3.8	Geospatial Enhancements .....	20
3.9	Entity Services API .....	20
3.10	New Stemming and Tokenization .....	21
3.11	Accurate Wildcard Expansion Options .....	21
3.12	mlcp Enhancements .....	22
	3.12.1 Support for SSL Connections .....	22
	3.12.2 Greater Control over Host Connections .....	22
	3.12.3 Redaction .....	23
	3.12.4 Batch Support for Server-Side Import Transformations .....	23
	3.12.5 Ability to Access and Modify Metadata in a Transformation .....	23
3.13	Java Client API Enhancements .....	23
	3.13.1 Bulk Asynchronous Data Movement .....	24
	3.13.2 Enhanced Temporal Document Support .....	24
	3.13.3 Security and Authentication Improvements .....	24
	3.13.4 Values Metadata Support .....	24
	3.13.5 Row-Based Search .....	24
	3.13.6 Geospatial Search Enhancements .....	25
3.14	Node.js Client API Enhancements .....	25
	3.14.1 Authentication and Connection Security .....	25
	3.14.2 Enhanced Temporal Document Support .....	25
	3.14.3 Values Metadata Support .....	25
	3.14.4 Geospatial Search Enhancements .....	26

3.14.5	Minimum Distance on Near Queries .....	26
3.15	Rest Client API Enhancements .....	26
3.15.1	Enhanced Temporal Document Support .....	26
3.15.2	Row-Based Data Evaluation .....	26
3.15.3	Point-in-Time Operations .....	27
3.15.4	cts:query Support .....	27
3.15.5	Values Metadata Support .....	27
3.16	Telemetry .....	27
3.17	XQuery 3.x Features .....	27
3.18	Query Console Enhancements .....	28
3.19	Application Display Environment Customization .....	28
3.20	Rolling Upgrades .....	29
3.21	Bi-temporal Compliance Enhancements .....	29
3.22	Secure Credentials .....	29
3.23	Certificate Authentication .....	29
3.24	New REST Management APIs .....	29
3.25	IPv6 Support .....	31
3.26	New Features in MarkLogic 9.0-2 .....	31
3.26.1	Geospatial Enhancements .....	31
3.26.1.1	Extended Units Support in Region Indexes .....	31
3.26.1.2	Crosses, Equals, and Touches Region Comparison Operators .....	31
3.26.1.3	New Functions for Probing Region Relationships .....	32
3.26.1.4	Support for Specifying Tolerance .....	32
3.26.2	Additional REST Client API Support for cts:query .....	33
3.26.3	REST Client API Supports QBE in a Combined Query .....	33
3.26.4	New Server-Side Transaction Controls .....	33
3.26.5	XCC: New Session Methods for Transaction Control .....	34
3.27	New Features in MarkLogic 9.0-3 .....	35
3.27.1	REST Management API: Advancing LSQT .....	35
3.27.2	REST Client API: Advancing LSQT .....	35
3.27.3	Ops Director .....	36
3.27.4	Amazon Web Services (AWS) 1-Click .....	36
3.27.5	Entity Services Enhancements .....	36
3.27.5.1	JSON Envelope Document Support .....	36
3.27.5.2	Support for Element Range Indexes .....	37
3.27.5.3	Entity Instance Namespace Support .....	37
3.27.5.4	Generated Code Refactoring .....	37
3.27.6	Search Result Sorting Enhancements .....	37
3.27.7	Redaction Built-in for Masking Numbers .....	37
3.27.8	Client APIs: Performance Improvements for JavaScript Extensions and Transforms .....	38
3.27.9	Java Client API: Values and Tuples Query Support for cts:query .....	38
3.28	New Features in MarkLogic 9.0-4 .....	38
3.28.1	Redaction: Support for Salting of Deterministic Masking Values .....	38
3.28.2	Redaction: New redact-datetime Built-In Function .....	39
3.28.3	Separate Download and Installation of Converters and Filters .....	39

3.28.4	Node.js Client API Enhancements .....	39
3.28.5	Client API Support for JavaScript Patch Content Constructors .....	40
3.28.6	Database Restore Enhancement .....	40
3.28.7	1-Click AWS Support .....	40
3.28.8	Restricted XPath Changes .....	40
3.28.9	Protected Path Sets Added to Element Level Security .....	41
3.29	New Features in MarkLogic 9.0-5 .....	41
3.29.1	MarkLogic Data Hub Framework .....	41
3.29.2	Entity Enrichment and Extraction Enhancements .....	42
3.29.3	Query Console: Explore Database by URI Pattern .....	42
3.29.4	Configuration Management API (CMA) XQuery and JavaScript Libraries 43	
3.29.5	Configuration Management API (CMA) REST Endpoints .....	43
3.29.6	Ops Director Enhancements .....	44
3.29.7	Monitoring History Enhancements .....	44
3.29.8	Tolerance Support in Geospatial Region Queries .....	44
3.29.9	Database Replication Enhancements .....	45
3.29.10	Fast Fail-over .....	45
3.30	New Features in MarkLogic 9.0-6 .....	45
3.30.1	Switching from Internal to External KMS .....	45
3.30.2	Multiple KMS Hosts for Failover .....	46
3.30.3	Geospatial Region Query Tolerance Improvements .....	46
3.30.4	Change in Multipart/Form-data Processing .....	46
3.31	New Features in MarkLogic 9.0-7 .....	46
3.31.1	Database Access with Granular Privileges .....	46
3.31.2	Forest Access with Granular Privileges .....	47
3.31.3	Query Console: Content Editing .....	47
3.31.4	Request Monitoring .....	47
3.31.5	MarkLogic Services in the Cloud .....	48
3.31.6	Java Client API: Load Balancer Improvements .....	48
3.32	New Features in MarkLogic 9.0-8 .....	48
3.32.1	Request Monitoring Enhancements .....	49
3.32.2	New Meters Available in Query Meters and Request Logs .....	49
3.32.3	Synchronize Keystore .....	50
3.32.4	New Client Issuer Authority Functions .....	50
3.32.5	Kerberos Authorization and LDAP Authentication Change .....	50
3.32.6	Logout Capability Added .....	50
3.33	New Features in MarkLogic 9.0-9 .....	51
3.33.1	Kerberos Authorization Parameter Change .....	51
3.33.2	Task Server Monitoring at Endpoint Level .....	51
3.33.3	MarkLogic SQL ORDER-BY Keyword .....	51
3.33.4	Change to ldapSearch Functionality .....	51
3.33.5	Amazon Linux 2 Support .....	51
3.33.6	Encryption Support for Thales HSM .....	51
3.33.7	MarkLogic Supports ECDH Key Exchange for SSL/TLS .....	52
3.34	New Features in MarkLogic 9.0-10 .....	52

3.34.1	OpenSSL Upgrade to Version 1.0.2s .....	52
3.34.2	Updates to the BiText Dictionaries .....	52
3.35	New Features in MarkLogic 9.0-11 .....	53
3.35.1	Request Monitoring Enhancements .....	53
3.35.2	Using MarkLogic Encryption with Microsoft Azure Key Vault .....	53
3.36	New Features in MarkLogic 9.0-12 .....	53
3.36.1	Rolling Upgrade Status Added .....	54
3.36.2	Other Changes .....	54
3.37	New Features in MarkLogic 9.0-13 .....	54
3.37.1	Packages by Linux Platform Updated .....	54
3.37.2	IAM Permissions Updated .....	54
4.0	Known Incompatibilities with Previous Releases .....	55
4.1	Incompatibilities (External) .....	56
4.1.1	The libnsl.so.1 library is required for CentOS 8 .....	56
4.2	Hadoop Connector Removed .....	56
4.3	HDFS Support Deprecated .....	56
4.4	Ops Director Deprecated .....	57
4.5	JavaScript: ValueIterator Replaced By Sequence .....	57
4.6	Database Stemming is Off, Word Searches On By Default .....	57
4.7	Collection Lexicon and Triple Index Enabled by Default .....	57
4.8	XCC .NET API No Longer Available .....	57
4.9	Changes in Semantic Query Behavior .....	58
4.9.1	Triple Index and SPARQL Engine Changes .....	58
4.9.2	Forest IDs Removed From sem:sparql Function .....	58
4.10	Triple Count Increased After Inserting Same Data Twice .....	58
4.11	Database max merge size Now Defaults to 48 GB .....	59
4.12	Changes to Range Index Reference Resolution .....	59
4.13	Default Stemming and Tokenization Libraries Changed for Most Languages .....	59
4.14	SQL DESCRIBE No Longer Supported by xdm:sql .....	61
4.15	Application-Specific Logging .....	61
4.16	Change to Classification of Some Special Symbol Tokens .....	61
4.17	Change to xdm:user-last-login .....	61
4.18	Changed Interfaces for xdm:document-insert and xdm:document-load .....	62
4.19	search:parse Returns a Different Type for cts:query Output Format .....	62
4.20	Default Client API Search Behavior Change on Port 8000 .....	62
4.21	JSON Property Scope and Container Queries Match Array Items Differently .....	62
4.22	REST Client API Incompatibilities .....	65
4.22.1	keyvalue Service Removed .....	65
4.22.2	Collections in Request Parameters are OR Related .....	65
4.22.3	Change in Default rest-reader and rest-writer Permissions .....	66
4.22.4	Default value of Document Management “repair” parameter changed .....	67
4.23	Java Client API Incompatibilities .....	67
4.23.1	Java Client API: Removal of Deprecated Interfaces .....	68
4.23.2	Java Client API: JAR File Name and Maven Artifact ID Change .....	70
4.23.3	Logging Turned Off by Default .....	71

4.24	Node.js Client API Incompatibilities .....	71
4.24.1	Changes to Return Value of documents.remove .....	71
4.24.2	Transaction Creation Returns an Object by Default .....	71
4.24.3	Default Search Result Slice is Zero-Based .....	71
4.25	Geospatial Region Accessors Can Now Return Double Values .....	72
4.26	User-Defined Function Plugins Must Be Recompiled .....	72
4.27	SLES 12 No Longer Supported .....	73
4.28	Solaris No Longer Supported .....	73
4.29	Nagios Plugin No Longer Supported .....	73
4.30	Application Builder and Information Studio No Longer Available .....	73
4.31	Admin Interface No Longer Selects a Default Schemas Database .....	74
4.32	Internal Security ON with External Security Object Behavior Change .....	74
4.33	REST Management API Changes in MarkLogic 9 .....	75
4.34	Configuration Packaging Format Incompatibilities .....	76
4.35	Java Client API 4.1.1 Incompatibilities .....	76
4.35.1	Load Balancer Configuration for DMDSDK Jobs .....	76
4.36	Incompatibilities Between 9.0-10 and 9.0-11 .....	76
4.36.1	The API <code>trgr:triggers-change-modules-database()</code> Only Has Two Parameters 76	
4.36.2	Entity Services Features Deprecated .....	77
4.37	Incompatibilities Between 9.0-8 and 9.0-9 .....	77
4.37.1	MarkLogic SQL ORDER-BY Keyword .....	77
4.38	Incompatibilities Between 9.0-5 and 9.0-6 .....	77
4.38.1	Changes to Accepted XML Character Set .....	78
4.39	Incompatibilities Between 9.0-4 and 9.0-5 .....	78
4.39.1	Configuration Manager Deprecated .....	79
4.39.2	Minimum Required Version of HDP is 2.6 .....	79
4.39.3	Reindex Recommended for Geospatial Region Indexes .....	79
4.39.4	Geospatial Region Query Results Might Differ .....	79
4.39.5	<code>return-query</code> Option Output Format Change .....	79
4.40	Incompatibilities Between 9.0-3 and 9.0-4 .....	80
4.40.1	Redaction: Deterministic Masking Values Differ .....	80
4.41	Incompatibilities Between 9.0-2 and 9.0-3 .....	80
4.41.1	Changes to Authentication Behavior with Client Certificate .....	80
4.41.2	XCC ContentSource.newSession Interface Change .....	81
4.41.3	Document Digest Authorization Behavior Changed in 9.0-3 .....	81
4.41.4	1-click AMIs, new compatible CloudFormation, and additional upgrade procedures 81	
4.41.5	<code>map:new</code> Retains Keys with Empty Values .....	81
4.42	Incompatibilities Between 9.0-1 and 9.0-2 .....	82
4.42.1	The <code>mlcp</code> Option <code>-tolerate_errors</code> is Ignored .....	82
4.42.2	Changes to <code>jsearch.facets</code> Output Structure .....	82
4.42.3	Array Type is Preserved in x509 Certificate with Array-Valued Properties 82	
4.42.4	Node.js Client API: <code>valuesBuilder.slice</code> is Now Zero-Based .....	83
4.42.5	Changes to <code>xdmp:update</code> XQuery Prolog Option .....	83

4.42.6	Java Client API 4.0.2 Ignores HttpClientConfigurator .....	84
4.43	MarkLogic 8 Incompatibilities .....	84
4.43.1	JSON Related Incompatibilities .....	85
4.43.1.1	Documents Created as JSON With MarkLogic 7 REST API or mlcp Must Be Converted to Native JSON 85	
4.43.1.2	json:unquotedString Primitive No Longer Available .....	87
4.43.1.3	xdmp:to-json and json:transform-to-json Now Returns a document-node() 87	
4.43.1.4	Search, Java, REST: json-key Is Now json-property in Options and Structured Query 87	
4.43.1.5	Java and REST: Specifying a Language for JSON Documents is Deprecated 88	
4.43.1.6	Java and REST: Default Path Language for JSON Document Patches is Now XPath 88	
4.43.1.7	Java and REST: New Restrictions on Patching JSON Content 89	
4.43.1.8	Java and REST: Transforms and Extensions That Manipulate JSON Must Be Rewritten 89	
4.43.1.9	Java and REST: JSON Array Items and Property Values No Longer Distinguishable in QBE 90	
4.43.1.10	Field Range Query and Field Value Query on JSON May Behave Differently 90	
4.43.2	Semantics Incompatibilities .....	91
4.43.2.1	Changed Function: sem:sparql .....	91
4.43.2.2	Changed Function: sem:sparql-values .....	92
4.43.2.3	Changed Function: sem:sparql-values .....	92
4.43.2.4	Deprecated Function: sem:sparql-triples .....	92
4.43.2.5	Changed Behavior: Graphs .....	92
4.43.3	REST and Java Client API Incompatibilities .....	92
4.43.3.1	Must Upgrade to Java Client API v3.0 .....	93
4.43.3.2	REST API Instance Must Use the Declarative Rewriter on the App Server 93	
4.43.3.3	Default Transaction Mode for the POST Method of Resource Service Extensions is Now Query 93	
4.43.3.4	REST API: Empty Bulk Read by Query Now Returns 200 Status 94	
4.43.3.5	Error Reporting Format and Detail Changes .....	94
4.43.3.6	Deprecated Interface: Keyvalue Queries .....	95
4.43.3.7	Transaction ID Format Has Changed .....	96
4.43.3.8	A Transaction Can No Longer Be Shared Across Users .....	96
4.43.3.9	Java: QBE Search Results No Longer Automatically Match the Query Format 96	
4.43.4	Document Library Services (DLS) Repositories Need To Perform A Bulk Upgrade Operation 97	
4.43.5	Linux Now Requires Red Hat 6 .....	99
4.43.6	mysql On Linux No Longer Ships With Server .....	99
4.43.7	Cyrillic Tokenization Changes .....	99



4.43.8	Application Builder Applications Must Be Re-Deployed in MarkLogic 8 ..	100
4.43.9	Application Builder and Information Studio Links Removed .....	100
4.43.10	Search API Incompatibilities .....	100
4.43.10.1	search:parse Output is Now Unannotated cts:query XML .....	101
4.43.10.2	Deprecated Option: extract-metadata .....	101
4.43.10.3	Deprecated Functions: search:unparse, search:remove-constraint	101
4.43.10.4	Structured Query: locks-query and properties-query Renamed ....	101
4.43.10.5	sort-order Query Option Requires an Index .....	102
4.43.11	Locks and Properties Query Built-In Functions Renamed .....	102
4.43.12	xdmp:uri-content-type Of an XML Document Now Returns application/ xml, Can Affect CPF Applications	102
4.43.13	xdmp:function Signature Change .....	103
4.43.14	Incompatibilities Between 8.0-5 and 8.0-6 .....	103
4.43.14.1	Terms Matched by additional-query Are Highlighted in Snippets	103
4.43.15	Incompatibilities Between 8.0-3 and 8.0-4 .....	103
4.43.15.1	xdmp.multipartDecode Now Returns a JSON Payload for Headers	104
4.43.15.2	In JavaScript, Some Thesaurus and Spelling Function Have Dif- ferent Return Type	104
4.43.15.3	xdmp.databaseRestoreStatus Now Returns an Object .....	104
4.43.15.4	Serialization Error Code Changes .....	105
4.43.15.5	Change to Required Java Version .....	105
4.43.15.6	Deprecated mlcp Command Line Options .....	105
4.43.15.7	REST APIs That Have JSON or XML Payloads Cannot Have Empty Payloads	105
4.43.15.8	Geospatial Namespace and Data Version Changes .....	106
4.43.16	Incompatibilities Between 8.0-2 and 8.0-3 .....	108
4.43.16.1	spell.suggestDetailed, xdmp.filesystemDirectory, and xdmp.en- codingLanguageDetect Now Return ValueIterator	109
4.43.16.2	xdmp.databaseRestoreStatus Now Returns an Array .....	109
4.43.16.3	xdmp.gssServerNegotiate Now Returns a JavaScript Object ..	109
4.43.16.4	Use of String Transaction Ids in Node.js To Be Deprecated ..	109
4.43.16.5	CDH 4.3 is No Longer a Supported Hadoop Distribution .....	109
4.43.16.6	Changes to How MarkLogic Locates Java and Hadoop Libraries for HDFS Forest Storage	110
4.43.17	Incompatibilities Between 8.0-1 and 8.0-2 .....	110
4.43.17.1	Array Input Differences in fn.distinctValues, fn.subsequence, and Other Functions	110
4.43.17.2	The Second Parameters of xdmp.eval, xdmp.invoke, xdmp.xque- ryEval, and xdmp.spawn Now Take a Single Object	111
4.43.17.3	extract-document-data Results Now Inline By Default .....	111
4.43.17.4	XCC v8.0-2 May Require Config Change When Used with Older	

	Versions of MarkLogic	111
	4.43.17.5 Client APIs: JavaScript Extension and Transform Error Reporting Convention Change	112
	4.43.17.6 Some JavaScript Built-In Functions that Returned XML Structures Now Return JSON Structures	112
4.44	MarkLogic 7 Incompatibilities	113
4.44.1	Incompatibilities Between MarkLogic 7.0-3 and 7.0-2	113
4.44.1.1	HDP No Longer a Supported Hadoop Platform	113
4.44.1.2	Java API: ContentVersionRequest Property Deprecated	113
4.44.1.3	REST API: content-versions Property Deprecated	114
4.44.1.4	REST API: JSON documents Cannot be Retrieved as XML	114
4.44.2	Float Precision Greater in 7.0-3	115
4.44.3	Incompatibilities Between MarkLogic 7.0-2 and 7.0-1	115
4.44.3.1	Change to JSON Output from the REST API	115
4.44.3.2	Changes to the MarkLogic Connector for Hadoop API	115
4.44.4	Incompatibilities Between MarkLogic 7.0-1 and MarkLogic 6	115
4.44.4.1	XQuery HTTP Client Built-In Functions Now Require a Privilege	116
4.44.4.2	HTTP Client Functions Are Now HTTP 1.1 Compliant	116
4.44.4.3	xdmp:get-request-username and xdmp:get-request-user Changes	117
4.44.4.4	Specifying a Forest Now Only Works With Strict Locking	117
4.44.4.5	Custom Dictionaries for Japanese and Chinese Languages Need to be Re-saved	117
4.44.4.6	Default Attributes on XML Copy Changes	118
4.44.4.7	Serialization of Alerting, Reverse, and Path Range Queries Change	119
4.44.4.8	Java and REST Client API Incompatibilities	119
4.44.4.9	Namespace Change for Properties Persisted Using JSON	121
4.44.4.10	mlcp Incompatibilities	127
4.44.4.11	REST Management API Version Incremented to v2	128
4.44.4.12	Changes to the Configuration Manager	131
4.44.4.13	xdmp:plan Now Requires a Privilege	131
4.44.4.14	fn:analyze-string Now Returns Output in a Different Namespace	131
5.0	Planning for Future Upgrades	133
5.1	Hadoop Connector Deprecated	133
5.2	HDFS Support Deprecated	133
5.3	Packaging API Deprecated	133
5.4	info and infodev APIs Deprecated	134
5.5	Annotated Query Output from search:parse Deprecated	134
5.6	Search API Grammar Customization Deprecated	135
5.7	Search API searchable-expression Deprecated	135
5.8	The mlcp Option -tolerate_errors Deprecated	135
5.9	xdmp:transaction-mode XQuery Prolog Option Deprecated	136

5.10	Deprecation of transaction-mode Option to xdm:eval .....	137
5.11	XCC Session.setTransactionMode is Deprecated .....	137
5.12	Java Client API 4.0.2 Deprecations .....	138
5.13	Java Client API 4.0.4 Deprecations .....	138
	5.13.1 NamespacesManager Interface Deprecated .....	139
	5.13.2 QueryBatcher.getQuerySuccessListeners Deprecated .....	139
5.14	REST Client API Namespace Configuration Deprecation .....	139
5.15	Configuration Packaging XQuery Library Deprecated .....	139
5.16	Configuration Manager Deprecated .....	140
6.0	Other Notes .....	141
6.1	Memory and Disk Space Requirements .....	141
6.2	Compatibility with XQuery Specifications .....	142
6.3	XQuery Extensions .....	142
6.4	SQL Queries .....	142
6.5	Documentation .....	143
6.6	Browser Requirements .....	146
6.7	Security: Prevent Abuse of System Entity Expansion .....	147
7.0	Technical Support .....	149
8.0	Copyright .....	151



## 1.0 Introduction

MarkLogic 9 is a major release of MarkLogic Server that includes many new features. The new features are described in “New Features in MarkLogic 9” on page 17. The following lists some of the major features with links to where they are described:

- [Template Driven Extraction \(TDE\)](#)
- [SQL Enhancements](#)
- [Optic API](#)
- [Enhanced Tiered Storage](#)
- [Encryption at Rest](#)
- [Element Level Security](#)
- [Redaction](#)
- [Geospatial Enhancements](#)
- [Entity Services API](#)
- [Bulk Asynchronous Data Movement](#)

For a description of these and many more new features, see “New Features in MarkLogic 9” on page 17.

If you are upgrading from MarkLogic 8, some applications will require minor changes to run correctly on MarkLogic 9. For details, see “Known Incompatibilities with Previous Releases” on page 55.

### 1.1 Bug Fixes

To review the list of bug fixes included between any two version of MarkLogic Server, click on the “Fixed Bugs” link (<https://help.marklogic.com/Bugtrack/List>) that is available on the MarkLogic Support Portal (<https://help.marklogic.com/>), and enter the version range you are interested in.



## 2.0 Installation and Upgrade

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

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

### 2.1 Supported Platforms

For a complete list of supported platforms, see [Supported Platforms](#) in the *Installation Guide*.

### 2.2 Supported Filesystems

For a complete list of supported filesystems, see [Supported Filesystems](#) in the *Installation Guide*.

### 2.3 Upgrade Support

This section describes upgrade support to MarkLogic 9. For details on installing MarkLogic Server and for the upgrade procedure, see the *Installation Guide*.

**Warning** MarkLogic Early Access does not support upgrade. This section describes upgrade for 9.0-1 and later.

Upgrading is supported from 7.0-6 or later. If you are running a release prior to 7.0, you must first upgrade to MarkLogic 7 or MarkLogic 8 before upgrading to MarkLogic 9. If you are upgrading a cluster, you must first upgrade the node in which the Security database forest is located before you upgrade other nodes in the cluster.

**Note:** MarkLogic Corporation strongly recommends performing a backup of your databases before upgrading to MarkLogic 9. Additionally, MarkLogic Corporation recommends that you first upgrade to the latest maintenance release of MarkLogic 7 or MarkLogic 8 before upgrading to MarkLogic 9.

An upgrade from MarkLogic 7 or MarkLogic 8 does not require a reindex. If you are upgrading from a previous release that does require a reindex and you choose not to reindex your databases, the database will run in 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 9 features (as well as earlier features depending upon which compatibility mode it runs) and may treat all content in the database as English language content. For details on database compatibility, see the *Installation Guide*.

**Note:** If you are upgrading clusters with DB replication configured, see [Upgrading Clusters Configured with Database Replication](#) in the *Database Replication Guide* for details.

MarkLogic 7 and later includes a new rebalancing feature with a more efficient document placement algorithm. Upon upgrade, databases from previous MarkLogic releases are set to use the `legacy` document assignment policy, which is the same as used in previous MarkLogic releases. If you do plan on reindexing an upgraded database, MarkLogic recommends that you consider setting your databases to use the new `bucket` document assignment policy. The `bucket` policy is more efficient for rebalancing your database across forests if you add or remove forests from your configuration. For more details, see [Database Rebalancing](#) in the *Administrator's Guide*.

There are some known incompatibilities between MarkLogic 8 and MarkLogic 9. You might need to make some minor code changes to your MarkLogic 8 applications before they can run correctly in MarkLogic 9. For details on the incompatibilities, see “Known Incompatibilities with Previous Releases” on page 55. For instructions on upgrading to MarkLogic 9, including information about database compatibility between MarkLogic 8 and MarkLogic 9, see the *Installation Guide*.

If you are upgrading to MarkLogic 9.0-4 or later, you may have to install MarkLogic Converters package separately. For more details, see [MarkLogic Converters Installation Changes Starting at Release 9.0-4](#) in the *Installation Guide*.

If you upgrade to MarkLogic 9.0-5 or later from an earlier version of MarkLogic 9 and you use a geospatial region index, you should reindex. For more details, see “Tolerance Support in Geospatial Region Queries” on page 44.



## 3.0 New Features in MarkLogic 9

This chapter describes the new features in MarkLogic 9.

- [Template Driven Extraction \(TDE\)](#)
- [SQL Enhancements](#)
- [Optic API](#)
- [Enhanced Tiered Storage](#)
- [Encryption at Rest](#)
- [Element Level Security](#)
- [Redaction](#)
- [Geospatial Enhancements](#)
- [Entity Services API](#)
- [New Stemming and Tokenization](#)
- [Accurate Wildcard Expansion Options](#)
- [mlcp Enhancements](#)
- [Java Client API Enhancements](#)
- [Node.js Client API Enhancements](#)
- [Rest Client API Enhancements](#)
- [Telemetry](#)
- [XQuery 3.x Features](#)
- [Query Console Enhancements](#)
- [Application Display Environment Customization](#)
- [Rolling Upgrades](#)
- [Bi-temporal Compliance Enhancements](#)
- [Secure Credentials](#)
- [Certificate Authentication](#)
- [New REST Management APIs](#)
- [IPv6 Support](#)
- [New Features in MarkLogic 9.0-2](#)
- [New Features in MarkLogic 9.0-3](#)
- [New Features in MarkLogic 9.0-4](#)

- [New Features in MarkLogic 9.0-5](#)
- [New Features in MarkLogic 9.0-6](#)
- [New Features in MarkLogic 9.0-7](#)
- [New Features in MarkLogic 9.0-8](#)
- [New Features in MarkLogic 9.0-9](#)
- [New Features in MarkLogic 9.0-10](#)
- [New Features in MarkLogic 9.0-11](#)
- [New Features in MarkLogic 9.0-12](#)
- [New Features in MarkLogic 9.0-13](#)

### 3.1 Template Driven Extraction (TDE)

MarkLogic 9 enables you to define a relational lens over your document data, so you can query parts of your data using SQL or the new Optic API. Templates let you specify which parts of documents make up rows in a view. You can also use templates to define a semantic lens, specifying which values from a document make up triples in the triple index.

For more details, see [Template Driven Extraction \(TDE\)](#) in the *Application Developer's Guide*.

### 3.2 SQL Enhancements

MarkLogic is a NoSQL database, where the unit of storage and indexing is a document. The document model makes it possible to express rich, related, varying structures — anything from a scientific journal article to a complex derivative trade. Many users want to view parts of these rich structures as though they were simple tables — to see the data in those documents through a relational lens. While it is possible to create and query SQL views with MarkLogic 8, MarkLogic 9 further enhances the SQL capabilities with a number of new features, including Templates, the new Optic API, and an updated ODBC driver.

In MarkLogic 9, you can define a template that specifies which parts of the document make up a row in a view, and then query that view from a server-side program with `xdmp:sql()` or `xdmp.sql()` or via ODBC. You can also query that view server-side from the new MarkLogic Optic API — a fluent JavaScript interface with the ability to perform joins and aggregates on views over documents.

### 3.3 Optic API

The Optic API blends the relational world with rich NoSQL document features by providing the capability to perform joins and aggregates over documents. One of the enabling features, Template Driven Extraction, makes it possible to create a relational lens over documents stored in MarkLogic by using templates to specify the parts of a document that make up a row in a view. You can access that data using the Optic API in XQuery, JavaScript, or Java.

For more details, see [Optic API for Multi-Model Data Access](#) in the *Application Developer's Guide*.

### 3.4 Enhanced Tiered Storage

MarkLogic 9 adds more flexibility to the Tiered Storage feature. First, it allows documents to be matched to a tier based on a query, in addition to the current solution that is based on a range index pair of boundary values. That means that any complex query can be used to create a policy to match data or metadata about documents in order to decide in which tier to save or move the document. Second, it allows tiered storage queries to match dates using age, for instance, you can write a query for high-end tier that matches documents created in the last 30 days, and another for mid-end tier that matches documents created between 1 and 3 years, rather than the fixed pair of boundary values in the current solution. Third, it removes the need of using super databases to create tiers, simplifying the life of developers, which were required to understand the tiering policy in order to write high performance queries or any update.

MarkLogic 9 adds more performance to the Tiered Storage feature. Now queries that use elements which are part of the tiering policy are optimized to run against the nodes that are more likely to have that data. That means that if you have a query that uses the same element as the tiering policy, for instance an element range query on create date > 30 days, the query engine will direct the search only to the data nodes that store that data, in this example high-end tier nodes.

### 3.5 Encryption at Rest

MarkLogic 9 introduces the ability to encrypt “data at rest” - data that is on media (on disk or in the cloud), as opposed to data that is being used in a process. Encryption can be applied to newly created files, configuration files, or log files. Existing data files can be encrypted by triggering a merge or re-index of the data.

For more information about using Encryption at Rest, see [Encryption at Rest](#) in the *MarkLogic Security Guide*.

**Note:** If a node in your cluster is offline for any reason, wait until the host comes back online to make any changes to your encryption at rest settings. Do not change your encryption settings while a host is offline.

### 3.6 Element Level Security

Element Level Security is an addition to the MarkLogic security model that allows you to specify more complex security rules for specific elements within documents. Element Level Security can be applied to either JSON or XML documents. Users without the appropriate permissions cannot view the secured element or JSON property.

Element Level Security can be used in addition to the existing document level security and compartment security. For details about using Element Level Security, see [Element Level Security](#) in the *Security Guide*.

### 3.7 Redaction

MarkLogic 9 introduces a new library module and `mlcp` command line option that enable you to redact sensitive data when extracting documents from the database. Redaction enables you to obscure or hide portions of a document using a rule-based read transformation.

Redaction is available through the following interfaces:

- The `mlcp -redaction` command line option.
- The `rdt:redact` XQuery function.
- The `rdt:redact` Server-Side JavaScript function.

For more details, see [Redacting Document Content](#) in the *Application Developer's Guide* and [Redacting Content During Export or Copy Operations](#) in the *mlcp User Guide*.

### 3.8 Geospatial Enhancements

MarkLogic 9 includes the following enhancements to geospatial operations:

- Geospatial region search, including a new geospatial region index type
- ETRS-89 coordinate system support
- Double precision coordinate support
- Support for kilometers, meters, and feet as units of measure

You can take advantage of these enhancements using all of the MarkLogic geospatial and search related interfaces, including `cts:search`, `search:search`, JSearch, and the REST, Java, and Node.js Client APIs.

**Note:** The performance of geospatial region queries for geographic coordinate systems (WGS84 and ETRS89) has not yet been fully optimized. Performance will be improved in a future release.

For details, see [Geospatial Search Applications](#) in the *Search Developer's Guide*.

### 3.9 Entity Services API

The Entity Services API enables you to quickly and easily model your business entities and relationships between them and then generate code and configuration artifacts that provide a framework for an entity based application.

The artifacts you generate with the Entity Services API make it easier to create entities from raw source data, query entities and the relationships between them, and manage your entities and models.

For details, see the *Entity Services Developer's Guide*.

### 3.10 New Stemming and Tokenization

In MarkLogic 9 the default tokenization and stemming code has been changed for all languages (except English tokenization). Some tokenization and stemming behavior will change between MarkLogic 8 and MarkLogic 9. We expect that in most cases results will be better in MarkLogic 9.

You can now customize stemming and tokenization for a language by selecting one of several built-in stemmer and lexer plug-ins, or by creating your own stemmer and lexer plug-ins in C++.

You can also configure a custom stemming and/or tokenization dictionary for any language. Previously, you could not distinguish between stemming and tokenization dictionaries, and you could not install a dictionary for an unsupported language.

For more details, see the following topics:

- [Tokenization and Stemming](#) in the *Search Developer's Guide*
- [User-Defined Lexer Plugins](#) in the *Search Developer's Guide*
- [Using a User-Defined Stemmer Plugin](#) in the *Search Developer's Guide*
- [Custom Dictionaries for Tokenizing and Stemming](#) in the *Search Developer's Guide*
- “Default Stemming and Tokenization Libraries Changed for Most Languages” on page 59

### 3.11 Accurate Wildcard Expansion Options

MarkLogic 9 adds three new cts:query constructor options for controlling wildcard expansion.

When evaluating a wildcarded search term, MarkLogic must sometimes make a tradeoff between speed and accuracy. MarkLogic attempts to expand a wildcard term from the lexicon. To prevent this expansion from taking too long, there is a limit on how many words MarkLogic will extract from the lexicon. When the limit is reached, MarkLogic falls back on alternative strategies that accurate for all possible terms containing wildcard characters.

For many searches, these inaccuracies are an acceptable tradeoff for the fast response of an unfiltered search. Where the inaccuracies are not acceptable, you can use the following new options to change the default behavior: `-lexicon-expansion-limit=N`, `-limit-check`, `-no-limit-check`.

For more details, see the function reference documentation for the following functions:

XQuery	Server-Side JavaScript
<code>cts:word-query</code>	<code>cts.wordQuery</code>
<code>cts:element-word-query</code>	<code>cts.elementWordQuery</code>
<code>cts:element-attribute-word-query</code>	<code>cts.elementAttributeWordQuery</code>
<code>cts:field-word-query</code>	<code>cts.fieldWordQuery</code>
<code>cts:json-property-word-query</code>	<code>cts.jsonPropertyWordQuery</code>
<code>cts:element-value-query</code>	<code>cts.elementValueQuery</code>
<code>cts:element-attribute-value-query</code>	<code>cts.elementAttributeValueQuery</code>
<code>cts:field-value-query</code>	<code>cts.fieldValueQuery</code>
<code>cts:json-property-value-query</code>	<code>cts.jsonPropertyValueQuery</code>

### 3.12 mlcp Enhancements

The following capabilities have been added to the mlcp command line tool:

- [Support for SSL Connections](#)
- [Greater Control over Host Connections](#)
- [Redaction](#)
- [Batch Support for Server-Side Import Transformations](#)
- [Ability to Access and Modify Metadata in a Transformation](#)

In addition the mlcp source code is now available for download through the GitHub `marklogic-contentpump` project. For details, see [Accessing the mlcp Source Code](#) in the *mlcp User Guide*.

#### 3.12.1 Support for SSL Connections

The mlcp command line tool can now connect to MarkLogic via an SSL (Secure Socket Layer) connections. For details, see [Connecting to MarkLogic Using SSL](#) in the *mlcp User Guide*.

#### 3.12.2 Greater Control over Host Connections

You can now specify multiple hosts for mlcp to connect to during `import`, `export`, and `copy` jobs. Used by itself, this feature enables mlcp to fall back an alternative host if the initial host is not available.

You can also use this capability in conjunction with the new `-restrict_hosts` option to prevent `mlcp` from connecting to any hosts except the ones on the initial host list.

For more details, see [Controlling How `mlcp` Connects to MarkLogic](#) in the *mlcp User Guide*.

### 3.12.3 Redaction

MarkLogic 9 introduces the ability to redact sensitive data when extracting documents from the database with the export or copy commands by specifying redaction rule collections in the `-redact` option.

For more details, see [Redacting Content During Export or Copy Operations](#) in the *mlcp User Guide* and [Redacting Document Content](#) in the *Application Developer's Guide*.

### 3.12.4 Batch Support for Server-Side Import Transformations

Previously, the batch size was always one when applying a server-side transformation during an `mlcp import` or `copy` job. This restriction has been lifted in MarkLogic 9. All documents in a batch are now transformed and inserted into the database as a single statement, greatly improving performance when using a transformation.

### 3.12.5 Ability to Access and Modify Metadata in a Transformation

Collections, permissions, document quality, and temporal collection specified by the client are now available to a transformation function via the context parameter. In addition, a transformation function can set collections, permissions, quality, temporal collection, and values metadata for its output document(s).

For details, see [Transforming Content During Ingestion](#) in the *mlcp User Guide*.

## 3.13 Java Client API Enhancements

The following capabilities have been added to the Java Client API:

- [Bulk Asynchronous Data Movement](#)
- [Enhanced Temporal Document Support](#)
- [Security and Authentication Improvements](#)
- [Values Metadata Support](#)
- [Row-Based Search](#)
- [Geospatial Search Enhancements](#)

### 3.13.1 Bulk Asynchronous Data Movement

The new Data Movement SDK feature of the Java Client API enables you to move large amounts of data into, out of, or within a MarkLogic cluster asynchronously. These interfaces leverage your entire cluster for scale-out performance. The feature supports streaming, asynchronous, long running data movement tasks.

For details, see the `com.marklogic.client.datamovement` package in the *Java Client API Documentation* and [Asynchronous Multi-Document Operations](#) in the *Java Application Developer's Guide*.

### 3.13.2 Enhanced Temporal Document Support

You can now pass a temporal document URI when creating, updating, and patching temporal documents. This is logical document URI in a temporal collection. For more details, see the methods of `com.marklogic.client.bitemporal.TemporalDocumentManager` that accept a `temporalDocumentURI` parameter and [Working with Temporal Documents](#) in the *Developing Applications With the Java Client API*.

If you have sufficient privileges, you can now wipe (completely remove) a temporal document using `TemporalDocumentManager.wipe`.

You can use the Java Client API to protect a temporal document against update, deletion, or wipe for a specified time. See `TemporalDocumentManager.protect`.

### 3.13.3 Security and Authentication Improvements

You can now use Kerberos or certificate-based authentication to authenticate with MarkLogic. You can also connect to MarkLogic via a SSL (Secure Socket Layer) connection.

For more details, see [Authentication and Connection Security](#) in the *Developing Applications With the Java Client API*.

### 3.13.4 Values Metadata Support

MarkLogic 9 adds the ability to associate key-value metadata with a document. You can use the Java Client API to add, update, delete, and search values metadata. For more details, see [Values Metadata](#) in the *Developing Applications With the Java Client API*.

### 3.13.5 Row-Based Search

You can use the Java Client API to execute a plan produced by the Optic API and receive row-based results in your Java application. For details, see [Optic Java API for Relational Operations](#) in the *Developing Applications With the Java Client API*.



### 3.13.6 Geospatial Search Enhancements

MarkLogic 9 includes enhancements to geospatial search such as region searches, double precision coordinates, additional coordinate systems. These features are exposed through the Java Client API. For more details, see “Geospatial Enhancements” on page 20, [Creating Region Queries Using the Client APIs](#) in the *Search Developer’s Guide*, and `StructuredQueryBuilder.geospatial` in the *Java Client API Documentation*.

### 3.14 Node.js Client API Enhancements

The following capabilities have been added to the Node.js Client API:

- [Authentication and Connection Security](#)
- [Enhanced Temporal Document Support](#)
- [Values Metadata Support](#)
- [Geospatial Search Enhancements](#)
- [Minimum Distance on Near Queries](#)

#### 3.14.1 Authentication and Connection Security

You can now use Kerberos or certificate-based authentication to authenticate with MarkLogic. You can also connect to MarkLogic via a SSL (Secure Socket Layer) connection.

For more details, see [Authentication and Connection Security](#) in the *Node.js Application Developer’s Guide*.

#### 3.14.2 Enhanced Temporal Document Support

You can now pass a temporal document URI when creating, updating, and patching temporal documents. This is the logical document URI in a temporal collection. For more details, see the methods in the `documents` namespace that accept a `temporalDocument` parameter and [Working with Temporal Documents](#) in the *Node.js Application Developer’s Guide*.

If you have sufficient privileges, you can now wipe (completely remove) a temporal document using `TemporalDocumentManager.wipe`.

#### 3.14.3 Values Metadata Support

MarkLogic 9 adds the ability to associate key-value metadata with a document. You can use the Node.js Client API to add, update, delete, and search values metadata. For more details, see the `metadataValues` metadata category in [Working with Metadata](#) in the *Node.js Application Developer’s Guide*.

### 3.14.4 Geospatial Search Enhancements

MarkLogic 9 includes enhancements to geospatial search such as region searches, double precision coordinates, additional coordinate systems. These features are exposed through the Node.js Client API. For more details, see “Geospatial Enhancements” on page 20, [Creating Region Queries Using the Client APIs](#) in the *Search Developer’s Guide*, and `queryBuilder.geospatialRegion` in the *Node.js Client API Reference*.

### 3.14.5 Minimum Distance on Near Queries

The `queryBuilder.near` method now accepts a minimum distance for near queries. Previously, you could only specify a maximum distance. For details, see the *Node.js Client API Reference* and `cts:near-query`.

## 3.15 Rest Client API Enhancements

The following capabilities have been added to the REST Client API:

- [Enhanced Temporal Document Support](#)
- [Row-Based Data Evaluation](#)
- [Point-in-Time Operations](#)
- [Values Metadata Support](#)
- [cts:query Support](#)

### 3.15.1 Enhanced Temporal Document Support

You can now pass a temporal document URI when creating, updating, and patching temporal documents. This is logical document URI in a temporal collection. For more details see [Working with Temporal Documents](#) in the *REST Application Developer’s Guide* and the `/v1/documents` methods that accept a `temporal-document` parameter.

If you have sufficient privileges, you can now wipe (completely remove) a temporal document using `DELETE /v1/documents?result=wiped`. For more details, see `DELETE:/v1/documents` in the *MarkLogic REST API Reference*.

The new `POST:/v1/documents/protection` method enables you to protect a temporal document against update, deletion, or wipe for a specified time period. For more details, see the *MarkLogic REST API Reference*.

### 3.15.2 Row-Based Data Evaluation

The new `/rows` service enables you to execute a plan produced by the Optic API and receive results in a variety of row-based formats. For more details, see:

- `GET:/v1/rows` in the *MarkLogic REST API Reference*
- `POST:/v1/rows` in the *MarkLogic REST API Reference*

- [Optic API for Multi-Model Data Access](#) in the *Application Developer's Guide*

### 3.15.3 Point-in-Time Operations

Most read and search operations now accept a timestamp request parameter that enables you to make successive requests that will be evaluated against the state of the database at a point in time. The timestamp can be obtained from the ML-Effective-Timestamp header returned by the same methods.

For more details, see [Performing Point-in-Time Operations](#) in the *REST Application Developer's Guide*.

### 3.15.4 cts:query Support

You can use a serialized cts:query in place of a structured query in `GET:/v1/search` and `POST:/v1/search`. For more details, see [Searching With cts:query](#) in the *REST Application Developer's Guide* and the *MarkLogic REST API Reference*.

### 3.15.5 Values Metadata Support

MarkLogic 9 adds the ability to associate key-value metadata with a document. You can use the REST Client API to add, update, delete, and search values metadata. You can use the new metadata category `metadata-values` when working with this type of metadata.

For more details, see [Working with Metadata](#) in the *REST Application Developer's Guide*.

## 3.16 Telemetry

Telemetry is part of our continuous effort to provide better and faster support by automating the data collection process required on most support tickets. When Telemetry is enabled, it collects, encrypts, and sends diagnostic and anonymized system-level information about a MarkLogic cluster to a secure MarkLogic destination.

The Telemetry feature collects only system level information, and sends it to a protected and secure location where it can only be accessed by the MarkLogic technical teams, to be used to facilitate troubleshooting and monitor performance. Telemetry does not collect any personally identifiable information, user data or application logs. See [Telemetry](#) in the *Monitoring MarkLogic Guide* for more information.

## 3.17 XQuery 3.x Features

MarkLogic now supports selected features from the XQuery 3.0 and XQuery 3.1 specifications. These features are only available when using the “1.0-ml” XQuery dialect.

For more details, see [XQuery 3.x Features](#) in the *XQuery and XSLT Reference Guide*.

### 3.18 Query Console Enhancements

MarkLogic 9 includes the following new features and enhancements in the Query Console application. For more details, see the *Query Console User Guide*.

- Support for profiling Server-Side JavaScript. The profiling output differs from the XQuery output in being sampling based.
- Auto-complete suggestions. As you type into a query, Query Console displays a list of functions, keywords, and in-scope variables suggests. For functions, you also see the function reference documentation.
- The editor now provides auto-closure of parentheses, quotes, braces, and other grouping symbols.
- Each query has now has its own tab for displaying results. The results of one query are not lost if you switch to another query and run it.
- Query results description differentiates between a result that is a single item versus a sequence of one item so that you can tell whether or not the query returns a sequence.
- Query execution time is displayed along with the results. For example, Query Console will report “Returned sequence of 2 items in 230ms”. The display also includes the time delta between the current and previous run.
- You can now select a content source and App Server independent of one another.
- MarkLogic 9 enhances the database Explorer to enable you to view metadata such as collections, permissions, key-value metadata, and document quality, along with the contents of a document. This feature is available when you select a document from the Explorer document list.
- You can reorder queries in a workspace and reorder tabs in the editor using drag-and-drop.
- Query results larger than 5M will be truncated.

### 3.19 Application Display Environment Customization

Administrators can now customize the display environment of MarkLogic applications such as Query Console and the Monitoring Dashboard in the following ways:

- Display a notification dialog when a user navigates to one of the MarkLogic application pages.
- Display a customized banner across the top of each page of the MarkLogic applications. This enables you to easily differentiate between environments such as production and staging.,

For more details, see [Configuring a MarkLogic Application Message and Banner](#) in the *Administrator's Guide*.

### 3.20 Rolling Upgrades

A rolling upgrade is one way to address the need for highly available clusters under heavy transaction loads to upgrade to a newer version of MarkLogic in a seamless manner. Hosts in a cluster are upgraded one by one, without incurring any downtime in availability or interruption of transactions.

For more details, see [Rolling Upgrades](#) in the *Administrator's Guide*.

### 3.21 Bi-temporal Compliance Enhancements

The MarkLogic bi-temporal data management feature has been enhanced to provide the option to store valid and system axes and archival information outside of temporal documents in metadata, rather than directly in the documents. Storing the axes times in metadata enables MarkLogic to update the axes timestamps without changing the documents and invoking reindexing.

The ability to update nodes in temporal documents has been added.

Temporal documents can be protected from certain temporal operations, such as update, delete or wipe for a specified period of time and then automatically archived to a WORM (Write Once Read Many) device.

MarkLogic also supports uni-temporal documents that have only a system axes.

### 3.22 Secure Credentials

Secure credentials enable a security administrator to manage credentials and to be made available to less privileged users for authentication to other systems without giving them access to the credentials themselves.

### 3.23 Certificate Authentication

Certificate-based authentication requires internal and external users and HTTPS clients to authenticate themselves to MarkLogic via a client certificate, either in addition to, or rather than a password.

### 3.24 New REST Management APIs

In MarkLogic 9, a number of REST Management APIs have been added. A few of the APIs have been updated as well. This table summarizes the changes.

Added	Updated
	GET:/manage/v2/credentials/properties
	PUT:/manage/v2/credentials/properties
GET:/manage/v2?view=describe	
GET:/manage/v2/credentials/secure	

Added	Updated
POST: /manage/v2/credentials/secure	
GET: /manage/v2/databases/{id name}/partition-queries	
POST: /manage/v2/databases/{id name}/partition-queries	
DELETE: /manage/v2/databases/{id name}/partition-queries/{partition-number}	
GET: /manage/v2/databases/{id name}/partition-queries/{partition-number}	
GET: /manage/v2/databases/{id name}/partition-queries/{partition-number}/properties	
PUT: /manage/v2/databases/{id name}/partition-queries/{partition-number}/properties	
PUT: /manage/v2/databases/{id name}/partition-queries/{partition-number}/properties	
POST: /manage/v2/hosts	
GET: /manage/v2/meters	
GET: /manage/v2/protected-paths	
POST: /manage/v2/protected-paths	
GET: /manage/v2/protected-paths/{id}/properties	
PUT: /manage/v2/protected-paths/{id}/properties	
DELETE: /manage/v2/protected-paths/{id name}	
GET: /manage/v2/protected-paths/{id name}	
GET: /manage/v2/query-rolesets	
POST: /manage/v2/query-rolesets	
DELETE: /manage/v2/query-rolesets/{id name}	
GET: /manage/v2/query-rolesets/{id name}	
GET: /manage/v2/query-rolesets/{id name}/properties	
PUT: /manage/v2/query-rolesets/{id name}/properties	
GET: /manage/v2/security	
POST: /manage/v2/security	
GET: /manage/v2/security/properties	
PUT: /manage/v2/security/properties	

### 3.25 IPv6 Support

MarkLogic version 9 supports IPv6 networking, including client connectivity to MarkLogic endpoints. All MarkLogic features and capabilities work seamlessly with IPv6 in addition to working with the IPv4 protocol.

The IPv6 protocol provides virtually limitless address space for use by government, universities, and private sector companies in accessing Internet services.

### 3.26 New Features in MarkLogic 9.0-2

MarkLogic version 9.0-2 contains the following new features:

- [Geospatial Enhancements](#)
- [Additional REST Client API Support for cts:query](#)
- [REST Client API Supports QBE in a Combined Query](#)
- [New Server-Side Transaction Controls](#)
- [XCC: New Session Methods for Transaction Control](#)

#### 3.26.1 Geospatial Enhancements

MarkLogic 9.0-2 contains the following enhancements related to geospatial region support:

- [Extended Units Support in Region Indexes](#)
- [Crosses, Equals, and Touches Region Comparison Operators](#)
- [New Functions for Probing Region Relationships](#)
- [Support for Specifying Tolerance](#)

##### 3.26.1.1 Extended Units Support in Region Indexes

You can specify additional units (feet, kilometers, meters) when creating a geospatial region index. Previously, you could only specify miles. The units are only meaningful when using a geodetic coordinate system such as WGS84.

For more details, see `admin:database-geospatial-region-path-index` in the *MarkLogic XQuery and XSLT Function Reference*.

##### 3.26.1.2 Crosses, Equals, and Touches Region Comparison Operators

You can now use the “crosses”, “equals” and “touches” operators for geospatial region queries and other region computation functions that previously accepted region comparison operators.

For example, you can use “crosses” as an operator name to functions such as `cts:geospatial-region-query` (XQuery) or `cts.geospatialRegionQuery` (JavaScript). You can also use `DE9IM_CROSSES` in query text passed to `cts:parse` (XQuery) or `cts.parse` (JavaScript).

### 3.26.1.3 New Functions for Probing Region Relationships

The following new functions are available for comparing geospatial region values:

- `geo:region-relate` (XQuery) and `geo.regionRelate` (JavaScript) - test whether two regions satisfy a relationship such as “crosses” or “overlaps”
- `geo:region-de9im` (XQuery) and `geo.regionDe9im` (JavaScript) - compute the DE9-IM intersection matrix between two regions.

For more details, see the function reference documentation in the *MarkLogic XQuery and XSLT Function Reference* or the *MarkLogic Server-Side JavaScript Function Reference*.

### 3.26.1.4 Support for Specifying Tolerance

Tolerance is the largest allowable variation in geometry comparisons. If the distance between two points is less than tolerance, then the two points are considered equal. For more details, see [Understanding Tolerance](#) in the *Search Developer’s Guide*.

The following functions now support a “tolerance” option. For more details, see the function reference documentation.

XQuery	Server-Side JavaScript
<code>geo:arc-intersection</code>	<code>geo.arcIntersection</code>
<code>geo:box-intersects</code>	<code>geo.boxIntersects</code>
<code>geo:circle-intersects</code>	<code>geo.circleIntersects</code>
<code>geo:polygon-intersects</code>	<code>geo.polygonIntersects</code>
<code>geo:complex-polygon-intersects</code>	<code>geo.complexPolygonIntersects</code>
<code>geo:region-intersects</code>	<code>geo.regionIntersects</code>
<code>geo:polygon-contains</code>	<code>geo.polygonContains</code>
<code>geo:complex-polygon-contains</code>	<code>geo.complexPolygonContains</code>
<code>geo:region-contains</code>	<code>geo.regionContains</code>
<code>geo:geohash-encode</code>	<code>geo.geohashEncode</code>



### 3.26.2 Additional REST Client API Support for cts:query

You can use a serialized cts:query in place of a structured query in the following methods:

- GET: /v1/search and POST: /v1/search
- GET: /v1/values/{name} and POST: /v1/values/{name}

On the GET methods, specify the cts:query as the value of the structuredQuery request parameter. On the POST methods, put the cts:query in the POST body.

In addition, you can include a cts:query serialized as XML or JSON in place of a structured query in a combined query. This is applicable to any method that accepts a combined query input, such as a POST request to /v1/search or /v1/values/{name}.

For more details, see [Searching With cts:query](#) in the *REST Application Developer's Guide* and the *MarkLogic REST API Reference*.

### 3.26.3 REST Client API Supports QBE in a Combined Query

You can now include a Query By Example (QBE) in a combined query when using the REST Client API. For more details, see [Specifying Dynamic Query Options with Combined Query](#) in the *REST Application Developer's Guide*.

### 3.26.4 New Server-Side Transaction Controls

You can now configure the commit mode (auto or explicit) and transaction type (query, update, or auto) independently when configuring a new transaction. This change manifests in the following ways:

- The `xdmp:update` XQuery prolog option accepts a new value, “auto”, which specifies that MarkLogic should determine the transaction/statement type (query or update) through static analysis. The pre-existing value “false” now means the transaction/statement type is query. Use this option plus `xdmp:commit` instead of the now-deprecated `xdmp:transaction-mode` prolog option.
- A new XQuery prolog option, `xdmp:commit`, has been added for specifying whether a main module should run as a single-statement, auto-commit transaction (the default) or a multi-statement explicit-commit transaction. Use this option plus `xdmp:update` instead of the now-deprecated `xdmp:transaction-mode` prolog option.
- New `commit` and `update` options have been added to the functions listed in the table below. Use these in preference to the `transaction-mode` option, which has been deprecated.

The following functions support the new commit and update options. For more details, see the function reference documentation for `xdmp:eval` (XQuery) or `xdmp.eval` (JavaScript).

XQuery	JavaScript
<code>xdmp:eval</code>	<code>xdmp.eval</code>
<code>xdmp:javascript-eval</code>	<code>xdmp.xqueryEval</code>
<code>xdmp:invoke</code>	<code>xdmp.invoke</code>
<code>xdmp:invoke-function</code>	<code>xdmp.invokeFunction</code>
<code>xdmp:spawn</code>	<code>xdmp.spawn</code>
<code>xdmp:spawn-function</code>	

For more details on the new capabilities, see [Understanding Transactions in MarkLogic Server](#) in the *Application Developer's Guide*.

For details on transitioning from the old transaction controls to the new ones, see the following topics:

- “Changes to `xdmp:update` XQuery Prolog Option” on page 83
- “`xdmp:transaction-mode` XQuery Prolog Option Deprecated” on page 136
- “Deprecation of transaction-mode Option to `xdmp:eval`” on page 137.

### 3.26.5 XCC: New Session Methods for Transaction Control

The following methods have been added to the `Session` class for configuring transactions and querying transaction configuration:

- `Session.setAutoCommit` and `Session.isAutoCommit`
- `Session.setUpdate` and `Session.getUpdate`

You should use these methods rather than `Session.setTransactionMode`, which has been deprecated. For details, see “XCC `Session.setTransactionMode` is Deprecated” on page 137.

`Session.setAutoCommit` controls whether requests submitted during the session run in a transaction with auto-commit semantics (the default) or explicit commit semantics. Executing a request with commit set to explicit starts a multi-statement transaction.

`Session.setUpdate` controls whether requests submitted during the session run in a query transaction, an update transaction, or if the transaction type should be automatically detected by MarkLogic through analysis of the submitted code. Auto detection is the default behavior.

Note that if you override the `Session` transaction configuration in an ad hoc query, the behavior differs depending on whether you configure the session using `setTransactionMode` or `setAutoCommit` and `setUpdate`. With `setAutoCommit` and `setUpdate`, the transaction configuration reverts to the `Session` settings once the transaction involving the override completes. With `setTransactionMode`, the override persists and affects future transactions unless you explicitly change it.

### 3.27 New Features in MarkLogic 9.0-3

MarkLogic version 9.0-3 contains the following new features:

- [REST Management API: Advancing LSQT](#)
- [REST Client API: Advancing LSQT](#)
- [Ops Director](#)
- [Amazon Web Services \(AWS\) 1-Click](#)
- [Entity Services Enhancements](#)
- [Search Result Sorting Enhancements](#)
- [Redaction Built-in for Masking Numbers](#)
- [Client APIs: Performance Improvements for JavaScript Extensions and Transforms](#)
- [Java Client API: Values and Tuples Query Support for `cts:query`](#)

#### 3.27.1 REST Management API: Advancing LSQT

You can use the following new method of the REST Management API to advance LSQT on a temporal collection:

```
POST /manage/v2/databases/{id|name}/temporal/collections?collection=collname
```

For more details, see

`POST: /manage/v2/databases/{id|name}/temporal/collections?collection={name}` in the *MarkLogic REST API Reference*.

#### 3.27.2 REST Client API: Advancing LSQT

You can now use the following new method of the REST Client API to advance LSQT on a temporal collection:

```
POST /v1/temporal/collections/{name}
```

For more details, see `POST: /v1/temporal/collections/{name}` in the *MarkLogic REST API Reference*.

### 3.27.3 Ops Director

Ops Director presents a consolidated view of your MarkLogic infrastructure via dashboards that streamline monitoring and troubleshooting of clusters with alerting, performance, and log data. For details, see the *Ops Director Guide*.

### 3.27.4 Amazon Web Services (AWS) 1-Click

MarkLogic is now available for deployment on Amazon Web Services (AWS) by means of AWS 1-Click. For details, see <http://developer.marklogic.com/products/cloud/aws>.

### 3.27.5 Entity Services Enhancements

The Entity Services API offers the following additional capabilities as of MarkLogic 9.0-3:

- [JSON Envelope Document Support](#)
- [Support for Element Range Indexes](#)
- [Entity Instance Namespace Support](#)
- [Generated Code Refactoring](#)

#### 3.27.5.1 JSON Envelope Document Support

You can now work with either XML or JSON envelope documents. Previously, only XML was supported.

As part of this change, some of the functions generated in instance converter modules have been extended to accept a format parameter. If you omit the format, these functions generate XML instances and envelopes, as they did prior to MarkLogic 9.0-3.

The following generated function interfaces now accept an optional format parameter, for some entity type *T*.

- `T:instance-to-canonical`
- `T:instance-to-envelope`

If you want to use JSON envelopes, you should regenerate your instance converter and version translator code.

For more details, see [Entity Instance Concepts](#) in the *Entity Services Developer's Guide*.

### 3.27.5.2 Support for Element Range Indexes

An entity type can now include specifications for entity properties to be backed by an element range index. Previously, you could only specify a path range index.

In an XML model descriptor, use the `path-range-index` and `element-range-index` elements to specify entity properties to be backed by an index. The `path-range-index` element is identical to the pre-existing `range-index` element.

In a JSON model descriptor, use the `pathRangeIndex` and `elementRangeIndex` properties to specify entity properties to be backed by an index. The `pathRangeIndex` property is identical to the pre-existing `rangeIndex` property.

For more details, see [Identifying Entity Properties for Indexing](#) in the *Entity Services Developer's Guide*.

### 3.27.5.3 Entity Instance Namespace Support

An entity type definition can now include a namespace URI and namespace prefix for XML instances of that type. The namespace binding is used when generating entity instances in XML, and when generating index configuration and query option artifacts.

For more details, see [Defining a Namespace URI for an Entity Type](#) in the *Entity Services Developer's Guide*.

### 3.27.5.4 Generated Code Refactoring

The instance converter and version translator code you can generate using Entity Services has been refactored to improve readability and ease of customization, and to accommodate both XML and JSON envelopes. Previously generated code will continue to work as-is.

For more details, see [Generating Code and Other Artifacts](#) in the *Entity Services Developer's Guide*.

### 3.27.6 Search Result Sorting Enhancements

The Search API and REST, Java, and Node.js Client APIs now support additional search result ordering choices through the `sort-order` query option.

Previously, you could only specify score ordering. You can now sort based on score, fitness, quality, and more. For details, see [sort-order](#) in the *Search Developer's Guide*.

### 3.27.7 Redaction Built-in for Masking Numbers

A new built-in redaction function, `redact-number`, is available as of MarkLogic 9.0-3. This function provides fine-grained control over the type, range, and format of masking values for numeric data. You can use `redact-number` with mlcp's redaction capability, the `rdt:redact` XQuery function, or the `rdt.redact` Server-Side JavaScript function.

For details, see `redact-number` in the *Application Developer's Guide*.

### 3.27.8 Client APIs: Performance Improvements for JavaScript Extensions and Transforms

Performance improvements have for Client API server-side transformations and extensions implemented in Server-Side JavaScript. This affects applications using transformations or extensions with the REST, Java, or Node.js Client APIs.

**Note:** Your applications will only benefit from the improved performance if you reinstall your Server-Side JavaScript transforms and extensions.

### 3.27.9 Java Client API: Values and Tuples Query Support for cts:query

As of Java Client API version 4.0.3, you can use a serialized cts:query as an additional constraint on a values or tuples query.

## 3.28 New Features in MarkLogic 9.0-4

MarkLogic 9.0-4 introduces the following new features:

- [Redaction: Support for Salting of Deterministic Masking Values](#)
- [Redaction: New redact-datetime Built-In Function](#)
- [Separate Download and Installation of Converters and Filters](#)
- [Node.js Client API Enhancements](#)
- [Client API Support for JavaScript Patch Content Constructors](#)
- [Database Restore Enhancement](#)
- [1-Click AWS Support](#)
- [Restricted XPath Changes](#)
- [Protected Path Sets Added to Element Level Security](#)

### 3.28.1 Redaction: Support for Salting of Deterministic Masking Values

As of MarkLogic 9.0-4, the `mask-deterministic` built-in redaction function supports two new options for specifying a salt value for masking value generation. Salting can provide a higher level of security. For more details, see the discussion of the `salt` and `extend-salt` options of [mask-deterministic](#) in the *Application Developer's Guide*.

Note that the introduction of these options changes the default behavior of deterministic masking value generation. For details, see “Redaction: Deterministic Masking Values Differ” on page 80.

### 3.28.2 Redaction: New redact-datetime Built-In Function

As of MarkLogic 9.0-4, you can use the `redact-datetime` built-in redaction function in your redaction rules. You can mask a `dateTime` value with a random `dateTime` value or using a picture string.

For more details, see `redact-datetime` in the *Application Developer's Guide*.

### 3.28.3 Separate Download and Installation of Converters and Filters

As of MarkLogic release 9.0-4, MarkLogic converters/filters are offered as a separate package (called *MarkLogic Converters* package) from MarkLogic Server package.

This change provides better flexibility and enables you to install/uninstall MarkLogic converters/filters separately from MarkLogic Server.

For more details, see [MarkLogic Converters Installation Changes Starting at Release 9.0-4](#) in the *Installation Guide*.

### 3.28.4 Node.js Client API Enhancements

The following new features are available as of Node.js Client API v2.1.1. Unless otherwise noted, MarkLogic 9 is required to take advantage of these changes.

- Support for Optic API queries. The Optic API enables you to query documents, triples, and lexicons to extract row oriented data from MarkLogic. For more details, see [Using the Optic API for Relational Operations](#) in the *Node.js Application Developer's Guide*.
- Point-in-time query support. Many read-only operations, such as document read and query, now support a timestamp parameter or option to enable successive requests to share a consistent view of the database at a given timestamp. For details, see [Performing Point-in-Time Operations](#) in the *Node.js Application Developer's Guide*.
- Partial update (document patch) improvements. Prior to MarkLogic 9.0-4, you could not perform the following operations in a patch. These limitations have been removed. For more details, see [Patching Document Content or Metadata](#) in the *Node.js Application Developer's Guide*.
  - Insert a property immediately under the root node of a JSON object node using “last-child” position.
  - Address items in nested arrays.
- Support for advancing LSQT for temporal collections. See the function `documents.advanceLsqt` in the *Node.js Client API Reference*.

### 3.28.5 Client API Support for JavaScript Patch Content Constructors

The partial update or “patch” feature of the REST and Java Client APIs now support replacement content constructor functions implemented in Server-Side JavaScript. Versions of MarkLogic prior to 9.0-4 only support XQuery implementations.

For more details, see [Writing an XQuery User-Defined Replacement Constructor](#) in the *REST Application Developer’s Guide*.

### 3.28.6 Database Restore Enhancement

As of MarkLogic release 9.0-4, you can restore a database from a backup, even if the number of database forests are asymmetrical to the backup forests. As a result, the Admin Interface and related API database restore functions have been changed.

For details, see [Restoring a Reconfigured Database](#) in the *Administrator’s Guide*.

### 3.28.7 1-Click AWS Support

MarkLogic now supports the 1-click launch option in AWS Marketplace. Because of this, the published MarkLogic AMIs will have data volume predefined.

### 3.28.8 Restricted XPath Changes

The following features impose restrictions on XPath expressions used in their configuration. The existence of the restrictions is unchanged. However, as of MarkLogic 9.0-4, the restrictions have been more formally defined for each feature and somewhat relaxed for some features.

The following list summarizes how the restrictions have changed for affected features. In all cases, the changes enable a larger subset of XPath and do not introduce backward incompatibilities.

- Indexable path expressions: The path expressions usable for defining path-based indexes can include function calls to selected functions in predicates.
- TDE context path expressions: The path expressions usable for defining path-based indexes can include function calls to selected functions in predicates. You can use “/” as a context path.
- Element Level Security protected paths: The path expressions usable for defining path-based indexes can include function calls to selected functions in predicates.
- Extracting document data from search results (extract-document-data query option), document patch feature of the REST Client APIs, Optic API xpath constructor: The path expressions usable for defining path-based indexes can include function calls to selected functions in predicates. You can use an unnamed node test in the leaf step of a path expression.

For more details, see [Restricted XPath](#) in the *XQuery and XSLT Reference Guide*.



### 3.28.9 Protected Path Sets Added to Element Level Security

As part of MarkLogic 9.0-4, Element Level Security now includes the protected path set feature. A protected path set is a way to allow multiple protected paths to cover and secure the same element, with both AND and OR relationships between the permissions. The information (the name of the protected path set) is simply a “tag” on the protected path definition. This enables multiple arbitrary security markings for an element.

For more details, see [Protected Path Sets](#) in the *Security Guide*.

## 3.29 New Features in MarkLogic 9.0-5

MarkLogic 9.0-5 introduces the following new features:

- [MarkLogic Data Hub Framework](#)
- [Entity Enrichment and Extraction Enhancements](#)
- [Query Console: Explore Database by URI Pattern](#)
- [Configuration Management API \(CMA\) XQuery and JavaScript Libraries](#)
- [Configuration Management API \(CMA\) REST Endpoints](#)
- [Ops Director Enhancements](#)
- [Monitoring History Enhancements](#)
- [Tolerance Support in Geospatial Region Queries](#)
- [Database Replication Enhancements](#)
- [Fast Fail-over](#)

### 3.29.1 MarkLogic Data Hub Framework

MarkLogic 9.0-5 introduces full support for MarkLogic Data Hub Framework (DHF) version 3.0. DHF is an Open Source data integration framework and set of tools that enable you to quickly integrate data from multiple sources into a single MarkLogic database, and then expose that integrated data through MarkLogic.

DHF includes both client-side libraries and tools for integration and modeling, and server-side framework support. To learn more about DHF, see the following:

- DHF documentation and tutorial: <https://marklogic.github.io/marklogic-data-hub/>
- DHF Open Source project: <https://github.com/marklogic/marklogic-data-hub>

To use DHF, you will need DHF version 3.0 or later, MarkLogic 9.0-5 or later, and Oracle Java 8 JRE (client-side). For details, see the DHF documentation.

### 3.29.2 Entity Enrichment and Extraction Enhancements

MarkLogic 9.0-5 introduces new built-in and library functions for entity enrichment and extraction. The new interfaces use entity dictionaries to identify entities. You can create an entity dictionary in several ways, including from a graph created from a SKOS ontology. For more details, see [Entity Extraction and Enrichment](#) in the *Search Developer's Guide*.

Use the following new functions for entity enrichment and extraction:

XQuery	Server-Side JavaScript
<code>entity:enrich</code>	<code>entity.enrich</code>
<code>entity:extract</code>	<code>entity.extract</code>
<code>cts:entity-highlight</code>	<code>cts.entityHighlight</code>
<code>cts:entity-walk</code>	<code>cts.entityWalk</code>

Use the following new functions for creating and managing entity dictionaries:

XQuery	Server-Side JavaScript
<code>entity:dictionary-insert</code>	<code>entity.dictionaryInsert</code>
<code>entity:dictionary-load</code>	<code>entity.dictionaryLoad</code>
<code>entity:skos-dictionary</code>	<code>entity.skosDictionary</code>
<code>cts:entity-highlight</code>	<code>cts.entityHighlight</code>
<code>cts:entity-dictionary</code>	<code>cts.entityDictionary</code>
<code>cts:entity-dictionary-get</code>	<code>cts.entityDictionaryGet</code>
<code>cts:entity-dictionary-parse</code>	<code>cts.entityDictionaryParse</code>

### 3.29.3 Query Console: Explore Database by URI Pattern

You can now apply a URI filter to the documents listed in the Explorer view. You can specify a single URI (“/my/interesting/document.xml”) or a wildcard expression (“/my/interesting/\*.xml”). For more details, see [Filtering the Explorer View by URI](#) in the *Query Console User Guide*.

### 3.29.4 Configuration Management API (CMA) XQuery and JavaScript Libraries

MarkLogic 9.0-5 introduces new library functions for configuration management.

The configuration management functions can be used to:

- retrieve a configuration of an individual resource, a set of resources, or a full cluster;
- generate a configuration from scenarios, such as High Availability (HA) scenario;
- apply a named configuration, overriding parameters and setting options.

Use the following new functions for configuration management:

XQuery	Server-Side JavaScript
<code>cma:generate-config</code>	<code>cma.generateConfig</code>
<code>cma:apply-config</code>	<code>cma.applyConfig</code>

For more details, see *MarkLogic XQuery and XSLT Function Reference* and *MarkLogic Server-Side JavaScript Function Reference*.

### 3.29.5 Configuration Management API (CMA) REST Endpoints

MarkLogic 9.0-5 introduces a new REST API for configuration management: Configuration Management API (CMA).

The Configuration Management API is a RESTful API that allows retrieving, generating, and applying configurations for MarkLogic clusters, databases, and application servers.

Use the following new endpoints for configuration management:

Endpoint	Description
GET /manage/v3	This endpoint enables retrieving configuration of an individual resource, a set of resources, or a full cluster. It also enables generating new configurations from scenarios.
POST /manage/v3	This endpoint enables applying named configurations to MarkLogic resources, overriding parameters and setting options. The configurations may be applied to an individual resource, a set of resources, or a full cluster.

For more details, see *MarkLogic REST API Reference*.

### 3.29.6 Ops Director Enhancements

MarkLogic 9.0-5 includes the new version of the Ops Director application: Ops Director 1.1-1.

The Ops Director 1.1-1 has the following new features and enhancements:

- Enhanced security: default authentication changed from “application level” to “digest”.
- Improved support for data integration with other applications. In particular, you may export all tables from the Manage view, the Support view, and the Console Settings view. The tables are exported into CSV files, which might be later on imported into other applications (e.g. Excel) for further data processing and analysis. Also, you may export all metrics from charts in the Analyze view.
- Improved Role Based Access Control (RBAC): added “inherited” access to resources driven by hierarchical relationships in resource groups.
- Implemented functionality for upgrading Ops Director and updating its configuration.
- Added UX enhancements, such as preserving view state between navigations, and UI performance improvements, such as enhanced rendering speed of the pages.

For more details, see the *Ops Director Guide*.

### 3.29.7 Monitoring History Enhancements

The Monitoring History application in MarkLogic 9.0-5 was enhanced with new metrics for XDQP Server Requests performance. For details, see [XDQP Server Requests Performance Data](#) in the *Monitoring MarkLogic Guide*.

### 3.29.8 Tolerance Support in Geospatial Region Queries

As of MarkLogic 9.0-5, geospatial region queries support a “tolerance” option.

Tolerance is the largest allowable variation in geometry comparisons. If the distance between two points is less than tolerance, then the two points are considered equal. For more details, see [Understanding Tolerance](#) in the *Search Developer’s Guide*.

This change affects the functions `cts:geospatial-region-query` (XQuery), `cts.geospatialRegionQuery` (JavaScript), the `geo-option` portion of a Search API `geo-region-path` constraint option, the `geo-option` portion of a `geo-region-path-query` or `geo-region-constraint-query`, and the options you can specify in a region query using the Java Client API and Node.js Client API.

Due to the improved tolerance support, when upgrading to MarkLogic 9.0-5 or later from an earlier version of MarkLogic 9, you might not get accurate results for geospatial region queries until you reindex.

### 3.29.9 Database Replication Enhancements

MarkLogic 9.0-5 and later versions enable you to more easily enable/disable or suspend/resume database replication.

Disabling database replication changes the database configuration, so the disabled state persists across restarts and failovers. Use suspend/resume when lag is high and you need to stop database replication quickly.

Suspending database replication for a forest does not change the configuration, so replication resumes after a restart or failover. Use enable/disable when you want to stop replication for a long period, such as when you move a replica site.

This capability is exposed in the following ways:

- In the Admin UI, navigate to the Database Replication settings for a database, click on the Summary tab, and use the enable, disable, suspend, or resume buttons of the Summary tab. For more details, see [Enabling, Disabling, Suspending, and Resuming Database Replication](#) in the *Database Replication Guide*.
- Use the XQuery functions `xmdp:forest-database-replication-suspend` and `xmdp:forest-database-replication-resume` or the Server-Side JavaScript functions `xmdp.forestDatabaseReplicationSuspend` and `xmdp.forestDatabaseReplicationResume`.

### 3.29.10 Fast Fail-over

If one or more host forests are configured for local-disk or shared-disk failover, you now have the option to failover those forests when you shut down the host. A new option, “Immediately fail over forests to replica hosts,” has been added to the Host Shutdown confirmation page to enable you to fail over the forests to replica hosts.

## 3.30 New Features in MarkLogic 9.0-6

MarkLogic 9.0-6 introduces the following new features:

- [Switching from Internal to External KMS](#)
- [Multiple KMS Hosts for Failover](#)
- [Geospatial Region Query Tolerance Improvements](#)
- [Change in Multipart/Form-data Processing](#)

### 3.30.1 Switching from Internal to External KMS

In MarkLogic 9.0-6, when switching from an external KMS to an internal KMS, encryption at rest does not limit the files that can be decrypted to only those encrypted with keys that are in the internal KMS. MarkLogic can decrypt any file encrypted with the internal MarkLogic KMS or an external KMS.

It is no longer necessary to decrypt your data before transitioning files encrypted with an external KMS to the internal MarkLogic KMS. In MarkLogic 9.0-5, you needed to decrypt all your data before switching from an external KMS back to internal KMS. This is no longer necessary.

### 3.30.2 Multiple KMS Hosts for Failover

In MarkLogic 9.0-6, you can now specify multiple hosts, multiple ports, and multiple KMIP credentials to connect to KMIP servers. MarkLogic can store more than one set of external KMIP server credentials, to be used in encrypting and decrypting data. In the case where a KMIP server is unavailable and the first specified KMIP server stops responding, MarkLogic will try to connect to each of the other hosts in the user-specified list until it succeeds.

### 3.30.3 Geospatial Region Query Tolerance Improvements

MarkLogic 9.0-6 includes refinements to tolerance support in geospatial region queries. Some tolerance related edge cases now produce more accurate results. You can only benefit from this change if you reindex.

If your region queries do not involve polygons that fall into these edge cases, you will not see any change in geospatial region query results even if you do reindex.

### 3.30.4 Change in Multipart/Form-data Processing

Starting with MarkLogic 9.0-6, forms submitted with blank fields result in field entries with blank string content (for example: ""). Previously, these empty form fields would not result in any field entry.

## 3.31 New Features in MarkLogic 9.0-7

MarkLogic 9.0-7 introduces the following new features:

- [Database Access with Granular Privileges](#)
- [Forest Access with Granular Privileges](#)
- [Query Console: Content Editing](#)
- [Request Monitoring](#)
- [MarkLogic Services in the Cloud](#)
- [Java Client API: Load Balancer Improvements](#)

### 3.31.1 Database Access with Granular Privileges

In MarkLogic 9.0-7, you can now enable a user to create, reindex, update properties, run operations on, or delete a database if they have the `manage` role and one of the following granular privileges:

```
http://marklogic.com/xdmp/privileges/admin/database/database-ID
```

```
http://marklogic.com/xdmp/privileges/admin/database/activity/database-ID
```

where *activity* is one of the following:

- backup
- merge
- restore
- clear
- index

### 3.31.2 Forest Access with Granular Privileges

In MarkLogic 9.0-7, you can now enable a user to create, update properties, run operations on, or delete a forest if they have the `manage` role and one of the following granular privileges:

`http://marklogic.com/xdmp/privileges/admin/forest`

`http://marklogic.com/xdmp/privileges/admin/forest/forest-ID`

### 3.31.3 Query Console: Content Editing

The database Explorer in Query Console now enables you to do the following operations without writing any code:

- Insert a new document into a database.
- Modify an existing document in the database.
- Delete one or selected documents from a database.

For more details, see [Editing Database Content](#) in the *Query Console User Guide*.

### 3.31.4 Request Monitoring

The Request Monitoring feature enables you to configure logging of information related to requests, including metrics collected during request execution. This feature lets you enable logging of internal preset metrics for requests on specific endpoints. You can also log custom request data by calling the provided Request Logging APIs. This logged information may help you evaluate server performance.

For more details, see [Endpoints and Request Monitoring](#) in the *Query Performance and Tuning Guide*.

### 3.31.5 MarkLogic Services in the Cloud

In MarkLogic 9.0-7 you can now access both Query Service and Data Hub Service as services on AWS. Query Service enables you to respond to varying query workloads for your MarkLogic cluster. This elastic query service eliminates the need to set up and manage the underlying infrastructure required to scale for capacity. MarkLogic Query Service work by attaching MarkLogic e-nodes to to your existing MarkLogic cluster, and auto-scaling as your workload requires. For more information, see <https://cloudservices.marklogic.com/> and <https://www.marklogic.com/product/marklogic-database-overview/query-service/>.

MarkLogic Data Hub Service is a cloud service providing the features of the Data Hub Framework in a managed instance available on AWS. The Data Hub Service enables you to create an Operational Data Hub in the cloud. The service is designed to be transparent, scaleable, and easy to deploy. For more details, see <https://www.marklogic.com/product/marklogic-database-overview/data-hub-service/>. To learn more about Data Hub Framework see the [Data Hub Framework website](#).

### 3.31.6 Java Client API: Load Balancer Improvements

As of Java Client API 4.1.1, you can specify a connection type when creating a `DatabaseClient`. The connection type defaults to `DIRECT`, meaning that direct connections to hosts in your MarkLogic cluster are possible. Specify the `GATEWAY` connection type when connecting to MarkLogic through a load balancer. A `GATEWAY` connection is required when using the Data Movement SDK with a load balancer.

For more details, see [Connecting Through a Load Balancer](#) in the *Java Application Developer's Guide*.

## 3.32 New Features in MarkLogic 9.0-8

MarkLogic 9.0-8 introduces the following new features:

- [Request Monitoring Enhancements](#)
- [New Meters Available in Query Meters and Request Logs](#)
- [Synchronize Keystore](#)
- [New Client Issuer Authority Functions](#)
- [Kerberos Authorization and LDAP Authentication Change](#)
- [Logout Capability Added](#)



### 3.32.1 Request Monitoring Enhancements

The following features have been added to Request Monitoring:

- **Request Monitoring on XDBC Server**

The Request Monitoring feature now enables you to setup and enable request monitoring for an endpoint, a main module, or globally on an XDBC Server.

- **Request Cancelling**

The Request Cancelling feature enables you to setup and enable request cancelling for an endpoint, a main module, or globally on an App Server or XDBC Server.

- **New Functions**

The following functions have been added to the Request Monitoring API:

- `xdmp:set-request-limit`
- `xdmp:request-status`

For more details, see [Endpoints and Request Monitoring](#) in the *Query Performance and Tuning Guide*.

### 3.32.2 New Meters Available in Query Meters and Request Logs

The following meters are collected on the server as of MarkLogic Server 9.0-8 and are available in `*.api` files, request logs, and in the output of `xdmp:query-meters()`:

Meter	Definition
<code>compressed-tree-size</code>	The aggregate size in bytes read from disk by unsuccessful compressed tree cache lookups. Each unsuccessful compressed tree cache lookup is followed by a disk access to load the compressed tree into the cache.
<code>list-size</code>	The aggregate size in bytes read from disk by unsuccessful list cache lookups. Each unsuccessful list cache lookup is followed by a disk access to load the search term list into the cache.
<code>read-locks</code>	The number of read locks.
<code>write-locks</code>	The number of write locks.
<code>compile-time</code>	The aggregate time spent compiling a module or a program, represented as a double-precision values in seconds.
<code>run-time</code>	The aggregate time spent evaluating or running a module or a program, represented as a double-precision values in seconds.

Meter	Definition
lock-time	The aggregate time forests spend waiting for transactional read and write locks, represented as a double-precision values in seconds. This time can exceed the run-time.
contemporaneous-timestamp-time	The time spent by queries waiting for the contemporaneous timestamp for which any transaction is known to have committed, represented as a double-precision values in seconds. When the multi-version concurrency control is set contemporaneous, queries can block waiting for the contemporaneous transactions to fully commit.
indexing-time	The indexing time of documents before they are inserted into database, represented as a double-precision values in seconds.
commit-time	The aggregate commit phase time, represented as a double-precision values in seconds.

### 3.32.3 Synchronize Keystore

Customers who use MarkLogic’s Encryption at Rest feature can now synchronize the enveloped keys on MarkLogic with their KMS. For more information, see “[Synchronizing the KMS Keys](#)” in the *Security Guide* or documentation for the function `xcmp:keystore-synchronize`.

### 3.32.4 New Client Issuer Authority Functions

The following XQuery functions have been added to MarkLogic Server:

- [admin:appserver-get-ssl-client-issuer-authority-verification](#)
- [admin:appserver-set-ssl-client-issuer-authority-verification](#)

The following JavaScript functions have been added to MarkLogic Server:

- [admin.appserverGetSslClientIssuerAuthorityVerification](#)
- [admin.appserverSetSslClientIssuerAuthorityVerification](#)

See the online documentation provided in the links for details.

### 3.32.5 Kerberos Authorization and LDAP Authentication Change

MarkLogic 9.0-8 has the parameter `ldap-remove-domain`, enabling you to choose whether or not to remove the domain name. The default value is `false`.

### 3.32.6 Logout Capability Added

MarkLogic 9.0-8 now has a “Logout” button in the upper right corner of the Admin Interface.

### 3.33 New Features in MarkLogic 9.0-9

MarkLogic 9.0-9 introduces the following new features:

- [Kerberos Authorization Parameter Change](#)
- [Task Server Monitoring at Endpoint Level](#)
- [MarkLogic SQL ORDER-BY Keyword](#)
- [Change to ldapSearch Functionality](#)
- [Amazon Linux 2 Support](#)
- [Encryption Support for Thales HSM](#)
- [MarkLogic Supports ECDH Key Exchange for SSL/TLS](#)

#### 3.33.1 Kerberos Authorization Parameter Change

In MarkLogic 9.0-9, the `ldap-remove-domain` parameter default value has been changed to `true` to keep behavior consistent with MarkLogic 9.0-6 and earlier.

#### 3.33.2 Task Server Monitoring at Endpoint Level

MarkLogic 9.0-9 supports endpoint monitoring at the Task Server level. For more information, see the *Query Performance and Tuning Guide*.

#### 3.33.3 MarkLogic SQL ORDER-BY Keyword

Starting in MarkLogic 9.0-9, `ORDER-BY` ordering on `NULL` is set to `NULLS LAST`. This changes default behavior for nulls in query results and therefore may cause backwards incompatibility for certain queries (unless the `Optic Nulls Smallest On trace` event is turned on).

#### 3.33.4 Change to ldapSearch Functionality

An `ldapSearch` call in JavaScript now returns an array if there are multiple values associated with a key. Previously it only returned a single value for every key, even if there were multiple values associated with that key.

#### 3.33.5 Amazon Linux 2 Support

MarkLogic Amazon Web Services (AWS) AMIs are now supported on Amazon Linux 2. This is the recommended base image on which to deploy customized MarkLogic images.

#### 3.33.6 Encryption Support for Thales HSM

With MarkLogic 9.0-9, MarkLogic supports the Thales nCipher nShield Connectd for Encryption at Rest for both Windows and Linux platforms. Thales HSM (Hardware Security Module) implements PKCS #11 (the IETF standard) and when configured, replaces our internal keystore (SoftHSM) for all encryption operations.

### 3.33.7 MarkLogic Supports ECDH Key Exchange for SSL/TLS

MarkLogic 9.0-9.1, ECDH is a supported cipher for SSL/TLS communication. SSL/TLS works if a ECDH cipher is specified.

## 3.34 New Features in MarkLogic 9.0-10

MarkLogic 9.0-10 introduces the following new features:

- [OpenSSL Upgrade to Version 1.0.2s](#)
- [Updates to the BiText Dictionaries](#)

### 3.34.1 OpenSSL Upgrade to Version 1.0.2s

OpenSSL has been upgraded to version 1.0.2s. For more information, please see the list of changes [here](#).

### 3.34.2 Updates to the BiText Dictionaries

Updates to the BiText Dictionaries:

In most cases, new entries/words come from neologisms or newly-coined words that have a significant number of occurrences in corpus (always including any necessary inflection). In some cases, new entries/words come from previously missing words (normally due to too low occurrences in corpus in previous iterations).

For English we have made a wider change: in our previous version, comparative and superlative forms of adjectives were their own lemmas (so, the lemma for “nicest” was “nicest”, and the lemma for “nicer” was “nicer”). Now we have linked those forms to the positive form, with that positive form as lemma (so, in that new version, the lemma for “nicest” and “nicer” is “nice”). Anyway, as we weren’t sure how this change could affect to your operations, we have prepared two versions of the updated English data:

**oldCS.txt**, containing the comparatives and superlatives of adjectives with the comparative or superlative as lemma.

**newCS.txt**, containing the comparatives and superlatives of adjectives with the positive form as lemma.

Below are the details of the changes per language (amount of newly added words):

English (ENG): Around 4,500 new entries added (89,248 è 93,827).

Arabic (ARB): Around 5,500 new entries added (24,060,893 è 24,066,479).

Dutch (NLD): Around 25,000 new entries added (376,185 è 401,741).

French (FRA): Around 700 new entries added to the “non-clitics” dictionary (375,713 è 376,449). Around 250 new entries added to the “clitics” dictionary (985,407 è 985,676).

German (DEU): Around 26,000 new entries added (1,132,410 è 1,158,531).

Italian (ITA): Around 900 new entries added (3,320,525 è 3,321,417).

Korean (KOR): Around 20,000 new entries added (5,013,103 è 5,032,966).

Norwegian (NOR): Around 400 new entries added to the Bokmål dictionary (NOB) (378,738 è 379,112). Around 28 entries added to the Nynorsk dictionary (NNO) (328,872 è 329,004).

Persian (FAS): Around 1,200 new entries added (403,038 è 404,283).

Portuguese (POR): Around 2,000 new entries added (10,861,645 è 10,863,756).

Russian (RUS): Around 1,200 new entries added to the “standard” dictionary (1,447,917 è 1,449,175). 16 entries added to the “orthographic yo/ye” dictionary (61,236 è 61,252).

Spanish (ESP): Around 7,500 new entries added (3,317,323 è 3,324,939).

Swedish (SWE): Around 1,000 new entries added (424,824 è 425,862).

### 3.35 New Features in MarkLogic 9.0-11

MarkLogic 9.0-11 introduces the following new features:

- [Request Monitoring Enhancements](#)
- [Using MarkLogic Encryption with Microsoft Azure Key Vault](#)

#### 3.35.1 Request Monitoring Enhancements

The logging for each request can be controlled by adding a threshold. The request information is logged only when the threshold conditions are met.

Request monitoring is now enabled by default on some MarkLogic application servers:

- The default App-Services application server, usually port 8000, for all query console-related requests, and for non-query console related requests with a runtime exceeding one second.
- The default Manage application server, usually port 8002, for requests that run longer than one second.

For more information, see the MarkLogic *Query Performance and Tuning Guide*.

#### 3.35.2 Using MarkLogic Encryption with Microsoft Azure Key Vault

Microsoft Azure Key Vault can encrypt your data in MarkLogic. Azure Key Vault is supported for customers running their cluster on Microsoft Azure. For more information, see the MarkLogic *Security Guide*.

### 3.36 New Features in MarkLogic 9.0-12

MarkLogic 9.0-12 introduces the following new features:

- [Rolling Upgrade Status Added](#)
- [Other Changes](#)

### 3.36.1 Rolling Upgrade Status Added

MarkLogic 9.0-12 now has an Upgrade tab in the Admin Interface. During an upgrade, click the Upgrade tab to view the upgrade status of each host in the cluster. For more details, see [Rolling Upgrade Status in Admin UI](#) in the *Administrator's Guide*.

### 3.36.2 Other Changes

Swap space is automatically configured when running MarkLogic Server on Amazon Web Services (AWS). Swap space is configured during the system startup process with the `MARKLOGIC_AWS_SWAP_SIZE` configuration variable. For more details, see [AWS Configuration Variables](#) and [Deployment and Startup](#) in the *MarkLogic Server on Amazon Web Services (AWS) Guide*.

The Managed Cluster feature supports SSL-enabled clusters. For details, see [The Managed Cluster Feature](#) in the *MarkLogic Server on Amazon Web Services (AWS) Guide*.

## 3.37 New Features in MarkLogic 9.0-13

MarkLogic 9.0-13 introduces the following new features:

- [Packages by Linux Platform Updated](#)
- [IAM Permissions Updated](#)

### 3.37.1 Packages by Linux Platform Updated

Updated the list of packages required for each supported Linux platform. For more details, see [Supported Platforms](#) and [Appendix: Packages by Linux Platform](#) in the *Installation Guide for All Platforms*.

### 3.37.2 IAM Permissions Updated

Updated the minimum required IAM permissions to create and delete a stack. For more details, see [Creating an IAM Role](#) in the *MarkLogic Server on Amazon Web Services (AWS) Guide*.

## 4.0 Known Incompatibilities with Previous Releases

Most applications implemented on previous versions of MarkLogic will run either without modifications or with very minor modifications in MarkLogic 9. This section summarizes product changes that can cause compatibility issues for applications developed using previous releases and includes the following topics:

- [Incompatibilities \(External\)](#)
- [Hadoop Connector Removed](#)
- [HDFS Support Deprecated](#)
- [Ops Director Deprecated](#)
- [JavaScript: ValueIterator Replaced By Sequence](#)
- [Database Stemming is Off, Word Searches On By Default](#)
- [Collection Lexicon and Triple Index Enabled by Default](#)
- [XCC .NET API No Longer Available](#)
- [Changes in Semantic Query Behavior](#)
- [Triple Count Increased After Inserting Same Data Twice](#)
- [Database max merge size Now Defaults to 48 GB](#)
- [Changes to Range Index Reference Resolution](#)
- [Default Stemming and Tokenization Libraries Changed for Most Languages](#)
- [SQL DESCRIBE No Longer Supported by xdm:sql](#)
- [Application-Specific Logging](#)
- [Change to Classification of Some Special Symbol Tokens](#)
- [Change to xdm:user-last-login](#)
- [Changed Interfaces for xdm:document-insert and xdm:document-load](#)
- [search:parse Returns a Different Type for cts:query Output Format](#)
- [Default Client API Search Behavior Change on Port 8000](#)
- [JSON Property Scope and Container Queries Match Array Items Differently](#)
- [REST Client API Incompatibilities](#)
- [Java Client API Incompatibilities](#)
- [Node.js Client API Incompatibilities](#)
- [Geospatial Region Accessors Can Now Return Double Values](#)
- [User-Defined Function Plugins Must Be Recompiled](#)

- [SLES 12 No Longer Supported](#)
- [Solaris No Longer Supported](#)
- [Nagios Plugin No Longer Supported](#)
- [Application Builder and Information Studio No Longer Available](#)
- [Admin Interface No Longer Selects a Default Schemas Database](#)
- [Internal Security ON with External Security Object Behavior Change](#)
- [REST Management API Changes in MarkLogic 9](#)
- [Configuration Packaging Format Incompatibilities](#)
- [Java Client API 4.1.1 Incompatibilities](#)
- [Incompatibilities Between 9.0-10 and 9.0-11](#)
- [Incompatibilities Between 9.0-8 and 9.0-9](#)
- [Incompatibilities Between 9.0-5 and 9.0-6](#)
- [Incompatibilities Between 9.0-4 and 9.0-5](#)
- [Incompatibilities Between 9.0-3 and 9.0-4](#)
- [Incompatibilities Between 9.0-2 and 9.0-3](#)
- [Incompatibilities Between 9.0-1 and 9.0-2](#)
- [MarkLogic 8 Incompatibilities](#)
- [MarkLogic 7 Incompatibilities](#)

## 4.1 Incompatibilities (External)

- [The libnsl.so.1 library is required for CentOS 8](#)

### 4.1.1 The libnsl.so.1 library is required for CentOS 8

For CentOS 8, MarkLogic has a dependency on `libnsl.so.1()`. You must install this library when using CentOS 8. You can either rely on `yum` to pull in the dependency automatically, or install it manually.

## 4.2 Hadoop Connector Removed

The Hadoop Connector was deprecated starting with MarkLogic release 9.0-12 and is removed from the product in MarkLogic 9.0-13.

## 4.3 HDFS Support Deprecated

HDFS support was deprecated in MarkLogic 9.0-12.



## 4.4 Ops Director Deprecated

The Ops Director feature is deprecated in MarkLogic 9.0-13.

## 4.5 JavaScript: ValueIterator Replaced By Sequence

The `ValueIterator` interface used to represent sequences of value in MarkLogic 8 has been replaced by the new `Sequence` interface. A `Sequence` is a JavaScript Iterable object. All functions which previously operated on or returned a `ValueIterator` now use a `Sequence` instead.

In many cases, this change is transparent to your code. However, code that depends on the following `ValueIterator` properties and methods must be changed:

- `ValueIterator.next` - Use a `for..of` loop to iterate over a `Sequence`. Use `fn.head` if you just want to pick off the first or only value in a `Sequence`.
- `ValueIterator.count` - Use `fn.count` instead.
- `ValueIterator.clone` - No longer needed. You can iterate over the same `Sequence` multiple times.

For more details, see [Sequence](#) in the *JavaScript Reference Guide* and [Sequence](#) in the *MarkLogic Server-Side JavaScript Function Reference*.

## 4.6 Database Stemming is Off, Word Searches On By Default

In MarkLogic 9, when you create a new database, the `stemmed searches` property is `off` by default. In MarkLogic 8 and earlier, the default is `basic`.

In MarkLogic 9, when you create a new database, `word searches` are enabled by default. In MarkLogic 8 and earlier releases, `word searches` were disabled by default.

To achieve the pre-MarkLogic 9 default behavior, configure your database to turn off `stemmed searches` and set `word searches` to `true`.

These changes only affect databases you create after upgrading to MarkLogic 9. Databases that exist when you upgrade will retain their previous settings.

## 4.7 Collection Lexicon and Triple Index Enabled by Default

In MarkLogic 9, a fresh install of MarkLogic 9 will enable the triple index and collection lexicon for all databases. Databases that exist when you upgrade to MarkLogic 9 will retain their previous settings. Any new databases created after upgrading to MarkLogic 9 will have the triple index and collection lexicon enabled.

## 4.8 XCC .NET API No Longer Available

The XCC .NET interfaces have been removed from MarkLogic. Current users of XCC .NET are encouraged to use the REST Client API to create equivalent interfaces. For details, see the *REST Application Developer's Guide*.

The XCC Java Library continues to be available.

## 4.9 Changes in Semantic Query Behavior

MarkLogic 9 introduces the following changes to the behavior of semantic queries:

- [Triple Index and SPARQL Engine Changes](#)
- [Forest IDs Removed From sem:sparql Function](#)

### 4.9.1 Triple Index and SPARQL Engine Changes

The triple index and SPARQL engine have been changed in MarkLogic 9. Now as part of a query, two triples are considered equal if their subjects, predicates, and objects compare as equal with the SPARQL “=” operator.

Previously, these triples would be treated as two different, non-identical triples:

```
sem:triple(xs:anyURI("http://a"), xs:anyURI("http://b"),
xs:anyURI("http://c")),
sem:triple("http://a", "http://b", "http://c").
```

In MarkLogic 9, values of type `xs:string` and `xs:anyURI` can compare equal if they have the same lexical form according to the SPARQL “=” operator. This functionality is called D-entailment (D as in datatype). The two triples in the example are now considered to be the same, and the de-duplication process removes the duplicate. If you return frequencies, you would see the second triple show up as one extra in the frequencies.

Also note that the query may return either `sem:triple(xs:anyURI("http://a"), xs:anyURI("http://b"), xs:anyURI("http://c"))`, or `sem:triple("http://a", "http://b", "http://c")` as the result since they are considered equal.

### 4.9.2 Forest IDs Removed From sem:sparql Function

The forest ID options for the `sem:sparql` function, which were part of MarkLogic 7, were removed in MarkLogic 8, but their functionality was kept for backwards compatibility. In MarkLogic 9, the forest ID option will no longer work with `sem:sparql`.

## 4.10 Triple Count Increased After Inserting Same Data Twice

During a rolling upgrade, while the cluster is in a mixed mode (not all hosts committed to the new version yet) the triple count of semantic triples may be increased after inserting the same data twice. During a rolling upgrade, the MarkLogic 9 triple index is not able to return triples in the correct order for MarkLogic 8 semantics. For this situation to occur, a user would need to have multiple triples that are identical except for the types of the values.

## 4.11 Database max merge size Now Defaults to 48 GB

The database parameter `max merge size` defaults to 48 GB (49152 MB) for new databases created in MarkLogic 9. Previously, the default was 32 GB. Existing databases will keep whatever `max merge size` setting is configured after upgrading to MarkLogic 9. The new setting reflects advances in storage systems and should be appropriate for most databases.

## 4.12 Changes to Range Index Reference Resolution

In MarkLogic 9, index reference resolution in range index construction, query constructors, and lexicon operations will now succeed in some cases that would have thrown an `XDMP-RIDXNOTFOUND` error in previous releases. If the index reference contains enough information to unambiguously match an existing index, MarkLogic will now do so. Previously, MarkLogic assumed default values for some “missing” index attributes, resulting in an error.

Code that previously succeeded will continue to do so. Some code that would previously have gotten an error will no longer do so.

For example, suppose you define a geospatial element range index on element `xs:QName("coords")` with type `long-lat-point` and coordinate system `wgs84`, and then refer to the index by just the element `QName(cts:element-geospatial-values(xs:QName("coord")))`. Previously, the reference would have been an error because the default point type (`point`) would have been assumed. Now, the reference resolves correctly as long as there is not a second geospatial element range index on the same `QName` with different coordinate system or point type.

When an ambiguous range index reference is detected, MarkLogic 9 throws one of the new exceptions `XDMP-RIDXAMBIGUOUS` (range index) or `XDMP-GIDXAMBIGUOUS` (geospatial index).

## 4.13 Default Stemming and Tokenization Libraries Changed for Most Languages

In MarkLogic 9 the default tokenization and stemming libraries have been changed for all languages (except English tokenization). Consequently, some tokenization and stemming behavior will change between MarkLogic 8 and MarkLogic 9. We expect that, in most cases, stemming and tokenization will be more precise in MarkLogic 9.

If you upgrade to MarkLogic 9 from an earlier version of MarkLogic, your installation will continue to use the legacy stemming and tokenization libraries as the language baseline. Any fresh installation of MarkLogic will use the new libraries. You can change the baseline configuration using `admin:cluster-set-language-baseline`.

**Note:** Changing the baseline requires a cluster-wide restart and a reindex to avoid stemming and tokenization anomalies.

**Note:** Use of the legacy libraries is deprecated. These libraries will be removed from MarkLogic in a future release.

Unless you use the legacy language baseline, reindexing is required for content in the following languages:

- Chinese
- Danish
- Dutch, if you want to query with compounding
- Finish
- German
- Hungarian
- Japanese
- Korean, unless you use compounding
- Norwegian (Bokmal and Nynorsk) if you want to query with compounding
- Norwegian (generic ‘no’ lang code), though use of generic ‘no’ is not recommended
- Romanian
- Russian
- Swedish, if you want to query with compounding
- Tamil
- Turkish

For other languages (except English), you might be able to avoid incompatibilities depending on the nature of your queries, but reindexing is still strongly recommended.

Tokenization and stemming are significantly different for Japanese. Tokenization is significantly different for Chinese (both simplified and traditional). The impact on other languages is more nuanced, but should lead to better results, overall. You might observe some relevance score changes on stemmed searches due to the higher degree of precision. If you require low-level details about the impact on a specific language and you have an active maintenance contract, you can contact MarkLogic Technical Support.

For more details on incompatibilities related to the changes to stemming and tokenization, see the Knowledge Base article entitled “MarkLogic Server v9 Tokenization and Stemming” from MarkLogic Technical Support, available at the following URL:

<https://help.marklogic.com/knowledgebase/article/View/484>

For more information about the new tokenization and stemming support, see “New Stemming and Tokenization” on page 21.

#### 4.14 SQL DESCRIBE No Longer Supported by xdm:sql

In MarkLogic 9, the SQL `DESCRIBE` function is not supported. MarkLogic does support `DESCRIBE` queries over ODBC, but not from `xdmp:sql()`.

#### 4.15 Application-Specific Logging

MarkLogic 9 now provides application-specific access to logging. Instead of logging all errors to `ErrorLog.txt`, the `xdmp:log` function now writes to an app-server-specific log file - for example `8000_ErrorLog.txt` or `8020_ErrorLog.txt`. Anything event running on the task server is logged to `TaskServer_ErrorLog.txt`. The `ErrorLog.txt` file is now for MarkLogic “system” logging only.

Splitting the log files in this manner enables customer log files (which could potentially contain confidential information) to be separated from system log files. This helps ensure the privacy of customer data when they interact with Support and use the new Telemetry feature.

In a Unix environment you can watch the multiple log files using something like this:

```
tail -f /path/file1 /path/file2 etc.
```

#### 4.16 Change to Classification of Some Special Symbol Tokens

The classification of some symbols has changed for purposes of tokenization, including the symbols in the following table. These changes can affect search results.

Symbols	Old Classification	New Classification
spacing accents (5E, 60, A8, AF, B4, B8)	punctuation	diacritic
copyright (A9) registered (AE) degree (B0)	punctuation	symbol
Spanish masculine/feminine ordinals (AA, BA)	punctuation	diacritic
superscript numbers (B2, B3, B9)	punctuation	diacritic
micro (B5)	punctuation	greek
fractions (BC, BD, BE)	punctuation	symbol

#### 4.17 Change to xdm:user-last-login

The `userid` parameter has been removed from `xdmp:user-last-login`. In MarkLogic 8 the `xdmp:user-last-login` took one parameter (`userid`). If the `userid` was different from the current user, the function returned an empty sequence. If the `userid` was the same as that of the current user, it only returned the last login information.

In MarkLogic 9 the function does not take a parameter. The `xdmp:user-last-login` function only returns the last login information for the current user.

#### 4.18 Changed Interfaces for `xdmp:document-insert` and `xdmp:document-load`

These two functions, `xdmp:document-insert` and `xdmp:document-load`, have significantly changed interfaces in MarkLogic 9. The interfaces will still work for purposes of backward compatibility, but are no longer documented.

#### 4.19 `search:parse` Returns a Different Type for `cts:query` Output Format

In MarkLogic 8, `search:parse` returned the XML serialization of a `cts:query` by default. You could also explicitly select this return type by specifying “`cts:query`” as the value of the `$output` parameter.

In MarkLogic 9, if you explicitly specify “`cts:query`” as the value of the `$output` parameter, you will now get a `cts:query` object instead of a serialized query. The default return type from `search:parse` is unchanged.

If your code explicitly sets the output type to “`cts:query`” and passes the output of `search:parse` straight through to functions such as `search:resolve` or `cts:search`, no code changes are required.

If your code explicitly sets the output type to “`cts:query`” and then transforms or traverses the output query XML, then you must change your code to use the new “`schema-element(cts:query)`” for the `$output` parameter of `search:parse`.

For example, the following call generates the XML serialization of a `cts:query` in MarkLogic 9:

```
search:parse("cat AND dog", (), "schema-element(cts:query)")
```

#### 4.20 Default Client API Search Behavior Change on Port 8000

The defined default search behavior for the Java, Node.js, and REST Client APIs is unfiltered search unless you override the default with your own options. Prior to MarkLogic 9, this default behavior was not honored when you used the Client APIs to interact with MarkLogic through the pre-defined App Server on port 8000.

As of MarkLogic 9, searches via the Client APIs on port 8000 will default to unfiltered search. To get the old behavior, use query options that explicitly specify filtered search.

#### 4.21 JSON Property Scope and Container Queries Match Array Items Differently

In MarkLogic 8, the query constructed by the XQuery function `cts:json-property-scope-query` and the Server-Side JavaScript function `cts.jsonPropertyScopeQuery` matched if it's criteria query matched anywhere within the configured scope. In MarkLogic 9, the behavior has changed for certain JSON property scope queries where the scope property value is an array of objects.

The behavior has changed for searches that have the following characteristics:

- The `query` parameter of the property scope query is an and-query. Such a scope query effectively finds co-occurrences of matches to the and-query criteria within the specified scope.
- In the document to be matched, the value of the scope property is an array of objects. The item type is determined by examining only the first item in the array.

Given this setup, in MarkLogic 9, the query only matches if all the sub-queries of the and-query occur in the same array item. In MarkLogic 8, the query matches even when the and-query matches occur in different array items. All other forms of JSON property scope query are unchanged.

This change makes it possible to construct a JSON property scope query that constrains matches to occurrences with a single array item. In MarkLogic 8, this was not possible.

This change also affects the `container-query` element of a structured query and QBE container queries.

The following example query is of the form affected by this change. It finds co-occurrences of “prop1” with value “value1” and “prop2” with “value2” within the scope of “root”. (The and-query criteria can be arbitrarily complex.)

Language	Example Query
XQuery	<pre>cts:json-property-scope-query(   "root",   cts:and-query((     cts:json-property-value-query("prop1", "value1"),     cts:json-property-value-query("prop2", "value2")   )) )</pre>
Server-Side JavaScript	<pre>cts.jsonPropertyScopeQuery(   'root',   cts.andQuery([     cts.jsonPropertyValueQuery('prop1', 'value1'),     cts.jsonPropertyValueQuery('prop2', 'value2')   ]) );</pre>

If you search the following documents with the preceding query, you get the results shown. The JSON properties that match the and-query criteria are shown in bold.

Sample Document	MarkLogic 8	MarkLogic 9
<pre>// (1) criteria met in different array items {"root": [   { <b>"prop1": "value1"</b>, "prop2": "v" },   { "prop1": "v", <b>"prop2": "value2"</b> } ]}</pre>	Match	No match
<pre>// (2) criteria met in the same array item {"root": [   { <b>"prop1": "value1"</b>, <b>"prop2": "value2"</b> },   { "prop1": "v", "prop2": "v" } ]}</pre>	Match	Match
<pre>// (3) criteria met in a child property {"root": {   "child": [     { <b>"prop1": "value1"</b>, "prop2": "v" },     { "prop1": "v", <b>"prop2": "value2"</b> }   ] }}</pre>	Match	Match

Document (1) does not match in MarkLogic 9 because both and-query criteria are not satisfied in the same array item. Document (2) matches in MarkLogic 9 because both and-query criteria are satisfied in the same array item. The Document (3) results are unaffected because the value of the “root” property is not an array of objects.

You can restore the MarkLogic 8 behavior by using an and-query of multiple JSON property scope queries. For example, the following query matches Document (1) in MarkLogic 9:

Language	Example Query
XQuery	<pre>cts:and-query((   cts:json-property-scope-query(     "root", cts:json-property-value-query("prop1", "value1")),   cts:json-property-scope-query(     "root", cts:json-property-value-query("prop2", "value2")) ))</pre>
Server-Side JavaScript	<pre>cts.andQuery([   cts.jsonPropertyScopeQuery(     'root', cts.jsonPropertyValueQuery('prop1', 'value1')),   cts.jsonPropertyScopeQuery(     'root', cts.jsonPropertyValueQuery('prop2', 'value2')) ]);</pre>



To ensure the modified queries match only when the and-query matches occur in the same instance of the scope property (“root” in the examples), wrap it in a JSON property scope query on the parent parent property. For example:

Language	Example Query
XQuery	<pre>cts:json-property-scope-query("parent-of-root",   cts:and-query((     cts:json-property-scope-query(       "root", cts:json-property-value-query("prop1", "value1")),     cts:json-property-scope-query(       "root", cts:json-property-value-query("prop2", "value2"))   )))</pre>
Server-Side JavaScript	<pre>cts.jsonPropertyScopeQuery('parentOfRoot',   cts.andQuery([     cts.jsonPropertyScopeQuery(       'root', cts.jsonPropertyValueQuery('prop1', 'value1')),     cts.jsonPropertyScopeQuery(       'root', cts.jsonPropertyValueQuery('prop2', 'value2'))   ]));</pre>

## 4.22 REST Client API Incompatibilities

MarkLogic 9 introduces the following backward incompatible changes to the REST Client API.

- [keyvalue Service Removed](#)
- [Collections in Request Parameters are OR Related](#)
- [Change in Default rest-reader and rest-writer Permissions](#)
- [Default value of Document Management “repair” parameter changed](#)

### 4.22.1 keyvalue Service Removed

The previously deprecated `/keyvalue` is no longer available. That is, the method GET `/v1/keyvalue` is no longer available.

To perform similar operations, use the `/search` service with a structured query or combined query, or the `/qbe` service.

### 4.22.2 Collections in Request Parameters are OR Related

When you specify multiple collections through the collection request parameter of the following methods, they are now OR related:

- GET and POST `/v1/search`
- GET and POST `/v1/qbe`

- GET and POST `/v1/values/name`

The new behavior is consistent with the semantics of `cts:collection-query` and the structured query `collection-query`.

To get AND semantics, construct your own `and-query` of collection URIs as part of a structured or combined query, or by defining a constraint in your query options.

The new behavior differs from previous behavior in the following way:

- In MarkLogic 8.0-6, the GET methods applied AND semantics to multiple collection parameters, while the POST methods applied OR semantics.
- In MarkLogic 8.0-5 and earlier, both GET and POST methods applied AND semantics to multiple collection parameters.

### 4.22.3 Change in Default rest-reader and rest-writer Permissions

In MarkLogic 9, when you write documents with the REST API using PUT `v1/documents`, the documents have the union of the following permissions:

- Any permissions specified in the request
- The default permissions for the user and document or, when update policy is set to `overwrite-metadata`, the existing permissions on the document
- Read permission for the rest-reader role and update permissions for the rest-writer role

In MarkLogic 10, when you write documents with the REST API using PUT `v1/documents`, the documents have the union of the following permissions:

- Any permissions specified in the request
- The default permissions for the user and document or, when update policy is set to `overwrite-metadata`, the existing permissions on the document

In other words, what has changed is that the rest-reader and rest-writer convenience roles no longer have any permissions on a document unless one of the following is true:

- The request specifies permissions for the rest-reader or rest-writer role
- The definition of the user grants default permissions to the rest-reader or rest-writer role

Since the rest-writer convenience role default permissions grant reader permission to the rest-reader role and update permissions to the rest-writer role, documents written by a user who has the rest-writer convenience role are readable by users with the rest-reader role and writable by users with the rest-writer role.

- In MarkLogic 9, a user given the rest-reader or rest-writer role has access to every document written with the REST API.

- In MarkLogic 10, a security model need not grant any role access to every document written with the REST API

The expected behavior is for explicit permissions to replace rest-reader and rest-writer permissions. This change has been made to MarkLogic 10 to provide more control over the security model.

#### **4.22.4 Default value of Document Management “repair” parameter changed**

The default value of the `repair` parameter in the Document Management PUT `/v1/documents` function has changed in the following way:

- In Marklogic 8.0 and earlier, the default value for `repair` is `full`.
- In Marklogic 9.0 and later, the default value for `repair` is `none`.

### **4.23 Java Client API Incompatibilities**

MarkLogic 9 introduces the following backwards incompatible changes to the Java Client API:

- [Java Client API: Removal of Deprecated Interfaces](#)
- [Java Client API: JAR File Name and Maven Artifact ID Change](#)
- [Logging Turned Off by Default](#)

### 4.23.1 Java Client API: Removal of Deprecated Interfaces

Most of the previously deprecated packages, interfaces, classes, and methods of the Java Client API have been removed. The following table lists what has been removed and provides guidance for modifying your code.

In the table below, “c.m.c.” is an abbreviation for “com.marklogic.client.”.

Removed Package, Class, or Method	Alternative
c.m.c.admin.ServerConfigurationManager: <ul style="list-style-type: none"> <li>• <code>getContentVersionRequests</code></li> <li>• <code>setContentVersionRequests</code></li> </ul>	Use the following equivalent methods instead. <ul style="list-style-type: none"> <li>• <code>getUpdatePolicy</code></li> <li>• <code>setUpdatePolicy</code></li> </ul>
c.m.c.admin.TransformExtensionsManager: <ul style="list-style-type: none"> <li>• <code>writeXqueryTransform</code></li> <li>• <code>writeXSLTransform</code></li> </ul>	Overloads of the listed methods that accept a <code>paramTypes</code> map have been removed.  Use one of the overloads that does not accept parameter metadata. This metadata is not required to install or use a transform.
c.m.c.admin.config	This package, its sub-packages, and all contained classes, interfaces, and types have been removed. This includes the <code>QueryOptions</code> and <code>QueryOptionsBuilder</code> classes.  Create query options using your preferred JSON or XML libraries, read options from a file, or build options as a string. Read and write options using appropriate handles, such as <code>DOMHandle</code> , <code>JacksonHandle</code> , <code>FileHandle</code> , or <code>StringHandle</code> .
c.m.c.document.JSONDocumentManager: <ul style="list-style-type: none"> <li>• <code>getLanguage</code></li> <li>• <code>setLanguage</code></li> </ul>	Remove your usage. Language specifications are not supported for JSON. Calls to these methods were ignored in MarkLogic 8.
c.m.c.io.QueryOptionsHandle	Create query options using your preferred JSON or XML libraries, read options from a file, or build options as a string. Read and write options using appropriate handles, such as <code>DOMHandle</code> , <code>JacksonHandle</code> , <code>FileHandle</code> , or <code>StringHandle</code> .  This class is no longer needed with the removal of <code>QueryOptions</code> and <code>QueryOptionsBuilder</code> .

Removed Package, Class, or Method	Alternative
<code>c.m.c.query:</code> <ul style="list-style-type: none"> <li>• <code>KeyValueDefinition</code></li> <li>• <code>KeyLocator</code></li> <li>• <code>ElementLocator</code></li> </ul>	<p>The key-value search capability has been removed. Build equivalent queries using a <code>StructuredQueryDefinition</code> (JSON property query or element query) or QBE.</p>
<code>c.m.c.query.QueryManager:</code> <ul style="list-style-type: none"> <li>• <code>newKeyValueDefinition</code></li> <li>• <code>newElementLocator</code></li> <li>• <code>newKeyLocator</code></li> </ul>	<p>The key-value search capability has been removed. Build equivalent queries using a <code>StructuredQueryDefinition</code> (JSON property query or element query) or QBE.</p>
<code>c.m.c.io.SearchHandle.forceDOM</code>	<p>Remove your usage. This setting was ignored in MarkLogic 8.</p>
<code>c.m.c.query.StructuredQueryBuilder:</code> <ul style="list-style-type: none"> <li>• <code>FragmentScope.DOCUMENT</code></li> </ul>	<p>Use <code>FragmentScope.DOCUMENTS</code>.</p>

Removed Package, Class, or Method	Alternative
StructuredQueryBuilder (nested classes): <ul style="list-style-type: none"> <li>• AndNotQuery</li> <li>• AndQuery</li> <li>• BoostQuery</li> <li>• CollectionConstraintQuery</li> <li>• CollectionQuery</li> <li>• CustomConstraintQuery</li> <li>• DirectoryQuery</li> <li>• DocumentFragmentQuery</li> <li>• DocumentQuery</li> <li>• ElementConstraintQuery</li> <li>• GeospatialConstraintQuery</li> <li>• LocksQuery</li> <li>• NearQuery</li> <li>• NotQuery</li> <li>• OrQuery</li> <li>• OrQuery</li> <li>• PropertiesConstraintQuery</li> <li>• PropertiesQuery</li> <li>• RangeConstraintQuery</li> <li>• TermQuery</li> <li>• ValueConstraintQuery</li> <li>• WordConstraintQuery</li> </ul>	Continue to construct structured queries using StructuredQueryDefinition methods, as before. The return type from the query builder methods is always StructuredQueryDefinition now. For example, StructuredQueryDefinition.and() used to return an AndQuery, but now returns a StructuredQueryDefinition.
StructuredQueryBuilder.elementConstraint	StructuredQueryBuilder.containerConstraint

#### 4.23.2 Java Client API: JAR File Name and Maven Artifact ID Change

The JAR file for the Java Client API is now named `marklogic-client-api-version.jar`. Previously, the name was `java-client-api-version.jar`. For example, the JAR file name for version 3.0.6 is `java-client-api-3.0.6.jar`, but the JAR file name for version 4.0.1 is `marklogic-client-api-4.0.1.jar`.

The Maven artifact ID for the Java Client API is now `marklogic-client-api` instead of `java-client-api`. Update your build dependency configuration files, such as `pom.xml` or `build.gradle`, accordingly.

### 4.23.3 Logging Turned Off by Default

The Logback library is no longer included in the Java Client API distribution. Logging is now off by default. To re-enable logging, include Logback or another slf4j JAR in your classpath.

## 4.24 Node.js Client API Incompatibilities

The following incompatible changes have been made to the Node.js Client API:

- [Changes to Return Value of documents.remove](#)
- [Transaction Creation Returns an Object by Default](#)
- [Default Search Result Slice is Zero-Based](#)

### 4.24.1 Changes to Return Value of documents.remove

The method `DatabaseClient.documents.remove` previously returned an object with both a "uri" property and a "uris" property that served the same purpose. The "uri" property has been removed. The return value now has the following form:

```
{ "uris": [uri1, uri2, ...], "removed": true }
```

The value of the "uris" property is now always an array. Previously, if you passed in just a single string as input, `remove` returned just the URI string, rather than an array of one item.

### 4.24.2 Transaction Creation Returns an Object by Default

Previously, `DatabaseClient.transactions.open` returned a transaction id string by default, and you could use the `withState` parameter to request a transaction object instead. The use of the string form was deprecated as of MarkLogic 8.

As of MarkLogic 9 and Node.js Client API v2.0.0, `DatabaseClient.transactions.open` defaults to returning a transaction object.

To force the previous behavior, set `withState` to `false` when creating a transaction. This setting is deprecated and will be removed in a future release.

### 4.24.3 Default Search Result Slice is Zero-Based

Previously, the slice clause on search (`queryBuilder.slice`) accepted a one-based starting position and a page length:

```
slice(oneBasedStart, PageLength)
```

As of MarkLogic 9 and Node.js Client API v2.0.0, `queryBuilder.slice` clause behaves like `Array.prototype.slice`. That is, it takes a zero-based starting position and the (zero-based) position after the last result to be retrieved. For example, the following slice call returns the first 5 results:

```
... .slice(0,5) ...
```

Also, you could previously use `slice(0)` to suppress the return of result documents and just retrieve an abbreviated summary. Now, you must include both the start and end positions for the same effect. For example: `slice(0,0)`.

To restore the legacy behavior, use `marklogic.setSliceMode`. Note, however, that this form is deprecated and will be removed in a later release.

**Note:** The semantics of `valuesBuilder.slice` are unchanged. This function still accepts as 1-based starting position and a page length.

## 4.25 Geospatial Region Accessors Can Now Return Double Values

The introduction of support for double precision coordinates in MarkLogic 9 means that some geospatial operations may return different results than in the past.

Previously, geospatial region accessor functions such as the XQuery function `cts:point-latitude` or the Server-Side JavaScript function `cts.pointLatitude` always returned float values. As of MarkLogic 9, these functions can return either single or double precision values, depending on the governing coordinate system. If you do not use a double precision coordinate system, you should not notice a difference.

The following functions are affected.

XQuery Function	JavaScript Function
<code>cts:point-latitude</code>	<code>cts.pointLatitude</code>
<code>cts:point-longitude</code>	<code>cts.pointLongitude</code>
<code>cts:box-west</code>	<code>cts.boxWest</code>
<code>cts:box-east</code>	<code>cts.boxEast</code>
<code>cts:box-south</code>	<code>cts.boxSouth</code>
<code>cts:box-north</code>	<code>cts.boxNorth</code>

For more details, see [How Precision Affects Geospatial Operations](#) in the *Search Developer's Guide*.

## 4.26 User-Defined Function Plugins Must Be Recompiled

The version of MarkLogic's C++ User-Defined Function (UDF) interface has been incremented to accommodate the following changes:

- The `marklogic::Point` class now accepts and returns longitude and latitude values as doubles instead of floats.



- Support for new types of UDF plugins in support of custom stemming and tokenization plugins.

You must recompile your UDF plugins for use with MarkLogic 9. UDFs from an earlier version of MarkLogic will not work in a mixed cluster prior to committing the new version. During a rolling upgrade, you must finish the upgrade and then recompile your UDFs for use with MarkLogic 9.

#### 4.27 SLES 12 No Longer Supported

In MarkLogic 9, the SUSE Linux Enterprise Server operating system is not supported. If you are using this discontinued platform, you will need to migrate your environment to a supported platform. For details on supported platforms, see “Supported Platforms” on page 15.

#### 4.28 Solaris No Longer Supported

In MarkLogic 9, the Solaris operating system is not supported. If you are using this discontinued platform, you will need to migrate your environment to a supported platform. For details on supported platforms, see “Supported Platforms” on page 15.

#### 4.29 Nagios Plugin No Longer Supported

Support for the Nagios monitoring plugin has been discontinued in MarkLogic 9.

You can create your own MarkLogic monitoring integration using the MarkLogic REST Management API; for details see the *Monitoring MarkLogic Guide*.

MarkLogic does not endorse or support any particular 3rd party monitoring integrations. However, the MarkLogic open source community includes integrations with 3rd party monitoring platforms such as New Relic and App Dynamics. For more details see:

- New Relic: <https://github.com/marklogic-community/newrelic-plugin>
- App Dynamics: <https://github.com/Appdynamics/marklogic-monitoring-extension>

In the future, Ops Director will be MarkLogic’s fully integrated monitoring solution. For details, see the *Ops Director Guide*.

#### 4.30 Application Builder and Information Studio No Longer Available

The Application Builder and Information Studio applications have been removed from MarkLogic. The info and infodev APIs remain, but they are deprecated; for details, see “info and infodev APIs Deprecated” on page 134.

An Application Builder application deployed on MarkLogic 8 should continue to work after the upgrade to MarkLogic 9. We do not support any mechanism to support redeploying such an app against a MarkLogic 9 node. The methods and modules needed for this process have been removed.

Porting an Application Builder application running in MarkLogic 7 to MarkLogic 9 is not supported by MarkLogic Server. If you have an active maintenance contract, you can contact MarkLogic Technical Support for guidance in porting this application.

### **4.31 Admin Interface No Longer Selects a Default Schemas Database**

As of MarkLogic 9, you can create a database without an associated schemas database. Earlier versions of MarkLogic required you to specify a schemas database when creating a new database. This change has no impact on users creating databases through the Admin API. However, this change affects the database creation page of the Admin Interface as follows:

The Admin Interface no longer automatically selects the pre-defined Schemas database as the schemas database when you create a database through the UI. Instead, you must explicitly select a schemas database or else you will create a database with no associated schemas database.

Depending on the uses to which you put the database, some operations might fail if there is no associated schemas database. For example, features such as temporal document management and template driven extraction require a schemas database to be associated with the content database.

### **4.32 Internal Security ON with External Security Object Behavior Change**

This section describes some changes related to external security support. You should be aware of these changes, but they do not introduce incompatibilities or necessitate a change to your application.

In MarkLogic 9, you can log into port 7001 on an LDAP account when “internal security” on appserver port 7001 is set to “true” or “false”. In previous versions of MarkLogic, you can only log into port 7001 on an LDAP account if “internal security” on appserver port 7001 is set to “false”.

As of MarkLogic 9, you can assign multiple external security objects to an App Server. When there are multiple external security objects assigned, a MarkLogic user is authenticated and assigned to an external security based on the order in which the external security objects are assigned. In previous version of MarkLogic, you could only assign one external security object to an App Server.

If internal security is enabled for an App Server, then when a user attempts to authenticate with MarkLogic, MarkLogic first checks to see if the user is in the security database. If so, then MarkLogic verifies the credentials against the security database. When this verification fails, the behavior of MarkLogic 8 and earlier versions differs from the behavior of MarkLogic 9 as follows:

- In MarkLogic 8 and earlier, if Security database verification fails, then the login attempt fails with an error.
- As of MarkLogic 9, if Security database verification fails, then MarkLogic attempts to authenticate the user against any external security objects assigned to the App Server. If

there are no external security objects assigned or if the user cannot be authenticated against any assigned external security objects, then the login fails with an error.

Thus, when internal security is enabled, there can be cases where a login will succeed in MarkLogic 9 that would have failed with earlier versions. The behavior is unchanged if there are no external security objects assigned to the App Server or internal security is disabled.

### 4.33 REST Management API Changes in MarkLogic 9

In MarkLogic 9, a number of REST Management API methods have been deleted and their functionality has been moved elsewhere. The following table lists the deleted methods and the new alternative.

Endpoint Deleted in 9.0	Alternative
DELETE /manage/v2/credentials	DELETE:/manage/v2/credentials/properties
GET /manage/v2/credentials	GET:/manage/v2/credentials/properties
PUT /manage/v2/credentials	PUT:/manage/v2/credentials/properties
GET /manage/v2/databases/{id name}/sub-databases	GET:/manage/v2/databases
POST /manage/v2/databases/{id name}/sub-databases	POST:/manage/v2/databases
GET /manage/v2/databases/{id name}/super-databases	GET:/manage/v2/databases
POST /manage/v2/databases/{id name}/super-databases	POST:/manage/v2/databases
DELETE /manage/v2/databases/{sub-id-or-name}/super-databases/{super-id-or-name}	DELETE:/manage/v2/databases/{id name}
GET /manage/v2/databases/{sub-id-or-name}/super-databases/{super-id-or-name}	GET:/manage/v2/databases/{id name}/properties
DELETE /manage/v2/databases/{super-id-or-name}/sub-databases/{sub-id-or-name}	DELETE:/manage/v2/databases/{id name}
GET /manage/v2/databases/{super-id-or-name}/sub-databases/{sub-id-or-name}	GET:/manage/v2/databases/{id name}/properties

### 4.34 Configuration Packaging Format Incompatibilities

Configuration packages created with the Packaging REST API and/or with the Configuration Packaging XQuery library module are not compatible with the new configuration format introduced in MarkLogic release 9.0-5.

These configuration packages cannot be used with the CMA REST API and/or with the Configuration Management API XQuery/JavaScript library modules.

For more details, see the following sections:

- “Configuration Management API (CMA) REST Endpoints” on page 43;
- “Configuration Management API (CMA) XQuery and JavaScript Libraries” on page 43;
- “Packaging API Deprecated” on page 133;
- “Configuration Packaging XQuery Library Deprecated” on page 139.

### 4.35 Java Client API 4.1.1 Incompatibilities

Version 4.1.1 of the Java Client API introduces the following incompatibilities with earlier Java Client API versions:

- [Load Balancer Configuration for DMSDK Jobs](#)

#### 4.35.1 Load Balancer Configuration for DMSDK Jobs

This change does not affect you if you do not connect to MarkLogic through a load balancer or if you do not use the Data Movement SDK (DMSDK) feature of the Java Client API.

You must change the configuration of any `DatabaseClient` objects used to connect to MarkLogic through a load balancer on behalf of a Data Movement SDK job. The following changes apply:

- You must add a connection type parameter value of `GATEWAY` when configuring a `DatabaseClient`.
- Do not use a `FilteredForestConfiguration` to configure connections through a load balancer.

For more details, see [Working with a Load Balancer](#) in the *Java Application Developer's Guide*.

### 4.36 Incompatibilities Between 9.0-10 and 9.0-11

- [The API `trgr:triggers-change-modules-database\(\)` Only Has Two Parameters](#)

#### 4.36.1 The API `trgr:triggers-change-modules-database()` Only Has Two Parameters

The API `trgr:triggers-change-modules-database()` only takes two parameters in 9.0-11.

1. the old database as `xs:anyAtomicType`. This can be an ID `unsignedLong` or a name `xs:string`
2. the new database as `xs:anyAtomicType`. This can be a database's ID `unsignedLong` or a database's name `xs:string`.

For more details and sample code, see our API documentation.

### 4.36.2 Entity Services Features Deprecated

These Entity Services features are deprecated and will be removed in a future release.

**Instance functions:** The instance functions, like `es.instanceConverterGenerate` or `es:instance-converter-generate` are deprecated.

**Model functions:** Of the `es.model/es:model` functions, only `es:model-validate` and `es:modelValidate` are used. The other `model` functions are deprecated.

**Attachment functions:** The attachment functions, like `es:add-attachments` and `es.addAttachments` are deprecated.

**Init functions:** The `init` functions, like `es:init-instance` are deprecated.

**Version translation functions:** The version translation functions, like `es:version-translator-generate` and `es.versionTranslatorGenerate` are deprecated.

**Utility functions:** The utility functions `es:extract-array` and `es:with-namespace` are no longer used and have been deprecated.

## 4.37 Incompatibilities Between 9.0-8 and 9.0-9

- [MarkLogic SQL ORDER-BY Keyword](#)

### 4.37.1 MarkLogic SQL ORDER-BY Keyword

Starting in MarkLogic 9.0-9, `ORDER-BY` ordering on `NULL` is set to `NULLS LAST`. This changes default behavior for nulls in query results and therefore may cause backwards incompatibility for certain queries (unless the `Optic Nulls Smallest On trace` event is turned on).

## 4.38 Incompatibilities Between 9.0-5 and 9.0-6

The following incompatibilities exist between MarkLogic 9.0-5 and MarkLogic 9.0-6:

- [Changes to Accepted XML Character Set](#)

### 4.38.1 Changes to Accepted XML Character Set

As of MarkLogic 9.0-6, parsing of XML documents with an XML declaration that explicitly specifies XML version 1.1 (`version="1.1"`) enforces the XML 1.1 character set. Consequently, you can now create content containing characters not accepted by XML 1.0.

Characters in the XML 1.1 “restricted character” ranges must be given as character entities. This enforcement applies to the following character ranges:

- 0x1-0x8
- 0xB-0xC
- 0xE-0x1F
- 0x7F-0x84
- 0x86-0x9F

The following character ranges that were previously disallowed are now accepted.

- 0x1-0x8
- 0xB-0xC
- 0xE-0x1F

Serializing content that contains characters allowed by XML 1.1 but not XML 1.0 can cause problems for client layers that cannot handle the XML 1.1 character set. You can control whether to serialize XML as version 1.0 or 1.1 by setting the output version in your App Server output options, using methods such as the following:

- **Admin Interface:** Set the output version to “1.0” under `Groups > yourGroup > App Servers > yourAppServer > Output Options`.
- **Admin API:** Call `admin:appserver-set-output-version` with a value of “1.0” for the `$value` parameter.
- **REST Management API:** Send a request to `PUT:/manage/v2/servers/{id|name}/properties` that sets the `output-version` property to “1.0”.

### 4.39 Incompatibilities Between 9.0-4 and 9.0-5

The following incompatibilities exist between MarkLogic 9.0-4 and MarkLogic 9.0-5:

- [Configuration Manager Deprecated](#)
- [Minimum Required Version of HDP is 2.6](#)
- [Reindex Recommended for Geospatial Region Indexes](#)
- [Geospatial Region Query Results Might Differ](#)

- [return-query Option Output Format Change](#)

### 4.39.1 Configuration Manager Deprecated

The Configuration Manager tool is deprecated starting with MarkLogic 9.0-5 and removed from MarkLogic Server in version 10.0-3.

### 4.39.2 Minimum Required Version of HDP is 2.6

As of MarkLogic 9.0-5, if you are using the Hortonworks Data Platform (HDP) with the following MarkLogic technologies, you must use HDP v2.6.

- mlcp
- HDFS
- The MarkLogic Connector for Hadoop

### 4.39.3 Reindex Recommended for Geospatial Region Indexes

This change affects only database configuration that define a geospatial region index.

MarkLogic 9.0-5 introduces support for a “tolerance” option on geospatial region queries. As a consequence, if you use geospatial region indexes, your region queries might not be accurate until you reindex. This is applicable whether or not you take advantage of the new tolerance option.

### 4.39.4 Geospatial Region Query Results Might Differ

As of MarkLogic 9.0-5, MarkLogic uses the coordinate system *default* tolerance when evaluating a geospatial region query. Previously, MarkLogic used the coordinate system *minimum* tolerance.

Consequently, some region queries might not give the same results after upgrading to MarkLogic 9.0-5. Use the “tolerance” option to adjust your region queries, if necessary.

To learn more about tolerance, see [Understanding Tolerance](#) in the *Search Developer’s Guide*.

### 4.39.5 return-query Option Output Format Change

This change only affects users of the REST, Java, and Node.js Client APIs who use the `return-query` query option and generate a search response summary in JSON.

In prior versions of MarkLogic, the `return-query` query option returned a serialized JSON structured query in the `query` property of the search response in JSON, but a serialized `cts:query` in an XML search response. As of MarkLogic 9.0-5, `return-query` always generates a serialized `cts:query`. The serialized JSON representation can be passed to `cts:query` or `cts.query` to reconstruct an in-memory query object on MarkLogic, just as you can with the XML serialization.

This change will not affect you if you do not use a JSON search response, do not use `return-query`, or do not have code that depends on the serialization produced by `return-query`.

## 4.40 Incompatibilities Between 9.0-3 and 9.0-4

The following incompatibilities exist between MarkLogic 9.0-3 and MarkLogic 9.0-4:

- [Redaction: Deterministic Masking Values Differ](#)

### 4.40.1 Redaction: Deterministic Masking Values Differ

MarkLogic 9.0-4 introduces support for salting in the generation of masking values by the `mask-deterministic` redaction function. This means the same text redacted with `mask-deterministic` in earlier versions of MarkLogic will not produce the same masking value by default when redacted using MarkLogic 9.0-4.

To preserve the previous behavior, modify your `mask-deterministic` rules to set the `salt` and `extend-salt` options to an empty value.

For more details, see [mask-deterministic](#) in the *Application Developer's Guide*.

## 4.41 Incompatibilities Between 9.0-2 and 9.0-3

The following incompatibilities exist between MarkLogic 9.0-2 and MarkLogic 9.0-3:

- [Changes to Authentication Behavior with Client Certificate](#)
- [XCC ContentSource.newSession Interface Change](#)
- [Document Digest Authorization Behavior Changed in 9.0-3](#)
- [1-click AMLs, new compatible CloudFormation, and additional upgrade procedures](#)
- [map:new Retains Keys with Empty Values](#)

### 4.41.1 Changes to Authentication Behavior with Client Certificate

MarkLogic Server 9.0-3 authenticates using both a client certificate and a username/password. This provides a greater level of security, by requiring that the user provide a client certificate that matches the specified user. In MarkLogic 9.0-2, if the password is incorrect, but the user has the correct client certificate, and the “`ssl require client certificate`” is false, authentication will succeed. In MarkLogic 9.0-3 if the password is incorrect, but the user has the correct client certificate and “`ssl require client certificate is false`”, authentication fails.

The “`ssl require client certificate`” is an appserver configuration. With the appserver, there are different authentication options: basic, digest, application-level, certificate, and kerberos-ticket. This change only applies to basic, digest or application-level authentication. This behavior only happens when “`ssl require client certificate`” is set to false



In previous versions, MarkLogic would accept username/password for any internal user, once MarkLogic verified that certificate was signed by the selected CA (certificate authority). MarkLogic 9.0-1 and 9.0-2 authentication is restricted to an internal user with a matching common name (CN). In MarkLogic 9.0-3, the username does not need to match the common name (CN).

The behavior in MarkLogic 9.0-3 is the same as it is in MarkLogic 8.0. For more information, see the [Certificate](#) topic in the *Security Guide*.

#### 4.41.2 XCC ContentSource.newSession Interface Change

The XCC method `ContentSource.newSession` has been extended to include a new overload that accepts a `char[]` value for the password parameter, rather than a `String`, to enable more secure handling of passwords.

As a consequence of this change, passing `null` in the password parameter of `newSession` now results in a compile time error because the compiler cannot determine the correct overload without type information. Use a typecast to disambiguate your call.

#### 4.41.3 Document Digest Authorization Behavior Changed in 9.0-3

Prior to MarkLogic 9.0-3, the server did not check the nonce for digest authentication. In MarkLogic 9.0-3, the server checks the nonce, verifies that it is a valid nonce, and verifies that the URI in the Authorization header is same as the request URI.

#### 4.41.4 1-click AMIs, new compatible CloudFormation, and additional upgrade procedures

In order to support 1-click deployment in AWS Marketplace, MarkLogic 9 AMIs have data volume pre-configured (device on `/dev/sdf`). To be compatible with 1-click AMIs, new CloudFormation templates are released on <https://developer.marklogic.com/products/cloud/aws>. To upgrade existing clusters, new 1-click compatible CloudFormation templates are required.

If custom templates or scripts are used, additional steps may be required to handle blank data volume that is part of the new marketplace AMIs. Please find more details about the upgrade procedure on <https://developer.marklogic.com/products/cloud/aws>.

#### 4.41.5 map:new Retains Keys with Empty Values

As of MarkLogic 9.0-4, the `map:new` function retains any input key-value pair whose value is an empty sequence. Previous versions of MarkLogic 9 and versions of MarkLogic 8 prior to 8.0-7 discarded such key-value pairs.

For example:

```
map:keys (map:new (map:entry ("noval", ())))  
  
(: Now returns "noval". Previously, returned an empty sequence. :)
```

## 4.42 Incompatibilities Between 9.0-1 and 9.0-2

The following incompatibilities exist between MarkLogic 9.0-1 and MarkLogic 9.0-2:

- [The mlcp Option `-tolerate\_errors` is Ignored](#)
- [Changes to `jsearch.facets` Output Structure](#)
- [Array Type is Preserved in x509 Certificate with Array-Valued Properties](#)
- [Node.js Client API: `valuesBuilder.slice` is Now Zero-Based](#)
- [Changes to `xdmp:update XQuery Prolog Option`](#)
- [Java Client API 4.0.2 Ignores `HttpClientConfigurator`](#)

### 4.42.1 The mlcp Option `-tolerate_errors` is Ignored

The `-tolerate_errors` option of the `mlcp import` command is deprecated. As of MarkLogic 9.0-2, `mlcp` ignores this option and always behaves as if `-tolerate_errors` is set to true. The option will be removed in a future release.

### 4.42.2 Changes to `jsearch.facets` Output Structure

In previous versions of MarkLogic, calling `jsearch.facets` always produced a JSON object for each facet, with object properties of the form `facetValue:count`. This structure prevented proper sorting of facet values.

As of MarkLogic 9.0-2, the structure of each facet is an array of arrays instead of a JSON object if and only if you include an explicit `orderBy` clause in your facet definition. If you do not use `orderBy`, the output is unchanged.

For more details, see [Sorting Facet Values with OrderBy](#) in the *Search Developer's Guide*.

### 4.42.3 Array Type is Preserved in x509 Certificate with Array-Valued Properties

In MarkLogic 9.0-1, if you use `xdmp.x509CertificateGenerate` to generate a certificate, and the configuration object includes array-valued properties, the array values were encoded as a single string. As of MarkLogic 9.0-2, the array type is preserved. This change applies to any Relative Distinguished Names (RDNs) within a Distinguished Name (DN), such as the `issuer` and `subject` DNs.

For example, in the following snippet, the `issuer.organizationName` property has an array value.

```
var certObj = {
  version: "2",
  serialNumber: "BA0195369CD6B679",
  issuer: {
    countryName: "US",
    stateOrProvinceName: "CA",
```

```

    localityName: "San Carlos",
    organizationName: ["MarkLogic", "Mark Logic"],
    organizationalUnitName: "Eng",
    emailAddress: "jdonner@marklogic.com",
    commonName: "JGD Certificate Authority",
  }, ...
};
var privateKey = ...;
xdmp.x509CertificateExtract (
  xdmp.x509CertificateGenerate(certObj, privateKey)
);

```

If you round trip the generated certificate through `xdmp.x509CertificateExtract`, you will see the following output for `issuer.organizationName` in MarkLogic 9.0-1 vs MarkLogic 9.0-2.

```

// MarkLogic 9.0-1
organizationName: "[\"MarkLogic\", \"Mark Logic\"]"

// MarkLogic 9.0-2 and later
organizationName: ["MarkLogic", "Mark Logic"]

```

If you do not have a certificate containing a multi-valued property, you will not notice any difference in behavior.

#### 4.42.4 Node.js Client API: `valuesBuilder.slice` is Now Zero-Based

Previously, the slice clause on values queries (`valuesBuilder.slice`) accepted a one-based starting position and a page length:

```
slice(oneBasedStart, PageLength)
```

As Node.js Client API v2.0.3, the `valuesBuilder.slice` clause behaves like `Array.prototype.slice`. That is, it takes a zero-based starting position and the (zero-based) position after the last result to be retrieved. For example, the following slice call returns the first 5 results:

```
... .slice(0,5) ...
```

To restore the legacy behavior, use `marklogic.setSliceMode`. Note, however, that this form is deprecated and will be removed in a later release.

#### 4.42.5 Changes to `xdmp:update XQuery Prolog Option`

Previously, setting the XQuery prolog option `xdmp:update` to “false” caused MarkLogic to automatically detect whether a module should be evaluated as an update or a query transaction.

As of MarkLogic 9.0-2, setting the option to “false” tells MarkLogic to treat the code as a query transaction. This could cause your program to get an error if you explicitly set the option to “false” and your code performs an update operation.

The new “auto” option value is equivalent to the previous behavior of “false”.

For related changes, see the following topics:

- “New Server-Side Transaction Controls” on page 33
- “xdmp:transaction-mode XQuery Prolog Option Deprecated” on page 136
- “Deprecation of transaction-mode Option to xdmp:eval” on page 137

#### 4.42.6 Java Client API 4.0.2 Ignores HttpClientConfigurator

As of version 4.0.2, the Java Client API uses OkHttp as its HTTP client for communicating with MarkLogic over HTTP. This change should be transparent to most applications, but imposes the following backward incompatibility on applications that customize their HTTP configuration:

Attaching a configurator based on `HttpClientConfigurator` to a `DatabaseClientFactory` object no longer has any effect on the HTTP configuration. Use the new `com.marklogic.client.extra.okhttpclient.OkHttpClientConfigurator` interface instead. The `HttpClientConfigurator` interface is deprecated and will be removed in a future release.

### 4.43 MarkLogic 8 Incompatibilities

This section describes the incompatibilities between MarkLogic 7 and MarkLogic 8. This is here just for convenience; for the MarkLogic 8 *Release Notes*, see [docs.marklogic.com/8.0/guide/relnotes](https://docs.marklogic.com/8.0/guide/relnotes). The following are the incompatibilities:

- [JSON Related Incompatibilities](#)
- [Semantics Incompatibilities](#)
- [REST and Java Client API Incompatibilities](#)
- [Document Library Services \(DLS\) Repositories Need To Perform A Bulk Upgrade Operation](#)
- [Linux Now Requires Red Hat 6](#)
- [mysql On Linux No Longer Ships With Server](#)
- [Cyrillic Tokenization Changes](#)
- [Application Builder Applications Must Be Re-Deployed in MarkLogic 8](#)
- [Application Builder and Information Studio Links Removed](#)
- [Search API Incompatibilities](#)
- [Locks and Properties Query Built-In Functions Renamed](#)
- [xdmp:uri-content-type Of an XML Document Now Returns application/xml, Can Affect CPF Applications](#)
- [xdmp:function Signature Change](#)
- [Incompatibilities Between 9.0-1 and 9.0-2](#)
- [Incompatibilities Between 8.0-5 and 8.0-6](#)

- [Incompatibilities Between 8.0-3 and 8.0-4](#)
- [Incompatibilities Between 8.0-2 and 8.0-3](#)
- [Incompatibilities Between 8.0-1 and 8.0-2](#)

### 4.43.1 JSON Related Incompatibilities

MarkLogic 9 includes Native JSON support. In MarkLogic 7, there was support for JSON via a set of libraries to convert between JSON and XML. If you are using the MarkLogic 7 JSON support, you will have to migrate your code to use the native JSON support. This should end up being more efficient, but will require you to do some minor code changes. This section lists the incompatibilities related to working with JSON documents:

- [Documents Created as JSON With MarkLogic 7 REST API or mlcp Must Be Converted to Native JSON](#)
- [json:unquotedString Primitive No Longer Available](#)
- [xdmp:to-json and json:transform-to-json Now Returns a document-node\(\)](#)
- [Search, Java, REST: json-key Is Now json-property in Options and Structured Query](#)
- [Java and REST: Default Path Language for JSON Document Patches is Now XPath](#)
- [Java and REST: Specifying a Language for JSON Documents is Deprecated](#)
- [Java and REST: New Restrictions on Patching JSON Content](#)
- [Java and REST: Transforms and Extensions That Manipulate JSON Must Be Rewritten](#)
- [Java and REST: JSON Array Items and Property Values No Longer Distinguishable in QBE](#)
- [Field Range Query and Field Value Query on JSON May Behave Differently](#)

#### 4.43.1.1 Documents Created as JSON With MarkLogic 7 REST API or mlcp Must Be Converted to Native JSON

In previous versions of MarkLogic, you can load JSON documents into MarkLogic using either the REST Client API or the Java Client API. When you do so, the documents are transformed and stored as XML, but are still queryable as and returned as JSON. In MarkLogic 8 this “XML facade” is no longer needed, and the REST and Java Client APIs in MarkLogic 8 do not do the translation to the XML facade anymore. Therefore, if you have any document that were loaded as JSON in MarkLogic 7 and earlier, you must convert those document to native JSON in order to query them as JSON using the REST API.

To help with the conversion, MarkLogic supplies a set of conversion scripts. These scripts do not handle every case, and for any case it does not handle you will have to convert the documents some other way. Because the conversion scripts do not handle all cases, it is very important to do a backup of your database before attempting the conversion. The scripts are located in the `Samples/migrate-scripts` directory under the MarkLogic installation directory (`/opt/MarkLogic` on Linux, `c:/Program Files/MarkLogic` on Windows, and `~/Library/MarkLogic` on Mac OS). The scripts require bash and curl, and on Windows platforms they also require cygwin.

For more details, see the `<marklogic-dir>/Samples/migrate-scripts/README` file.

Generally, the conversion scripts perform the following:

- Converts the documents to native JSON.
- Updates index configurations to reference the JSON content.
- Updates existing saved search options. This includes changing references to the JSON basic namespace to reference the new JSON content and changes occurrences of `json-key` to `json-property`.
- Updates any alerting rules that reference the JSON content.

The scripts do not upgrade your application code. The types of things you will need to change in your application include:

- Modify client code to update structured queries, combined queries, and query options that reference `json-key` and anything in the `http://marklogic.com/xdmp/json/basic` namespace. For more details see “Search, Java, REST: `json-key` Is Now `json-property` in Options and Structured Query” on page 87.
- Rewrite and reinsert any server-side code (for example, transformations or custom constraints) that operated over the internal XML representation of your JSON documents. For details see “Java and REST: Transforms and Extensions That Manipulate JSON Must Be Rewritten” on page 89.
- Modify client code that relies on the `/v1/keyvalues` endpoint for key/value searches over JSON.
- Modify client code to update patch specifications over JSON. For details see “Java and REST: New Restrictions on Patching JSON Content” on page 89.
- Review index settings and queries over document properties, update as needed (because you can no longer have JSON properties).

The basic steps to upgrade your MarkLogic 7 or earlier JSON to native JSON in MarkLogic 8 are as follows:

1. Backup the database in which your JSON documents exist.
2. Make copies and edit the connection details and other information in the various configuration files in the `Samples/migrate-scripts/conf` directory. These files have details about your configuration and index settings.
3. Run the `Samples/migrate-scripts/migrate` script.
4. Test your results. Make sure the index changes that the scripts made match your newly converted JSON data. It is especially important to review path and fields indexes to make sure they are including the same content in the converted JSON as they were previously.

If you have problems upgrading your application and you have an active maintenance contract, you can contact MarkLogic Technical Support.

#### 4.43.1.2 json:unquotedString Primitive No Longer Available

MarkLogic 7 had a primitive to convert an XQuery string to an unquoted String called `json:unquotedString`. In MarkLogic 8, that function is no longer available because it is no longer needed, as MarkLogic 8 has much more extensive support for JSON. For details on working with JSON in MarkLogic, see [Working With JSON](#) in the *Application Developer's Guide*.

#### 4.43.1.3 xdmp:to-json and json:transform-to-json Now Returns a document-node()

In MarkLogic 8, the `xdmp:to-json` function returns a `document-node()`; previously, it returned a string. If you have code that expects a JSON string, you might need to modify your code to perform XPath on the `document-node()` to get the JSON node (which will serialize into a string); depending on what your code does, you might or might not need to do this. For example:

```
(: 7.0 :)
xdmp:to-json(("a", fn:false()))
=> ["a", false]

(: 8.0 :)
xdmp:to-json(("a", fn:false()))/node()
=> ["a", false]
```

**Note:** The `json:transform-to-json` function uses `xdmp:to-json`, so it also returns a `document-node()` in MarkLogic 8.

#### 4.43.1.4 Search, Java, REST: json-key Is Now json-property in Options and Structured Query

All occurrences of `json-key` in query options and structured query are now `json-property`. If you use the constructs listed below to search JSON documents by key/property name, you will need to modify your query options (`search:options`, in XML) or structured queries.

The following query options are affected. For details, see `search:search` or [Appendix: Query Options Reference](#) in the *Search Developer's Guide*.

- `container-constraint`
- `extract-metadata`
- `range-constraint`
- `sort-order`
- `value-constraint`
- `word-constraint`

The following structured query components are affected. For more details, see the structured query [Syntax Reference](#) in the *Search Developer's Guide*.

- `container-query`
- `range-query`
- `value-query`
- `word-query`

The corresponding Java Client API structured query builder method name has also changed. `StructuredQueryBuilder.JSONKey` is now `StructuredQueryBuilder.JSONProperty`.

#### 4.43.1.5 Java and REST: Specifying a Language for JSON Documents is Deprecated

Previously, you could specify a language when ingesting JSON documents. This was only possible because JSON documents were represented internally as XML.

This parameter is now deprecated and will be ignored when present. This affects the following interfaces:

- **Java:** `JSONDocumentManager.setLanguage` and `JSONDocumentManager.getLanguage` are deprecated. Calling `setLanguage` has no effect.
- **REST:** The `lang` request parameter of `PUT:/v1/documents` is deprecated and will be ignored.
- **REST:** The `lang` request parameter of `POST:/v1/documents` (all variants) is deprecated and will be ignored.

#### 4.43.1.6 Java and REST: Default Path Language for JSON Document Patches is Now XPath

This section applies to applications that use the Java Client API or REST Client API patch (partial update) feature on JSON documents.

Previously, `JSONPath` was the default path language for identifying the target of update operations in a JSON patch. `XPath` is now the default path language for both XML and JSON patches.

Use of `JSONPath` is now deprecated. To convert your JSON patches to use `XPath` expressions instead of `JSONPath`, see [Traversing JSON Documents Using XPath](#) in the *Application Developer's Guide*.



To continue using JSONPath, you can explicitly override the default path language in one of the following ways:

- For a raw JSON patch, include a `pathlang` property as the sibling of the top level `patch` property. For example:

```
{ "pathlang": "jsonpath",  
  "patch": [ ... ] }
```

Raw patches are used by `POST:/v1/documents` and can be used with the Java Client API method `DocumentManager.patch`.

- When using the Java Client API, use `DocumentPatchBuilder.pathLanguage` to set the path language to JSONPath, as shown in the following example:

```
DocumentPatchBuilder patchBldr = docMgr.newPatchBuilder();  
patchBldr.pathLanguage(PathLanguage.JSONPATH);
```

#### 4.43.1.7 Java and REST: New Restrictions on Patching JSON Content

The native JSON document model imposes some new restrictions on partial updates to JSON documents. Therefore, some patch operations that were previously supported will now be rejected or produce different results. For details, see [Limitations of JSON Path Expressions](#) in the *REST Application Developer's Guide* and [Traversing JSON Documents Using XPath](#) in the *Application Developer's Guide*.

For example, you cannot construct patch path expressions that address anonymous nodes. In the JSON document model, object nodes and array nodes are anonymous. The name in a property name-value pair addresses the value(s), not the containing node.

This means you cannot use “last-child” position to insert a new property or value under the root node of a document or in an array because the parent node is anonymous and cannot be selected by the context expression of the insert operation. Similarly, you cannot address an array node that is nested inside another array (`[1, [2, 3], 4]`) because it is unnamed.

You can no longer replace an entire property (name-value pair) in a single `patch replace` operation for the same reason. To replace a property, you must delete it and then insert a new one.

#### 4.43.1.8 Java and REST: Transforms and Extensions That Manipulate JSON Must Be Rewritten

This section applies to REST Client API and Java Client API applications that use content transformations, resource service extensions, and other server-side code to manipulate the XML representation of JSON documents.

Previously, content transformations, resource service extensions, and other server-side code manipulated JSON content as XML in the `http://marklogic.com/xdmp/json/basic` namespace. Now, such code must operate on JSON document nodes instead of XML.

For example, previously, the JSON data { "key": "value" } was represented in the database as XML of the following form, and this is what your server-side transforms and extensions worked with:

```
<json type="object" xmlns="http://marklogic.com/xdmp/json/basic">
  <key type="string">value</key>
</json>
```

Thus, to access the value of the “key” property in XQuery, you could use a path expression like this following:

```
$someDocument/json:json/json:key
```

With a native JSON document you reference the same data using the following path expression:

```
$someDocument/key
```

You should understand the native JSON document model before rewriting your code. For details, see [Working With JSON](#) in the *Application Developer's Guide*.

#### 4.43.1.9 Java and REST: JSON Array Items and Property Values No Longer Distinguishable in QBE

This change may affect applications that use QBE to search JSON documents using the Java Client API or REST Client API.

Previously, you could construct a QBE that included a word or value query explicitly scoped to an array item or the . Now, it is not possible to distinguish between an array item and a property value.

For example, given a document with the following context:

```
{ "notArray": "blue", "array": ["azure", "blue"] }
```

Previously, the following QBE would only match the occurrence of the value "blue" in the "array" property. Now it matches both the occurrence in "array" and the occurrence in "notArray".

```
{"$query": [ {"$value": ["blue"]} ] }
```

This is because array nodes are unnamed in the native JSON document model, so they cannot be explicitly identified in a QBE.

#### 4.43.1.10 Field Range Query and Field Value Query on JSON May Behave Differently

This difference only applies to applications using field range queries or field value queries on JSON documents.

JSON and XML are not indexed in exactly the same way. Some of the indexing differences affect the behavior of field range queries and field value queries over JSON. Since JSON documents were previously stored as XML, this means your field range queries and field value queries over JSON may behave differently.

For example, previously you could construct a field value query for “John Smith” that would match the following document by defining a field on the `name` property that excluded the `middle` property.

```
{ "name": {  
  "first": "John",  
  "middle": "NMI",  
  "last": "Smith"  
}}
```

This was possible because the document was represented as XML and the text nodes in the field were concatenated together so that the field value in the above document was “John Smith”. In native JSON documents, this concatenation does not occur, and the values of the equivalent field are “John” and “Smith”. To get the same effect now, you would have to use a construct such as a near query.

For more details, see [Creating Indexes and Lexicons Over JSON Documents](#) and [How Field Queries Differ Between JSON and XML](#) in the *Application Developer's Guide*.

## 4.43.2 Semantics Incompatibilities

MarkLogic 8 introduces a number of new and changed Semantic features. This section describes those and includes the following changes that might cause incompatibilities:

- [Changed Function: `sem:sparql`](#)
- [Changed Function: `sem:sparql-values`](#)
- [Changed Function: `sem:sparql-values`](#)
- [Deprecated Function: `sem:sparql-triples`](#)
- [Changed Behavior: Graphs](#)

### 4.43.2.1 Changed Function: `sem:sparql`

In MarkLogic 8, the signature of `sem:sparql` has changed for the last parameters. The fourth parameter is now a `sem:store*`, where previously there were two parameters at the end, one to specify a `cts:query` and another to specify the forest ID. The old signature is still available, but is deprecated. If you have any code from MarkLogic 7 that uses the fourth or fifth arguments to `sem:sparql`, it will still work in MarkLogic 8, but you should migrate that code to use the new signature using `sem:store` as the fourth parameter.

#### 4.43.2.2 Changed Function: `sem:sparql-values`

In MarkLogic 8, `sem:sparql-values` now serializes a string as a `cts:word-query` when used as part of an argument. In other words, string values will be passed as `cts:query` arguments. This is a change from MarkLogic 7.

#### 4.43.2.3 Changed Function: `sem:sparql-values`

The `sem:sparql-values` function no longer accepts `forest-id` as an option. This is an incompatibility with MarkLogic 7 functionality.

#### 4.43.2.4 Deprecated Function: `sem:sparql-triples`

The `sem:sparql-triples` function has been deprecated in favor of `sem:in-memory-store` in MarkLogic 8. See documentation for details - [Querying Triples in Memory](#) in the *Semantics Developer's Guide*.

#### 4.43.2.5 Changed Behavior: Graphs

In MarkLogic 8, graph documents containing metadata are created when triples are ingested, whether they are ingested using SPARQL Update, `mlcp`, or SPARQL endpoints over REST. These named graphs inherit the permissions of the user, unless specified as part of the ingest process. The graph permissions are stored along with other metadata in the graph document.

When loading triples with `mlcp`, if the `output-permissions` parameter is set when loading RDF, the graph will inherit the default permissions just as it would in a `sem:sparql-update` operation. If the `output_collection` parameter is set when loading RDF, the graph document is only created for the first collection specified (because a managed triple can only belong to one graph).

In an upgraded system, graph metadata for managed triples created in MarkLogic 7 will be created the first time you add triples to the graph or modify triples in the graph. The graph will either have the user's permissions or the permissions specified as part of the operation. Make sure the document permissions of documents containing managed triples created in MarkLogic 7 are passed into `sem:sparql-update` as `default-permissions` to ensure the graph metadata created is consistent with the permissions for existing managed triples created in MarkLogic 7.

### 4.43.3 REST and Java Client API Incompatibilities

This section covers the incompatibilities for the REST API between MarkLogic 7 and MarkLogic 8. that are not related to JSON. If you work with JSON documents using the REST Client API or Java Client API, you should also see “JSON Related Incompatibilities” on page 85.

This section covers the following incompatibilities:

- [Must Upgrade to Java Client API v3.0](#)
- [REST API Instance Must Use the Declarative Rewriter on the App Server](#)
- [Default Transaction Mode for the POST Method of Resource Service Extensions is Now Query](#)

- [REST API: Empty Bulk Read by Query Now Returns 200 Status](#)
- [Error Reporting Format and Detail Changes](#)
- [Deprecated Interface: Keyvalue Queries](#)
- [Transaction ID Format Has Changed](#)
- [A Transaction Can No Longer Be Shared Across Users](#)
- [Java: QBE Search Results No Longer Automatically Match the Query Format](#)

#### 4.43.3.1 Must Upgrade to Java Client API v3.0

You cannot use earlier versions of the Java Client API with MarkLogic 8. Update your application to use version 3.0 or later.

#### 4.43.3.2 REST API Instance Must Use the Declarative Rewriter on the App Server

In MarkLogic 7, App Servers that are REST API Instances used a different URL rewriter than in MarkLogic 8. In MarkLogic 8, the App Server is configured to use the declarative rewriter. If you had not modified anything in your REST API Instance App Server setup, the upgrade to MarkLogic 8 will reconfigure your App Server to use the new rewriter. If, however, you have modified something in setup to use a different rewriter, then you will have to make similar changes to the new setup (or consider not using those changes in the REST API Instance). For details on the declarative rewriter, see [xWeb Services](#) in the *Application Developer's Guide*.

#### 4.43.3.3 Default Transaction Mode for the POST Method of Resource Service Extensions is Now Query

In MarkLogic 7, the POST method of a resource service extension was always executed in update mode. In MarkLogic 8, POST methods in a single-statement transaction are executed in query mode. However, within a multi-statement transaction, they are in update mode. The new way is safer because, generally speaking, if a function is not doing an update, it is much more efficient for it to run as a query.

If you have a resource service extension that requires the transaction mode to be update, you need to modify the extension code to add an annotation to the function to explicitly force it into update mode. For example, to modify an extension that is a GET, add an annotation like the following to your get function:

```
declare %rapi:transaction-mode("update") function testrstxn:get (
  $context, $params) {
  <Your code goes here>
};
```

The annotation `%rapi:transaction-mode("update")` forces the function to run as an update. For details, see [Controlling Transaction Mode](#) in the *REST Application Developer's Guide*.

#### 4.43.3.4 REST API: Empty Bulk Read by Query Now Returns 200 Status

Previously, performing a bulk read by retrieving all documents that match a query would return a 404 Not Found response status if no documents matched the query. As of MarkLogic 8, such a request returns a 200 OK status with an empty response body.

This change applies to the following methods, when the Accept header is set to multi-part/mixed and the request does not ask for a search result summary in addition to the matching documents.

- GET: /v1/search
- POST: /v1/search
- GET: /v1/qbe
- POST: /v1/qbe

For more details on these interfaces, see [Reading Multiple Documents Matching a Query](#) in the *REST Application Developer's Guide*.

#### 4.43.3.5 Error Reporting Format and Detail Changes

The changes in this section might affect your application if either of the following is true:

- Your REST or Java client application directly manipulates error details returned by MarkLogic through a REST API instance. This is unlikely for Java applications because they receive such errors as Java exceptions.
- Your application includes content transformations or resource service extensions that report errors to the client.

MarkLogic 8 introduces the following incompatible changes to error reporting for users of the

- [Error Format Defaults to JSON and is a REST Instance Creation Property](#)
- [Error Detail Element and Property Names Have Changed](#)
- [Use RESTAPI-SRVEXERR to Report Errors from Transforms and Extensions](#)

##### **Error Format Defaults to JSON and is a REST Instance Creation Property**

In previous versions, the default format for error messages returned by the REST Client API was XML, and you could change it by setting the `error-format` instance configuration property. As of MarkLogic 8, the default error format for new REST instances is JSON. You can now specify the format when you create the REST instance, and you can subsequently change it using the Admin Interface, `admin:appserver-set-default-error-format`, or the REST Management API.

This change has the following implications:

- The Java `setErrorFormat` and `getErrorFormat` methods of `com.marklogic.client.admin.ServerConfigurationManager` have been removed. Set the error message format when creating the REST instance instead.

- The REST GET and PUT `/v1/config/properties/error-format` methods are no longer available. Set the error message format when creating a REST instance instead.
- You cannot include an `error-format` XML element or JSON property in the payload to `PUT:/v1/config/properties`.
- The payload for `POST:/v1/rest-apis` can now include an `error-format` XML element or JSON property. This is an attribute of the App Server.
- You can use the `Accept` or `X-Error-Accept` HTTP headers to override the default error format for a particular request. For details, see [Error Reporting](#) in the *REST Application Developer's Guide*.

To set the error format when creating an instance, set the `error-format` configuration property. For details, see [Creating an Instance](#) in the *REST Application Developer's Guide*.

### **Error Detail Element and Property Names Have Changed**

The error detail returned by the REST Client API has changed in the following ways. For examples of the new format, see [Error Reporting](#) in the *REST Application Developer's Guide*.

- XML: The root element of the error detail is an `<error-response>` element in the namespace `http://marklogic.com/xdmp/error`. Previously, it was an `<error/>` element in the namespace `http://marklogic.com/rest-api`.
- JSON: The top level property name is now `errorResponse`. Previously, it was `error`. Child property names that previously used dashes to separate “words” now use camel case. For example, `message-code` is now `messageCode` and `status-code` is now `statusCode`.

### **Use RESTAPI-SRVEXERR to Report Errors from Transforms and Extensions**

Resource service extensions and transformations previously reported errors to the client using `RESTAPI-EXTNERR` and could specify a response payload in JSON or XML. The expected error response content type was controlled by a caller supplied parameter. This parameter is now ignored, and you should use `RESTAPI-SRVEXERR` instead of `RESTAPI-EXTNERR`. Your payload must be compatible with the MIME type expected by the caller, which can vary, so it is best to restrict the error detail to text. For details, see [Reporting Errors](#) in the *REST Application Developer's Guide*.

### **4.43.3.6 Deprecated Interface: Keyvalue Queries**

This topic applies to applications that use the Java Client API class `KeyValueQueryDefinition` or the REST Client API method `GET:/v1/keyvalue`. These interfaces are now deprecated.

You can use Query By Example (QBE) or structured query to perform the same kind of search.

For example, to search for a JSON property named “author” with the value “Mark Twain”, use a QBE such as the following:

```
{ "$query": { "author": "Mark Twain" } }
```

The following is a similar search for an XML element:

```
<q:qbe xmlns:q="http://marklogic.com/appservices/querybyexample">
  <q:query>
    <author>Mark Twain</author>
  </q:query>
</q:qbe>
```

With structured query, use `value-query` or `container-query`.

For details, see the following references:

- [Searching Using Query By Example](#) in the *Search Developer's Guide*.
- [Searching Using Structured Queries](#) in the *Search Developer's Guide*.
- `RawQueryByExampleDefinition` or `StructuredQueryBuilder` in the [Java Client API javadoc](#).
- `GET:/v1/search` or `POST:/v1/search` in the *MarkLogic REST API Reference*.

#### 4.43.3.7 Transaction ID Format Has Changed

Previously the transaction ids created using `DatabaseClient.openTransaction` (Java) or `POST:/v1/transactions` (REST) were of the form `hostId_transactionId`. The `hostId` segment has now been dropped.

As long as your application treats the transaction id as a “black box”, this change is transparent.

#### 4.43.3.8 A Transaction Can No Longer Be Shared Across Users

This change only affects Java Client API and REST Client API applications that use multi-statement transactions and share the resulting transaction id across multiple users.

Previously, it was possible to create a multi-statement transaction as one MarkLogic user and then perform operations within the transaction as another user. This is no longer possible. Now, all operations within a transaction must be performed as the same user who created the transaction.

#### 4.43.3.9 Java: QBE Search Results No Longer Automatically Match the Query Format

Previously, using `QueryManager.search` with a Query By Example (QBE) automatically returned results in the same format as the query. That is, XML results were returned for an XML QBE, and JSON results were returned for a JSON QBE. Now, you must explicitly request JSON.

For example, if the `qbe` variable in the following statement contains a JSON QBE, then previously you would receive JSON results. Now, you will receive XML by default, instead.

```
queryMgr.search(qbe.newStringHandle()).get();
```

To achieve the same result as before, explicitly set the format on the result handle to JSON. For example:



```
queryMgr.search(qbe.newStringHandle().withFormat(Format.JSON)).get();
```

#### 4.43.4 Document Library Services (DLS) Repositories Need To Perform A Bulk Upgrade Operation

MarkLogic 8 includes an enhancement to Document Library Services to make it significantly faster for large DLS repositories. This enhancement requires some metadata changes to the documents under DLS control.

If you have any DLS repositories created in MarkLogic 7 or earlier, you must first set compatibility mode for your repository, and then upgrade the documents in the repository, and finally set the repository to upgraded. The upgrade process will touch all of the documents under DLS control, so it will take a while, depending on the size of your DLS application. If you do not perform at least the first part of this upgrade, DLS functions might produce incorrect results in MarkLogic 8.

Because this is an upgrade that touches a large number of documents, MarkLogic strongly recommends that you first back up your database and that you thoroughly test your process on a development system before upgrading your production DLS repository.

To upgrade existing DLS repositories, perform the following steps:

1. Back up your database containing the documents under DLS control.
2. As either a user with the `admin` role, set compatibility mode for your DLS repository by running the following XQuery against your DLS database (for example, in Query Console):

```
xquery version "1.0-ml";

import module namespace dls = "http://marklogic.com/xdmp/dls"
  at "/MarkLogic/dls.xqy";

dls:set-upgrade-status(fn:false())
```

3. As a user with the `admin` role, run the following XQuery against your DLS database (for example, in Query Console):

```
xquery version "1.0-ml";
(: This starts a task which will run for a time proportional to
   the number of documents you have under DLS control. The
   function returns immediately though. It is safe to rerun
   this function if it is stopped or fails for any reason
   such as a system restart. :)

import module namespace dls = "http://marklogic.com/xdmp/dls"
  at "/MarkLogic/dls.xqy";

dls:start-upgrade()
```

If you stop the upgrade (for example, if the server is restarted or if there are errors in the upgrade that you have fixed), you can restart the upgrade at any time by running the above query.

4. You can check the progress at any time by looking at the `upgrade-task-status.xml` document, ad in the following XQuery:

```
xquery version "1.0-ml";
(:
   this document is updated every few minutes to show the
   progress of the upgrade
   :)
fn:doc("http://marklogic.com/dls/upgrade-task-status.xml")
```

5. You can check the progress at any time by running the following XQuery:

```
xquery version "1.0-ml";

import module namespace dls = "http://marklogic.com/xdmp/dls"
  at "/MarkLogic/dls.xqy";

dls:latest-validation-results()
```

6. The `dls:latest-validation-results` output has an element names `dls:validation-status`. When the value of that element is completed, the process is complete.

7. When you are satisfied that the process has completed (for example, if the previous step shows the process is complete), the set the upgrade status to true by running the following, either as a user with the `admin` role or a user with the `dls-admin` role:

```
xquery version "1.0-ml";

import module namespace dls = "http://marklogic.com/xdmp/dls"
  at "/MarkLogic/dls.xqy";

dls:set-upgrade-status(fn:true())
```

Once the process is successful, you can use DLS as usual.

It is also possible to run DLS in compatibility mode without running the upgrade on the repository (by not running the upgrade portion of the above procedure), but MarkLogic strongly recommends performing this upgrade procedure. If you have any problems or questions and you have an active maintenance contract, you can contact MarkLogic Technical Support.

#### 4.43.5 Linux Now Requires Red Hat 6

MarkLogic 8 Linux platforms now require Red Hat 6, and will no longer work on Red Hat 5; they will fail to install on Red Hat 5. So if you are running MarkLogic 7 on Red Hat 5, you will have to migrate that environment to Red Hat 6 in order to use MarkLogic 8.

#### 4.43.6 mlsq1 On Linux No Longer Ships With Server

On Linux platforms, the `mlsq1` utility no longer is packaged with the MarkLogic Server rpm binary. To get `mlsq1` on Linux now, you must install the ODBC Driver for Linux. For details on the ODBC Driver for Linux, see [Configuring the ODBC Driver on Linux](#) in the *SQL Data Modeling Guide*.

#### 4.43.7 Cyrillic Tokenization Changes

The tokenization rules for Cyrillic script has changed such that mixed letter and number tokens are handled consistently with Latin letter and number tokens. That is, the following now tokenizes as a single token (previously it was two tokens):

```
&#x401;1 (Cyrillic A + the digit 1)
```

If you have Cyrillic content in your database, you should reindex the database or reload the Cyrillic content so that it is properly tokenized and indexed in accordance with the new rules. If you do not trigger a retokenization of any existing Cyrillic content, certain queries may behave inconsistently between old content and newly loaded content.

### 4.43.8 Application Builder Applications Must Be Re-Deployed in MarkLogic 8

Applications deployed with Application Builder in MarkLogic 7 and earlier will not work correctly in MarkLogic 8 until you fully re-deploy them.

If you have any applications built using Application Builder, you must remove the code from the modules database of the deployed application and then re-deploy the application. Specifically, for each Application Builder application that is upgraded from MarkLogic 7 or earlier:

1. Back up the application in case you need to restore it to its previous state.
2. In the modules database for the application, either clear the database or delete the following directories:

```
/application  
/Default  
/marklogic.rest.resource  
/marklogic.rest.transform
```

As well as the document at the following URI

```
/index.html
```

3. Go into Application Builder and re-deploy the application.
4. If you have extended your application with any customizations, re-deploy those customizations.
5. Test your re-deployed application.

### 4.43.9 Application Builder and Information Studio Links Removed

The links in the navigation bar from Query Console and other tools no longer have links to Application Builder and Information Studio. To use these applications, enter the URL directly (for example, <http://localhost:8000/appservices>).

### 4.43.10 Search API Incompatibilities

The following incompatible changes have been made to the Search API:

- [search:parse Output is Now Unannotated cts:query XML](#)
- [Deprecated Option: extract-metadata](#)
- [Deprecated Functions: search:unparse, search:remove-constraint](#)
- [Structured Query: locks-query and properties-query Renamed](#)
- [sort-order Query Option Requires an Index](#)

#### 4.43.10.1 `search:parse` Output is Now Unannotated `cts:query` XML

Previously, `search:parse` produced XML representing an annotated `cts:query` that could be passed directly to `search:unparse`. As of MarkLogic 9, the default output from `search:parse` does not include annotations, so it cannot be passed to `search:unparse`.

You can get annotated output from `search:parse` in MarkLogic 9 using the `$output` parameter as shown below:

```
search:parse("myQueryText", options, "cts:annotated-query")
```

#### 4.43.10.2 Deprecating Option: `extract-metadata`

The `extract-metadata` query option is now deprecated. Use `extract-document-data` instead. For details, see `search:search` or [Extracting a Portion of Matching Documents](#) in the *Search Developer's Guide*.

The new `extract-document-data` option does not support extracting specific metadata properties, but properties are available in other ways, such as using `xdmp:document-properties` or by fetching all properties metadata through one of the client APIs.

#### 4.43.10.3 Deprecating Functions: `search:unparse`, `search:remove-constraint`

The `search:unparse` and `search:remove-constraint` functions are now deprecated. If you need to deconstruct a query, modify it, and “put it back together”, use structured query or `cts:query` with `search:resolve` instead.

For example, if you previously did something similar to the following:

```
let $ctsquery := search:parse("myQueryString", $options)
(: ...modify $ctsquery... :)
return search:search(search:unparse($ctsquery), $options)
```

Then you can achieve the same result doing the following:

```
let $ctsquery := search:parse("myQueryString", $options)
(: ...modify $ctsquery... :)
return search:resolve($ctsquery, $options)
```

To generate a structured query instead of a `cts:query`, set the third parameter of `search:parse` to `"search:query"`:

```
search:parse("myQueryString", $options, "search:query")
```

#### 4.43.10.4 Structured Query: `locks-query` and `properties-query` Renamed

The following structured query elements have been renamed to more accurately reflect their purpose:

- `locks-query` is now `locks-fragment-query`
- `properties-query` is now `properties-fragment-query`

For details, see [locks-fragment-query](#) and [properties-fragment-query](#) in the *Search Developer's Guide*.

#### 4.43.10.5 `sort-order` Query Option Requires an Index

If you use the `sort-order` query option to sort search results by something other than score, such as an XML element, JSON property, or field, then the database configuration must include a range index on the entity used for a sort key. Failing to create such an index causes `SEARCH-BADORDERBY` to be thrown when searching. Previously, the index requirement was not enforced.

For details, see [sort-order](#) in the *Search Developer's Guide*.

#### 4.43.11 Locks and Properties Query Built-In Functions Renamed

The following built-in functions related to locks and document properties queries have been renamed to more accurately reflect their purpose.

- `cts:locks-query` is now `cts:locks-fragment-query`
- `cts:locks-query-query` is now `cts:locks-fragment-query-query`
- `cts:properties-query` is now `cts:properties-fragment-query`
- `cts:properties-query-query` is now `cts:properties-fragment-query-query`

#### 4.43.12 `xdmp:uri-content-type` Of an XML Document Now Returns `application/xml`, Can Affect CPF Applications

In 8.0, the `xdmp:uri-content-type` function returns `application/xml`. In 7.0, it returns `text/xml`. Where MarkLogic previously returned `text/xml` for an xml document, it now returns `application/xml`. If you have applications that are expecting `text/xml`, you will either need to change the application to accept `application/xml` or you will have to modify your program to send a content type of `text/xml` (for example, using `xdmp:set-response-content-type`). MarkLogic will still accept `text/xml` for xml documents (in addition to `application/xml`). Similarly, JSON documents return `application/json` but MarkLogic accepts either `application/json` or `text/json`.

As a consequence of this change, if you have a CPF application that relies on the mimetype `text/xml` to identify an XML document, you must change that application to instead rely on `application/xml`. In the case of a CPF application, you will have to modify your CPF pipelines and change any `text/xml` mimetypes (that were referring to XML documents) to `application/xml`. If you are using the default pipelines, then reinstalling CPF for the database should correct this incompatibility. Otherwise, those CPF applications will not trigger the change actions for XML documents. Similar changes are required for CPF applications that rely on `text/json` to identify JSON documents; they need the mimetype changed to `application/json`.

### 4.43.13 `xdmp:function` Signature Change

To allow for anonymous functions in Server-Side JavaScript, the first argument of the `xdmp:function` built-in function takes an `xs:QName?` (zero or 1 QNames) in 8.0; previously, it took an `xs:QName` (exactly one QName). If you have code that does not allow for zero or 1 QNames, you will need to modify that code.

Also, if you are relying on function mapping with `xdmp:function`, you can no longer use function mapping with it (because it no longer takes a singleton). If you have code that function maps with `xdmp:function`, you will need to rewrite it to not use function mapping (by calling `xdmp:function` in the return of a FLOWR statement for each item in your sequence of QNames, for example).

### 4.43.14 Incompatibilities Between 8.0-5 and 8.0-6

The following incompatibilities exist between MarkLogic 8.0-5 and MarkLogic 8.0-6:

- [Terms Matched by additional-query Are Highlighted in Snippets](#)

#### 4.43.14.1 Terms Matched by additional-query Are Highlighted in Snippets

Previously, the Search API documentation stated that terms matched by the query specified in an `additional-query` query option were not highlighted in search result snippets. This is no longer the case.

You should expect any terms matched by the `additional-query` option to be included in highlighted sections of snippets.

### 4.43.15 Incompatibilities Between 8.0-3 and 8.0-4

There are a few incompatibilities made to the Server-Side JavaScript implementation in 8.0-3. The following are the incompatibilities:

- [xdmp.multipartDecode Now Returns a JSON Payload for Headers](#)
- [In JavaScript, Some Thesaurus and Spelling Function Have Different Return Type](#)
- [xdmp.databaseRestoreStatus Now Returns an Object](#)
- [Serialization Error Code Changes](#)
- [Change to Required Java Version](#)
- [Deprecated mlcp Command Line Options](#)
- [REST APIs That Have JSON or XML Payloads Cannot Have Empty Payloads](#)
- [Geospatial Namespace and Data Version Changes](#)

#### 4.43.15.1 `xdmp.multipartDecode` Now Returns a JSON Payload for Headers

In 8.0-4, the Server-Side JavaScript `xdmp.multipartDecode` function returns the headers (in the first item of the returned `ValueIterator`) as a JSON array; previously, it was returned as an XML element. If you have code that is expecting the XML element, you need to either rewrite your code to parse the JSON array or use the XQuery version (`xdmp:multipart-decode`).

#### 4.43.15.2 In JavaScript, Some Thesaurus and Spelling Function Have Different Return Type

In 8.0-4, there are changes to the thesaurus and spelling function modules to make them more friendly to JavaScript users. If you have JavaScript code that imports these XQuery libraries, then the following functions now return JavaScript Object by default:

- `spell.makeDictionary`
- `thsr.lookup`
- `thsr.queryLookup`

For each of these function, you can set a new optional parameter to change its output. By default, these functions return XML output in XQuery and JavaScript Objects in JavaScript. If you have existing JavaScript code that uses these functions, you either need to add the new option to your code specifying the XML output or rework your code to accept the returned JavaScript Object. For details on these functions, see the API documentation for each function. In XQuery, the functions behave the same way they did in previous versions by default, but now allow you to specify the optional parameter to output JavaScript Objects instead of XML, if you so choose.

#### 4.43.15.3 `xdmp.databaseRestoreStatus` Now Returns an Object

In 8.0-4, the Server-Side JavaScript API `xdmp.databaseRestoreStatus` now returns an `Object`; previously, it returned an `Array`. The `Array` that was previously returned is now the value of the `"forest"` key, and there is also a `"status"` key containing the current status of the restore. The `xdmp.databaseBackupStatus` also contains this new `"status"` key, but the return type is not changed. Similarly, the XQuery counterparts to these APIs (`xdmp:database-backup-status` and `xdmp:database-restore-status`) also contain information about the status, but their signature has not changed. If you have code that relies on any of the old behavior, you will need to modify that code to work with the changed output.



#### 4.43.15.4 Serialization Error Code Changes

In 8.0-4, the names of some error exception codes for serialization, as well as the message text, have changed as shown in the following table:

Old Error Code	New Error Code
SER-DOCTYPESYSTEM2	SER-DOCTYPESYSTOPLVTEXT
SER-DOCTYPESYSTEM3	SER-DOCTYPESYSMULTIELEMROOT
SER-STANDALONE2	SER-STANDALONETOPLVTEXT
SER-STANDALONE3	SER-STANDALONEMULTIELEMROOT
SER-STANDALONE4	SER-STANDALONEOMITXMLDEC

If you have code that does a try/catch looking for one of the old exceptions or the old text, or if you have tests that use the old exceptions as keys, you will have to rewrite that code to look for the new exception.

#### 4.43.15.5 Change to Required Java Version

The following tools and libraries that depend on a Java Runtime Environment (JRE) now require at least Java 7, rather than Java 6:

- mlcp
- MarkLogic Connector for Hadoop
- Java Client API
- XCC for Java (XCC/J)

#### 4.43.15.6 Deprecated mlcp Command Line Options

The `-aggregate_uri_id` and `-delimited_uri_id` command line options are now deprecated. Use the more general `-uri_id` instead.

#### 4.43.15.7 REST APIs That Have JSON or XML Payloads Cannot Have Empty Payloads

Starting in 8.0-4, any of the REST APIs that specifies a JSON or and XML `content-type` for its payload cannot have an empty payload. Previously, it allowed an empty payload. REST calls that specify a JSON or XML `content-type` with an empty payload throw a `MANAGE_EMPTYPAYLOAD` exception beginning in 8.0-4. For example, `POST:/admin/v1/init` previously allowed an empty payload with a JSON or XML content, but requires the payload in 8.0-4.

If you have code that does not send a payload, then you must either add an empty JSON or XML document as the payload or change the `content-type` to one that allows an empty payload (for example, `text/plain`).

#### 4.43.15.8 Geospatial Namespace and Data Version Changes

The following changes have been made to some of the geospatial built-in and library functions in 8.0-4:

- [GML and KML Library Modules Moved to a New Namespace](#)
- [Some Built-In Geospatial Functions Moved to geo Namespace](#)
- [Older GML and KML Versions Deprecated](#)

##### **GML and KML Library Modules Moved to a New Namespace**

The GML library module is now in a different namespace, and the library module no longer uses the same namespace as GML and KML data.

Module	Old Module Namespace	New Module Namespace
GML	<a href="http://www.opengis.net/gml">http://www.opengis.net/gml</a>	<a href="http://marklogic.com/geospatial/gml">http://marklogic.com/geospatial/gml</a>
KML	<a href="http://earth.google.com/kml/2.0">http://earth.google.com/kml/2.0</a>	<a href="http://marklogic.com/geospatial/kml">http://marklogic.com/geospatial/kml</a>

You must update your module import declarations in XQuery or require statements JavaScript to use the new namespace. In addition, since the module and the XML data of the same format no longer use the same namespace, you may need to change the namespace prefix you use for the module.

The following code snippets demonstrate the required changes for KML and GML in XQuery. The namespace URI in the module import declaration is changed to use the new URI, and the module namespace prefix is changed to distinguish names in the module from names in the data.

Old	New
<pre>xquery version "1.0-ml"; import module namespace kml =   "http://earth.google.com/kml/2.0" at   "/MarkLogic/geospatial/kml.xqy";  kml:box(   &lt;kml:LatLongBox&gt;...&lt;/kml:LatLongBox&gt;)</pre>	<pre>xquery version "1.0-ml"; import module namespace <b>geokml</b> =   "http://marklogic.com/geospatial/kml"   at "/MarkLogic/geospatial/kml.xqy"; declare namespace kml =   "http://www.opengis.net/kml/2.2";  geokml:box(   &lt;kml:LatLongBox&gt;...&lt;/kml:LatLongBox&gt;)</pre>
<pre>xquery version "1.0-ml"; import module namespace gml =   "http://www.opengis.net/gml"   at "/MarkLogic/geospatial/gml.xqy";  gml:box(   &lt;gml:Envelope&gt;...&lt;/gml:Envelope&gt;)</pre>	<pre>xquery version "1.0-ml"; import module namespace <b>geogml</b> =   "http://marklogic.com/geospatial/gml"   at "/MarkLogic/geospatial/gml.xqy"; declare namespace gml =   http://www.opengis.net/gml/3.2";  geogml:box(   &lt;gml:Envelope&gt;...&lt;/gml:Envelope&gt;)</pre>

**Some Built-In Geospatial Functions Moved to geo Namespace**

The following geospatial built-in functions are now deprecated and will be removed in a future release. You should use the corresponding built-in function with a “geo” prefix instead. For example, use `geo:distance` instead of `cts:distance` in XQuery, and use `geo.distance` instead of `cts.distance` in Server-Side JavaScript.

XQuery	Server-Side JavaScript
<code>cts:approx-center</code>	<code>cts.approxCenter</code>
<code>cts:arc-intersection</code>	<code>cts.arcIntersection</code>
<code>cts:bearing</code>	<code>cts.bearing</code>
<code>cts:box-intersects</code>	<code>cts.boxIntersects</code>
<code>cts:circle-intersects</code>	<code>cts.circleIntersects</code>
<code>cts:complex-polygon-contains</code>	<code>cts.complexPolygonContains</code>
<code>cts:complex-polygon-intersects</code>	<code>cts.complexPolygonIntersects</code>
<code>cts:destination</code>	<code>cts.destination</code>
<code>cts:distance</code>	<code>cts.distance</code>

XQuery	Server-Side JavaScript
<code>cts:parse-wkt</code>	<code>cts.parseWkt</code>
<code>cts:polygon-contains</code>	<code>cts.polygonContains</code>
<code>cts:polygon-intersects</code>	<code>cts.polygonIntersects</code>
<code>cts:region-contains</code>	<code>cts.regionContains</code>
<code>cts:region-intersects</code>	<code>cts.regionIntersects</code>
<code>cts:shortest-distance</code>	<code>cts.shortestDistance</code>
<code>cts:to-wkt</code>	<code>cts.toWkt</code>

### **Older GML and KML Versions Depreciated**

As of MarkLogic 8.0-4, the default KML data version is 2.2 and the default GML data version is 3.2. Older versions are deprecated. You should convert your data.

To continue querying older versions of data with `geokml:geospatial-query` or `geogml:geospatial-query`, specify the namespace URI of the older version in your query constructor. For example:

```
GML
Old: gml:geospatial-query($regions, $options, $weight)
New: geogml:geospatial-query(
    $regions, $options, $weight, "http://www.opengis.net/gml")

KML
Old: kml:geospatial-query($regions, $options, $weight)
New: geokml:geospatial-query(
    $regions, $options, $weight, "http://earth.google.com/kml/2.0")
```

### **4.43.16 Incompatibilities Between 8.0-2 and 8.0-3**

There are a few incompatibilities made to the Server-Side JavaScript implementation in 8.0-3. The following are the incompatibilities:

- [spell.suggestDetailed, xdmp.filesystemDirectory, and xdmp.encodingLanguageDetect Now Return ValueIterator](#)
- [xdmp.databaseRestoreStatus Now Returns an Array](#)
- [xdmp.gssServerNegotiate Now Returns a JavaScript Object](#)
- [Use of String Transaction Ids in Node.js To Be Deprecated](#)
- [CDH 4.3 is No Longer a Supported Hadoop Distribution](#)
- [Changes to How MarkLogic Locates Java and Hadoop Libraries for HDFS Forest Storage](#)

#### 4.43.16.1 `spell.suggestDetailed`, `xdmp.filesystemDirectory`, and `xdmp.encodingLanguageDetect` Now Return `ValueIterator`

In 8.0-3, the `spell.suggestDetailed`, `xdmp.filesystemDirectory`, and `xdmp.encodingLanguageDetect` functions now return results in a `ValueIterator`. In 8.0-2 and earlier, they returned results in an `Array`. If you have any code that expects an `Array`, you must rework that code to accept a `ValueIterator`; for example, you can use `xdmp.arrayValues` to convert the `ValueIterator` result into an `Array` result. For details on the `ValueIterator` JavaScript type, see [ValueIterator](#) in the *JavaScript Reference Guide*.

#### 4.43.16.2 `xdmp.databaseRestoreStatus` Now Returns an `Array`

In 8.0-3, the `xdmp.databaseRestoreStatus` function now returns a JavaScript `Array`. Previously, it returned an `ArrayNode`. If you have any code that expects the `ArrayNode`, you must rework that code to accept an `Array`.

#### 4.43.16.3 `xdmp.gssServerNegotiate` Now Returns a JavaScript Object

In 8.0-3, the `xdmp.gssServerNegotiate` function (used for kerberos GSS authentication) returns a JavaScript object. Previously, it returned an XML element.

#### 4.43.16.4 Use of String Transaction Ids in Node.js To Be Deprecated

The Node.js interfaces that open, manipulate, or pass around transaction ids now accept either a simple id (as before) or a transaction object. You obtain a transaction object rather than an id by passing `true` in for the new `withState` parameter of `DatabaseClient.transactions.open`. For example:

```
// old forms: return a transaction id, for backward compatibility
db.transactions.open();
db.transactions.open({timeLimit: limit, transactionName: name});

// new forms: return transaction object; preferred.
db.transactions.open(true);
db.transactions.open({withState: true, ...});
```

In a future release, the default of `withState` will be changed to `true` and the use of transaction ids will be deprecated.

If you do not modify your code to use transaction objects rather than ids, the session affinity required by multi-statement transactions is not guaranteed to be properly preserved in all cases.

For details, see [Managing Transactions](#) in the *Node.js Application Developer's Guide*.

#### 4.43.16.5 CDH 4.3 is No Longer a Supported Hadoop Distribution

The MarkLogic features and tools that rely on Hadoop no longer support CDH 4.3. Instead, use one of the newer supported versions of Hadoop.

This change affects MarkLogic Content Pump (mlcp), the MarkLogic Connector for Hadoop, and forest storage on HDFS.

#### 4.43.16 Changes to How MarkLogic Locates Java and Hadoop Libraries for HDFS Forest Storage

When you use HDFS for forest storage, MarkLogic must be able to find a suitable Java installation and Hadoop HDFS libraries. The algorithm for where MarkLogic looks has changed substantially as of 8.0-3.

You should now make the Hadoop libraries available to your MarkLogic hosts using one of the new Hadoop HDFS client bundles. You must unpack one of these bundles under `/opt`, `/usr`, or `/space`. For details, see [HDFS Storage](#) in the *Query Performance and Tuning Guide*.

MarkLogic now locates a suitable Java installation using a different algorithm. You might need to either change the path to your JDK or set `JAVA_HOME` in the MarkLogic startup environment. You can now make `JAVA_HOME` available to MarkLogic in a way that is preserved across upgrades, using `/etc/marklogic.conf`. For details, see [HDFS Storage](#) in the *Query Performance and Tuning Guide*.

#### 4.43.17 Incompatibilities Between 8.0-1 and 8.0-2

There are a few incompatible changes made to the Server-Side JavaScript implementation in 8.0-2. The following are the incompatibilities:

- [Array Input Differences in `fn.distinctValues`, `fn.subsequence`, and Other Functions](#)
- [The Second Parameters of `xdmp.eval`, `xdmp.invoke`, `xdmp.xqueryEval`, and `xdmp.spawn` Now Take a Single Object](#)
- [extract-document-data Results Now Inline By Default](#)
- [XCC v8.0-2 May Require Config Change When Used with Older Versions of MarkLogic](#)
- [Client APIs: JavaScript Extension and Transform Error Reporting Convention Change](#)
- [Some JavaScript Built-In Functions that Returned XML Structures Now Return JSON Structures](#)

##### 4.43.17.1 Array Input Differences in `fn.distinctValues`, `fn.subsequence`, and Other Functions

In 8.0-2, the Server-Side JavaScript functions `fn.distinctValues` and `fn.subsequence` behave differently from 8.0-1 if you pass in an array. In 8.0-1, these functions will extract the values out of the array to make multiple inputs, one for each item in the array. In 8.0-2, these functions treat the array as a single item. Therefore, to get the same behavior in 8.0-2, you need to call `xdmp.arrayValues` on the array before you pass it into these functions. For example:

```
fn.distinctValues([1, 1, 2]);  
// returns 1, 2 in 8.0-1  
// returns [1, 1, 2] in 8.0-2
```

To modify the above code in 8.0-2 to return the same answer as in 8.0-1:

```
fn.distinctValues(xdmp.arrayValues([1, 1, 2]));  
// returns 1, 2 in 8.0-2
```

The nature of the change is that there are fewer times when MarkLogic coerces an array into its values in 8.0-2 than there were in 8.0-1. The newer behavior is more natural to JavaScript developers.

In addition to `fn.distinctValues` and `fn.subsequence`, this applies to any JavaScript function that takes a `ValueIterator` as input to a parameter, including: `fn.exactlyOne`, `fn.head`, `fn.insertBefore`, `fn.reverse`, `fn.unordered`, `fn.empty`, `fn.remove`, `fn.reverse`, `fn.zeroOrOne`, `fn.oneOrMore`, `fn.deepEqual`, `fn.count`, `cts.contains`, `xdmp.setSessionField`, `xdmp.setServerField`, and others.

#### 4.43.17.2 The Second Parameters of `xdmp.eval`, `xdmp.invoke`, `xdmp.xqueryEval`, and `xdmp.spawn` Now Take a Single Object

In 8.0-2, the `vars` parameter to the Server-Side JavaScript functions `xdmp.eval`, `xdmp.invoke`, `xdmp.xqueryEval`, and `xdmp.spawn` has been simplified so it only takes a single Object. If you have any code that you used in 8.0-1 that passes the external variables (`vars` parameter) as an array of Objects, as an array of strings, or as a `ValueIterator`, you must rewrite that code in 8.0-2 so it passes a single Object.

#### 4.43.17.3 `extract-document-data` Results Now Inline By Default

The query option `extract-document-data` previously caused `search:search` and `search:resolve` to return a sequence consisting of the `search:response` and the extracted documents. Now, this option returns the extracted documents embedded in the `search:response` as `search:extracted` elements.

When you use the option with the REST Client API, the `/v1/search` service returns the extracted content inside the search response if the Accept header MIME type is `application/xml` or `application/json`. The extracted content is still returned as individual documents when the Accept header MIME type is `multipart/mixed` (a multi-document read).

For details, see [Extracting a Portion of Matching Documents](#) in the *Search Developer's Guide*.

#### 4.43.17.4 XCC v8.0-2 May Require Config Change When Used with Older Versions of MarkLogic

This change affects XCC applications that set the transaction mode (`Session.setTransactionMode`) or transaction time limit (`Session.setTransactionTimeout`) and use XCC v8.0-2 or later with MarkLogic Server 8.0-1 or earlier.

If your XCC application meets the above conditions, you must set the property `xcc.txn.compatible` to `true`. If you do not do so, an exception is raised if you set the transaction time limit or set the transaction mode to a value other than `Session.TransactionMode.AUTO`.

You can set this system property on the java command line with an argument of the following form:

```
java -Dxcc.txn.compatible=true
```

You can also set the property programmatically by calling `System.setProperty`.

#### 4.43.17.5 Client APIs: JavaScript Extension and Transform Error Reporting Convention Change

The following change affects applications that use the REST Client API, Java Client API, or Node.js Client API and that raise `RESTAPI-SRVEXERR` from a server-side JavaScript extension, transform, or custom snippet, or other custom code.

In MarkLogic Server 8.0-1, when calling `fn.error` to raise `RESTAPI-SRVEXERR`, error detail such as the response status code and status text are passed to `fn.error` as an array. Starting with MarkLogic Server 8.0-2, the error details must be passed as a sequence. You can use `xdmp.arrayValues` to convert the array to a sequence.

For example, if you previously had an `fn.error` call in an extension similar to the following:

```
fn.error(null, 'RESTAPI-SRVEXERR',
         [statusCode, statusMsg, body])
```

Then you should now wrap the array that is the 3rd argument to `fn.error` in a call to `xdmp.arrayValues`, similar to the following:

```
fn.error(null, 'RESTAPI-SRVEXERR',
         xdmp.arrayValues([statusCode, statusMsg, body]))
```

#### 4.43.17.6 Some JavaScript Built-In Functions that Returned XML Structures Now Return JSON Structures

The following APIs return JSON output in 8.0-2.

- `cts.plan`
- `cts.relevanceInfo`
- `xdmp.zipManifest`
- `xdmp.userExternalSecurity`
- `xdmp.userLastLogin`
- `xdmp.databasePathNamespaces`
- `cts.distinctiveTerms`
- `cts.cluster`
- `cts.train`
- `cts.classify`
- `cts.thresholds`



In 8.0-1, these APIs returned the same XML structures that their XQuery counterparts return. If you have any JavaScript code that relies on the XML structures, you will need to modify that code to use the new JSON output.

## 4.44 MarkLogic 7 Incompatibilities

MarkLogic 9 allows you to upgrade either from MarkLogic 5, MarkLogic 6, or MarkLogic 7. If you are upgrading from 4.2, you must first upgrade to at least MarkLogic 5, and there are some known incompatibilities between 4.2 and 5.0 that are documented in the 5.0 *Release Notes*. If you are upgrading from MarkLogic 7, you can skip this section. For convenience, the incompatibilities between MarkLogic 6 and MarkLogic 7 are repeated here, and are as follows:

- [Incompatibilities Between MarkLogic 7.0-3 and 7.0-2](#)
- [Incompatibilities Between MarkLogic 7.0-2 and 7.0-1](#)
- [Incompatibilities Between MarkLogic 7.0-1 and MarkLogic 6](#)

### 4.44.1 Incompatibilities Between MarkLogic 7.0-3 and 7.0-2

This section describes the incompatibilities between MarkLogic 7.0-3 and 7.0-2.

- [HDP No Longer a Supported Hadoop Platform](#)
- [Java API: ContentVersionRequest Property Deprecated](#)
- [REST API: content-versions Property Deprecated](#)
- [REST API: JSON documents Cannot be Retrieved as XML](#)
- [Float Precision Greater in 7.0-3](#)

#### 4.44.1.1 HDP No Longer a Supported Hadoop Platform

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).

#### 4.44.1.2 Java API: ContentVersionRequest Property Deprecated

The REST server configuration property `ContentVersionRequest` is now deprecated. If you currently use `com.marklogic.admin.ServerConfigurationManager.setContentVersionRequest()` and the related `ServerConfiguration.Policy` type, you should modify your application to use `com.marklogic.admin.ServerConfigurationManager.setUpdatePolicy()` and `ServerConfigurationManager.UpdatePolicy` instead. You should also change calls to `getContentVersionRequest()` into call to `getUpdatePolicy()`.

The table below shows the correspondence between `Policy` values and `UpdatePolicy` values. The behavior is unchanged with respect to these values.

If you used <code>Policy</code> Value...	Then use <code>UpdatePolicy</code> Value...
NONE	MERGE_METADATA
REQUIRED	VERSION_REQUIRED
OPTIONAL	VERSION_OPTIONAL

#### 4.44.1.3 REST API: content-versions Property Deprecated

The REST instance configuration property `content-versions` is now deprecated, in favor of the new `update-policy` property.

If you currently set `content-versions`, you should modify your application to use the `update-policy` configuration property instead. For example, if you set this property using `PUT /v1/config/properties` OR `PUT /v1/config/properties/content-versions`, then you should now use `PUT /v1/config/properties` OR `PUT /v1/config/properties/update-policy` to set `update-policy` instead.

Similarly, you should modify any code that reads the configuration properties to expect `update-policy` instead of `content-versions`.

The table below shows the correspondence between `content-versions` values and `update-policy` values. The behavior is unchanged with respect to these values.

If you used <code>Policy</code> Value...	Then use <code>UpdatePolicy</code> Value...
none	merge-metadata
required	version-required
optional	version-optional

#### 4.44.1.4 REST API: JSON documents Cannot be Retrieved as XML

In previous versions, you could use the REST Client API to retrieve the internal XML representation of a JSON document using `GET /v1/documents` and specifying `XML` in the `format` request parameter or `application/xml` in the `Accept` header.

As of MarkLogic 7.0-3, JSON documents are always returned as JSON. You can still examine the internal representation of a JSON document using XQuery or the Query Console database explorer.

#### 4.44.2 Float Precision Greater in 7.0-3

In MarkLogic 7.0-3, the numeric precision of a float has increased. For example:

```
xs:float(3435.99884)
(: Returns 3436 in 7.0-2,
   returns 3435.999 in 7.0-3 :)
```

In most cases, this will not cause an incompatibility, as it is just returning a more precise number. But if you have application logic that relies on the old behavior, you will have to make changes to account for the greater precision.

#### 4.44.3 Incompatibilities Between MarkLogic 7.0-2 and 7.0-1

This section describes the incompatibilities between MarkLogic 7.0-1 and 7.0-2.

- [Change to JSON Output from the REST API](#)
- [Changes to the MarkLogic Connector for Hadoop API](#)

##### 4.44.3.1 Change to JSON Output from the REST API

The JSON output from the manage REST resource addresses with the `metrics` view has been changed and you may need to change any custom clients that consume this JSON data.

##### 4.44.3.2 Changes to the MarkLogic Connector for Hadoop API

`com.marklogic.mapreduce.MarkLogicDocument` is now an interface instead of a class. The previous functionality of `MarkLogicDocument` is provided by the new class `com.marklogic.mapreduce.DatabaseDocument`.

Modify your code and job configuration to use `DatabaseDocument` instead of `MarkLogicDocument`.

#### 4.44.4 Incompatibilities Between MarkLogic 7.0-1 and MarkLogic 6

- [XQuery HTTP Client Built-In Functions Now Require a Privilege](#)
- [HTTP Client Functions Are Now HTTP 1.1 Compliant](#)
- [xdmp:get-request-username and xdmp:get-request-user Changes](#)
- [Specifying a Forest Now Only Works With Strict Locking](#)
- [Custom Dictionaries for Japanese and Chinese Languages Need to be Re-saved](#)
- [Default Attributes on XML Copy Changes](#)

- [Serialization of Alerting, Reverse, and Path Range Queries Change](#)
- [Java and REST Client API Incompatibilities](#)
- [Namespace Change for Properties Persisted Using JSON](#)
- [mlcp Incompatibilities](#)
- [REST Management API Version Incremented to v2](#)
- [Changes to the Configuration Manager](#)
- [xdmp:plan Now Requires a Privilege](#)
- [fn:analyze-string Now Returns Output in a Different Namespace](#)

#### 4.44.4.1 XQuery HTTP Client Built-In Functions Now Require a Privilege

The HTTP client XQuery APIs (`xdmp:http-delete`, `xdmp:http-get`, `xdmp:http-post`, `xdmp:http-put`, and so on) now require a privilege. Previously, these functions did not require a privilege. To support these privileges, the following privileges are added to MarkLogic 9:

- `http://marklogic.com/xdmp/privileges/xdmp-http-get`
- `http://marklogic.com/xdmp/privileges/xdmp-http-head`
- `http://marklogic.com/xdmp/privileges/xdmp-http-options`
- `http://marklogic.com/xdmp/privileges/xdmp-http-delete`
- `http://marklogic.com/xdmp/privileges/xdmp-http-post`
- `http://marklogic.com/xdmp/privileges/xdmp-http-put`

Additionally, the `network-access` role has been added, which contains all of these privileges.

If you have code that accesses these functions with a user that does not have the `admin` role, you will have to add these privileges (or the `network-access` role) to the set of privileges inherited by your users. For details on adding privileges, see [Security Administration](#) in the *Administrator's Guide*. For details on security, see *Security Guide*.

#### 4.44.4.2 HTTP Client Functions Are Now HTTP 1.1 Compliant

The XQuery HTTP client functions (`xdmp:http-delete`, `xdmp:http-get`, `xdmp:http-post`, `xdmp:http-put`, and so on) now fully implement HTTP 1.1, including the use of connection keep-alives and built-in decoding of chunked transfer encoded responses. In most cases, this will not cause any incompatible behavior, but if your code tried to perform some work to do HTTP 1.1 features such as taking a chunked response as a large binary and decoding the chunks in XQuery, then that code might behave differently in MarkLogic 7. If you have such code, you might need to rework it in MarkLogic 7.

### 4.44.4.3 `xdmp:get-request-username` and `xdmp:get-request-user` Changes

The `xdmp:get-request-username` and `xdmp:get-request-user` functions have changed slightly in MarkLogic 7.

In previous releases, `xdmp:get-request-username` always returned the user in the `Authorization` HTTP header; in MarkLogic 7, if you are using application-level authentication, then it returns the user from the last successful call to `xdmp:login` (when using any other authentication scheme, it is unchanged from previous releases and returns the user in the `Authorization` HTTP header).

In previous releases, `xdmp:get-request-user` returned the ID of the current user; in MarkLogic 7, if you are using application-level authentication, then it returns the user ID from the last successful call to `xdmp:login` (when using any other authentication scheme, it now returns the user ID corresponding to the user in the `Authorization` HTTP header). If you want the old behavior of returning the ID of the current user, use the new function `xdmp:get-current-userid`.

If you have any code that uses the `xdmp:get-request-username` or `xdmp:get-request-user` functions, you might need to change it to work with the current behavior.

### 4.44.4.4 Specifying a Forest Now Only Works With Strict Locking

The various loading APIs (for example, `xdmp:document-load` and `xdmp:document-insert`) have an option to specify the forest in which a document is loaded. In MarkLogic 7, those forest options are only available with the `locking` parameter on the database set to `strict` (the default is `fast`). If you try to specify forest placement with `fast` locking, it will throw an `XDMP-PLACEKEYSLOCKING` exception. Previously, this operation would be allowed.

In most cases, it is not recommended to specify a placekey, as MarkLogic does a good job of distributing data. If you want to continue to use the forest placekeys when loading or updating, you must change your `locking` parameter to `strict` or change your code to no longer use forest placekeys. You can also consider using Tiered Storage to control what data goes into which forests. For details on Tiered Storage, see [Tiered Storage](#) in the *Administrator's Guide*.

### 4.44.4.5 Custom Dictionaries for Japanese and Chinese Languages Need to be Re-saved

In MarkLogic 7, there is a change to the way custom dictionaries are stored for languages that use a unified `cd` (Japanese and both Simplified and Traditional Chinese). The custom dictionaries for those languages now have separate files for tokenization and for stemming. If you are using these languages and you have created a custom dictionary, you must re-save the custom dictionary in MarkLogic 7, which will save it in the new format.

For example, if you have a Japanese language custom dictionary, run the following XQuery program to re-save the custom dictionary:

```
import module namespace
  cdict="http://marklogic.com/xdmp/custom-dictionary"
  at "/MarkLogic/custom-dictionary.xqy";
```

```
cdict:dictionary-save("ja", cdict:dictionary-read("ja"))
```

This will re-save the dictionary to the new format.

#### 4.44.4.6 Default Attributes on XML Copy Changes

In MarkLogic 7, the behavior of default attributes has changed when you are copying XML nodes that have an in-scope schema with default attributes. In MarkLogic 6, the default attributes were applied in the copied node. In MarkLogic 7, the default attributes are still applied in the data model of the copied node, but they will only be serialized if `default-attributes=yes` is set in the serialization options (for example, the `xdmp:output` option). This is true for copied nodes in both XQuery (for example, see below) and XSLT (for example, `xsl:copy`).

The following shows an example using the XHTML schema, which is always in-scope.

```
xquery version "1.0-ml";
declare option xdmp:output "default-attributes=no";
(: the above xdmp:output declaration is the default :)

<x>{xdmp:unquote(' <html xmlns="http://www.w3.org/1999/xhtml">
  <body>
    <p>hello</p>
  </body>
</html>
')}</x>
=>
In MarkLogic 7 returns: (no default attributes on html element)
```

```
<x>
  <html xmlns="http://www.w3.org/1999/xhtml">
    <body>
      <p>hello</p>
    </body>
  </html>
</x>
```

In MarkLogic 6 returns: (added default attribute version on html element)

```
<x>
  <html version="-//W3C//DTD XHTML 1.1//EN"
    xmlns="http://www.w3.org/1999/xhtml">
    <body>
      <p>hello</p>
    </body>
  </html>
</x>
```

If the `xdmp:output` option is set to `default-attributes=yes`, then the attributes will continue to be serialized in MarkLogic 7. Furthermore, if the copied node is used in another schema that has different default values for the default attributes, then the resulting default value for those attributes comes from the copied node, not the new schema. This change is not likely to impact very many applications, and the new behavior is what most people would expect, but if you have code that relies on the old behavior, you will need to explicitly set the `xdmp:output` option is set to `default-attributes=yes`.

#### 4.44.4.7 Serialization of Alerting, Reverse, and Path Range Queries Change

The serialization of path range index queries has changed to use `cts:path-expression` rather than `cts:path` as the element name. This affects not only stored path range index queries, but also stored Alerting queries and reverse queries. All such queries must be transformed to the new serialization and reloaded.

The following XSLT stylesheet makes the required change:

```
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
                xmlns:cts="http://marklogic.com/cts"
                version="2.0">
  <xsl:template match="cts:path">
    <xsl:element name="cts:path-expression">
      <xsl:copy-of select="namespace::*|@*|node()" />
    </xsl:element>
  </xsl:template>
</xsl:stylesheet>
```

You can apply the stylesheet to affected documents using a query similar to the following:

```
xquery version "1.0-ml";
for $query in fn:collection("my-stored-queries")
return xdmp:node-replace($query, xdmp:xslt-eval($stylesheet, $query))
```

#### 4.44.4.8 Java and REST Client API Incompatibilities

Unless otherwise noted, the following changes can affect applications using either the REST or Java Client API:

- [Content Transformations on JSON Documents Operate on XML](#)
- [Resource Service Extensions Moved to Modules Database on Upgrade](#)
- [JSON Key Name Change for System Properties](#)
- [Java Batch Example Package Name Change](#)

#### **Content Transformations on JSON Documents Operate on XML**

This topic applies to REST and Java Client API applications that include custom content transformations for JSON documents.

JSON documents are stored in MarkLogic Server as XML. In MarkLogic 6, content transformation functions were invoked on JSON documents after conversion from JSON to XML. In MarkLogic 7, content transformations are applied before this conversion.

If you created a content transformation for MarkLogic 6 that expects JSON input, you must modify your implementation to expect XML in the `http://marklogic.com/xdmp/json/basic` namespace instead of JSON text. For details on the XML representation, see [Working With JSON](#) in *Application Developer's Guide*.

Transformations invoked during ingestion must similarly be modified to produce XML that conforms to the internal representation for JSON documents. Transformations invoked during document retrieval can either produce conforming XML or JSON text as when generating JSON output.

For details on the new expectations, see [Expected Input and Output](#) in *REST Application Developer's Guide*.

### **Resource Service Extensions Moved to Modules Database on Upgrade**

This topic applies to REST and Java Client API applications running on a REST API instance created with MarkLogic 6.0-1 or later that use resource service extensions.

REST API instances created with MarkLogic Server version 6.0-1 stored resource service extensions in the Extensions database. REST API instances created with MarkLogic Server version 6.0-2 or later install resource service extensions in the modules database associated with the instance.

When you upgrade from MarkLogic 6 to MarkLogic 7 or later, any resource service extensions stored in the Extensions database on behalf of a 6.0-1 REST API instance are automatically migrated from the Extensions database to the modules database for you.

If your resource extension has dependent libraries or other assets that you installed in the Extensions database, you should migrate them by re-installing them using the new `/ext` service. For details, see [Managing Dependent Libraries and Other Assets](#) in *REST Application Developer's Guide*.

### **JSON Key Name Change for System Properties**

This topic applies to REST applications that reference document properties using the JSON representation.

When you retrieve document properties in JSON format, protected system properties such as last-modified are now enclosed in an object with the key `$ml.prop`. In MarkLogic 6, such properties were immediate children of the `properties` container. User-defined property naming is unchanged.

The example output below shows the change in how the last-modified property value is returned by a request of the form `GET /v1/documents?uri=your-uri&category=properties&format=json:`



MarkLogic 6	MarkLogic 7
<pre>{ "properties": {   "last-modified": "value" }}</pre>	<pre>{ "properties": {   "\$ml.prop": {     "last-modified": "value"   } }}</pre>

#### **Java Batch Example Package Name Change**

The Java example application in the package `com.marklogic.client.example.batch` is now in the package `com.marklogic.client.example.extension`. If your application references any classes or interfaces from this package, you must change your package imports and recompile.

#### **4.44.4.9 Namespace Change for Properties Persisted Using JSON**

This topic applies to REST and Java Client API applications that insert or update document properties using JSON and either of the following are true:

- You define indexes based on JSON property keys.
- Your application includes a content transformation, resource extension, or other code that manipulates JSON property metadata in its XML representation.

In MarkLogic 6, inserting or updating user-defined document properties using JSON stored the properties as XML elements in the namespace `http://marklogic.com/json`. In MarkLogic 7, such properties are stored as elements in the namespace `http://marklogic.com/xdmp/json/basic`. If you insert or update a property using JSON using MarkLogic 7, the new namespace is used.

If you convert properties in the old namespace to the new namespace, then all pre-existing and future properties will use the new namespace. If you do not perform such a conversion, then you should adapt your indexes and queries to accommodate both namespaces because pre-existing properties will be in the MarkLogic 6 namespace while new or updated properties will be in the MarkLogic 7 namespace.

After upgrading to MarkLogic 7, use one of the following solutions to adapt your content and application to this change:

- [Modify Pre-Existing Propertie](#)
- [Modify Queries to Use Both Namespaces](#)
- [Create a Field That Spans Both Namespaces](#)

**Modify Pre-Existing Propertie**

Use a query similar to the one in this section to change the namespace of all properties in the old namespace. The example query uses `xdmp:spawn` to perform the property update in batches, thereby avoiding overly large transactions.

The procedure below assumes you have an App Server, such as a REST API instance, associated with your content database. If you are not familiar with Query Console, see *Query Console User Guide*.

The following procedure walks you through installing a transform query in the modules root of an App Server attached to an affected database, and then running the query to modify a specific property.

1. Save the following transform query to a file, such as `update-props.xqy`. You will changed the bolded text in the next step.

```
xquery version "1.0-ml";

declare namespace prop    = "http://marklogic.com/xdmp/property";
declare namespace old-ns = "http://marklogic.com/json";

declare default function namespace
  "http://www.w3.org/2005/xpath-functions";
declare option xdmp:mapping "false";

declare function local:transform(
  $elem as element()
) as element()
{
  let $elem-name :=
    if (exists($elem/self::old-ns:*))
    then QName("http://marklogic.com/xdmp/json/basic",
              local-name($elem))
    else node-name($elem)
  return element {$elem-name} {
    $elem/@*,
    for $child in $elem/node()
    return
      typeswitch ($child)
      case element() return local:transform($child)
      default return $child
  }
};

let $max    := 100
let $batch :=
  subsequence (
    cts:search(collection(),
               cts:properties-query(
                 cts:element-query(xs:QName("old-ns:my-prop"),
                                   cts:and-query(()))
               )
  )
```

```

        )
      ),
      1,
      $max
    )/document-uri(.)
  return
  if (empty($batch)) then ()
  else (
    for $doc-uri in $batch
    let $current-props := xdm:document-properties($doc-uri) /
      prop:properties/* except prop:last-modified
    let $modified-props := $current-props/local:transform(.)
    return xdm:document-set-properties($doc-uri, $modified-props),
    if (count($batch) lt $max) then ()
    else xdm:spawn("/some/path/update-props.xqy")
  )

```

2. Modify the saved query to match your environment by making the following changes:
  - a. Change occurrences of “my-prop” to the name of an affected property in your content.
  - b. Change the module path in the `xdmp:spawn` call to the path where you will install the query in your modules root.
3. Install the saved file in the modules database or modules root of your App Server. Install the module under the path you chose in Step 2b.
  - a. To install the modules database of a REST API instance using the REST API, see [Managing Dependent Libraries and Other Assets](#) in *REST Application Developer’s Guide*.
  - b. To install in the modules database of a REST API instance using Java API, see [Managing Dependent Libraries and Other Assets](#) in *Java Application Developer’s Guide*.
  - c. To install in the modules database using XQuery, run a query similar to the following in Query Console, after modifying the filesystem path and database URI. The file must be accessible to MarkLogic Server. Run the query with the modules database as the content source.

```

xquery version "1.0-m1";
xdmp:document-load(
  "/filesystem/path/update-props.xqy",
  <options xmlns="xdmp:document-load">
    <uri>/some/path/update-props.xqy</uri>
  </options>
)

```

4. Run the transform query using Query Console.

- a. Create a new query in Query Console with the following contents:

```
xquery version "1.0-ml";
xdmp:invoke("/my.domain/update-props.xqy")
```
  - b. Modify the module URI in the `xdmp:spawn` call to the URI under which you installed the module in Step 3.
  - c. In the Query Console Content Source dropdown, select the source that corresponds to your content database and the modules database in which you installed the transform query in Step 3.
  - d. Click the Run button to perform the transformation.
5. If your database is large, you might exceed the maximum number of spawned queries. If this happens, wait for the previous spawns to complete, and then run the query again.

You must also modify any range indexes, queries, or query options for this property that depend on the old namespace. Change occurrences of `http://marklogic.com/json` to `http://marklogic.com/xdmp/json/basic`.

For example, if in MarkLogic 6 you defined an element range index over the property `my-property` in the namespace `http://marklogic.com/json`, modify the your index configuration to use the namespace URI `http://marklogic.com/xdmp/json/basic`.

Similarly, if you have queries or constraint definitions using the old namespace, change them to use the new namespace. The table below contains an example of how to modify a properties constraint to use the new namespace:

Version	Example
MarkLogic 6	<pre data-bbox="553 495 1373 753">&lt;options xmlns="http://marklogic.com/appservices/search"&gt;   &lt;constraint name="my-prop"&gt;     &lt;range type="xs:string"&gt;       &lt;element name="my-property"         ns="http://marklogic.com/json" /&gt;       &lt;fragment-scope&gt;properties&lt;/fragment-scope&gt;     &lt;/range&gt;   &lt;/constraint&gt; &lt;/options&gt;</pre>
MarkLogic 7	<pre data-bbox="553 789 1373 1047">&lt;options xmlns="http://marklogic.com/appservices/search"&gt;   &lt;constraint name="my-prop"&gt;     &lt;range type="xs:string"&gt;       &lt;element name="my-property"         ns="http://marklogic.com/xdmp/json/basic" /&gt;       &lt;fragment-scope&gt;properties&lt;/fragment-scope&gt;     &lt;/range&gt;   &lt;/constraint&gt; &lt;/options&gt;</pre>

### **Modify Queries to Use Both Namespaces**

You can modify your queries and query options to accommodate both namespaces. For example, if your query options already include a constraint on the old namespace, add a constraint for the new namespace, and then use an OR query to apply both constraints.

The following example uses a JSON structured or-query to demonstrate this technique. It assumes the existence of an element range index on the property in the new namespace.

```
{ "search": {
  "options": {
    "constraint": [
      {
        "name": "old-prop",
        "range": {
          "type": "xs:string",
          "element": {
            "name": "my-property",
            "ns": "http://marklogic.com/json"
          },
          "fragment-scope": "properties"
        }
      },
      {
        "name": "new-prop",
```



Old	New
<pre>{ "search": {   "options": {     "constraint": {       "name": "my-prop",       "range": {         "type": "xs:string",         "element": {           "name": "my-property",           "ns": "http://marklogic.com/json"         },         "fragment-scope": "properties"       }     }   } }, "query" : {   "range-constraint-query": {     "constraint-name": "my-prop",     "value": "the value"   } } }</pre>	<pre>{"search": {   "options": {     "constraint": {       "name": "my-prop",       "range": {         "type": "xs:string",         "field": {           "name": "my-new-property"         },         "fragment-scope": "properties"       }     }   },   "query" : {     "range-constraint-query": {       "constraint-name": "my-prop",       "value": "the value"     }   } }</pre>

#### 4.44.4.10mlcp Incompatibilities

MarkLogic Content Pump (mlcp) version 1.1 includes the following changes that potentially affect compatibility for users of mlcp version 1.0-\*

- [Compressed Input Default URI Includes Input Filename](#)
- [Default Character Encoding Changed to UTF-8](#)

##### **Compressed Input Default URI Includes Input Filename**

The following change might affect you if you use mlcp to load documents from compressed files (-input\_compressed). Documents created with mlcp v1.0 used the following default URI template:

```
/path/inside/zip/filename
```

Starting with version 1.1, the default URI template is:

```
/compressed-file-path/path/inside/zip/filename
```

This change can affect the -output\_uri\_replace patterns needed to create documents with the same URIs as those created with mlcp 1.0.

For details, see [Default Document URI Construction](#) in the *mlcp User Guide*.

##### **Default Character Encoding Changed to UTF-8**

The default content character encoding when importing and exporting documents with mlcp has changed to UTF-8. Previously, mlcp used the platform default encoding for the host which mlcp was running.

You can get the previous behavior by using the following option setting:

```
-content-encoding system
```

#### 4.44.4.11 REST Management API Version Incremented to v2

The REST Management API version has been incremented to v2. The v1 services are no longer available. If you send a request using a v1 URL, MarkLogic Server responds with status code 410 (Gone) and a `MANAGE-UNSUPPORTEDVERSION` error.

Requests that use `LATEST` as the version when constructing requests will continue to work. However, you may need to make other changes due to the behavior changes between v1 and v2.

The incompatibilities between v1 and v2 are detailed in the remainder of this section

- [View Parameter Required Instead of Path Steps](#)
- [JSON Output Includes Units](#)
- [Element/Key Name Changes in Status Views](#)
- [Changes to the Management API Plugins](#)
- [Changes to the Packaging API](#)

##### **View Parameter Required Instead of Path Steps**

In previous releases, there were two ways to access some views: by path step or using the `view` request parameter. In MarkLogic 7, the path step form of URL has been removed. You must now use the `view` parameter. For example, MarkLogic 6 supported the following two ways of requesting database status, where `version` is `v1` or `LATEST`:

```
GET /manage/version/databases/{id/name}/status
GET /manage/version/databases/{id/name?view=status
```

In MarkLogic 7, only the second form is supported, as shown in the example below, where `version` is `v2` or `LATEST`.

```
GET /manage/version/databases/{id/name?view=status
```

This change applies to the `config`, `counts`, `edit`, and `status` views for clusters, databases, forests, groups, hosts, and servers. The table below lists the affected GET methods and the equivalent MarkLogic 7 URL.



Previous Form	MarkLogic 7 Form
/manage/v1/clusters/{id name}/config	/manage/v2/clusters/{id name}?view=config
/manage/v1/clusters/{id name}/status	/manage/v2/clusters/{id name}?view=status
/manage/v1/databases/{id name}/config	/manage/v2/databases/{id name}?view=config
/manage/v1/databases/{id name}/counts	/manage/v2/databases/{id name}?view=counts
/manage/v1/databases/{id name}/edit	/manage/v2/databases/{id name}?view=edit
/manage/v1/databases/{id name}/status	/manage/v2/databases/{id name}?view=status
/manage/v1/forests/{id name}/config	/manage/v2/forests/{id name}?view=config
/manage/v1/forests/{id name}/counts	/manage/v2/forests/{id name}?view=counts
/manage/v1/forests/{id name}/edit	/manage/v2/forests/{id name}?view=edit
/manage/v1/forests/{id name}/status	/manage/v2/forests/{id name}?view=status
/manage/v1/groups/{id name}/config	/manage/v2/groups/{id name}?view=config
/manage/v1/groups/{id name}/counts	/manage/v2/groups/{id name}?view=counts
/manage/v1/groups/{id name}/edit	/manage/v2/groups/{id name}?view=edit
/manage/v1/groups/{id name}/status	/manage/v2/groups/{id name}?view=status
/manage/v1/hosts/{id name}/config	/manage/v2/hosts/{id name}?view=config
/manage/v1/hosts/{id name}/counts	/manage/v2/hosts/{id name}?view=counts
/manage/v1/hosts/{id name}/edit	/manage/v2/hosts/{id name}?view=edit
/manage/v1/hosts/{id name}/status	/manage/v2/hosts/{id name}?view=status
/manage/v1/servers/{id name}/config	/manage/v2/servers/{id name}?view=config
/manage/v1/servers/{id name}/edit	/manage/v2/servers/{id name}?view=edit
/manage/v1/servers/{id name}/status	/manage/v2/servers/{id name}?view=status

### **JSON Output Includes Units**

JSON output for the various GET methods now includes units for metrics that were previously flat key-value pairs. The following template summarizes the difference:

- MarkLogic 6: "key" : *the-value*
- MarkLogic 7: "key" : { "units" : "*the-units*", "value" : *the-value* }

For example, in MarkLogic 6, elapsed-time was reported as follows:

```
"elapsed-time" : "0.023453"
```

In MarkLogic 7, this field include a unit, as shown in this example:

```
"elapsed-time" : { "units" : "sec", "value" : "0.023453" }
```

### **Element/Key Name Changes in Status Views**

The following XML element/JSON key name changes have been made:

- Many element/key names with a “total-” prefix in status views no longer use this prefix. This change affects the `load-detail` and `rate-detail` sections of the status views for clusters, databases, forests, groups, and hosts.
- The `on-disk-size` element/key of the databases status view has been renamed to `data-size`. This change affects `GET /manage/version/databases/{id|name}?view=status` output.

For example, in earlier releases, the `load-detail` section of the report returned by `GET /manage/LATEST/hosts/{id|name}?view=status` includes a `total-query-read-load` element/key. The equivalent data is now named `query-read-load`.

### **Changes to the Management API Plugins**

There have been some changes to the Management API plugins for MarkLogic 9, so if you have a plugin written against previous versions, depending upon the plugin, you might need to update it.

In previous releases (MarkLogic 5 and MarkLogic 6), plugins are installed in the following folder:

```
Assets/plugins/marklogic/manage/v1
```

Plugins installed to this directory will continue to work with the following exceptions:

- Previous releases support 3 types of plugins: "format", "resource", and "view extender". In MarkLogic 7:
  - No changes are necessary for "format" plugins, these will continue to work as before with no changes to the plugin module or on the client side.
  - New plugins installed in MarkLogic 7 or later should be created in:

```
Assets/plugins/marklogic/manage/extensions
```

- For "resource" plugins, clients must change the version step to `LATEST` (from `v1`). For example, `/manage/LATEST/myplugin`. No changes are required to the plugin code.
- Support is no longer available for the third type of plugin ("view extender"). If you have a plugin of this type, you can rewrite it as a resource plugin in MarkLogic 7 to provide the same functionality.

For details on extending the Management API with plugins, see [Extending Management API with Plugins](#) in the *Monitoring MarkLogic Guide*.

### **Changes to the Packaging API**

The Packaging REST API has changed for MarkLogic 7. Applications written using the MarkLogic 6 Packaging REST API (v1) must be rewritten to work with the MarkLogic 7 Packaging REST API (v2).

The differences between the v1 and v2 versions of the Packaging REST API are summarized in the table below.

MarkLogic 6	MarkLogic 7
<code>/v1/list/package/appserver={name}</code> (GET)	<code>/v2/servers/{name}?view=package</code> (GET) <code>/v2/packages/{pkgname}</code> (POST)
<code>/v1/list/package/database={name}</code> (GET)	<code>/v2/databases/{name}?view=package</code> (GET) <code>/v2/packages/{pkgname}</code> (POST)
<code>/v1/package/compare</code> (GET)	<code>/v2/packages/{pkgname}?view=differences</code> (GET)
<code>/v1/package/install</code> (POST)	<code>/v2/packages/{pkgname}/install</code> (POST)
<code>/v1/package-tickets/revert</code> (POST)	<code>/v2/tickets/{ticketnumber}/revert</code> (PUT)

#### **4.44.4.12 Changes to the Configuration Manager**

The Configuration Manager Packaging feature has been more tightly integrated with the Configuration Manager. Rather than exporting configurations to an XML file, as in MarkLogic 6, configurations in MarkLogic 7 are now exported, as XML, to a ZIP file.

You can import a configuration saved by MarkLogic 6 into MarkLogic 7. However, you cannot import a configuration saved in MarkLogic 7 into MarkLogic 6.

#### **4.44.4.13 `xdmp:plan` Now Requires a Privilege**

In MarkLogic 7, the `xdmp:plan` function requires a privilege (<http://marklogic.com/xdmp/privileges/xdmp-plan>). Previously, it did not require a privilege. If you have any code that calls `xdmp:plan` that is run by non-privileged users, you will need to add the privilege to a role that the users are granted (or to a role that they already have).

#### **4.44.4.14 `fn:analyze-string` Now Returns Output in a Different Namespace**

In MarkLogic 7, the `fn:analyze-string` function returns an XML node in the <http://www.w3.org/2005/xpath-functions> namespace, which is the namespace specified by the working draft of the latest XQuery specification. Previously, this function was not fully specified and it returned XML in the <http://www.w3.org/2009/xpath-functions/analyze-string> namespace. If you have code that is expecting output in the old namespace, you will need to change that code to expect the new namespace.



## 5.0 Planning for Future Upgrades

This chapter provides guidelines and warnings for preparing for changes expected in a future release, such as announcements of deprecated interfaces. You are not required to make changes related to the topics in this section at this time, but you should plan to do so in the future.

- [Hadoop Connector Deprecated](#)
- [HDFS Support Deprecated](#)
- [Packaging API Deprecated](#)
- [info and infodev APIs Deprecated](#)
- [Annotated Query Output from search:parse Deprecated](#)
- [Search API Grammar Customization Deprecated](#)
- [Search API searchable-expression Deprecated](#)
- [The mlcp Option -tolerate\\_errors Deprecated](#)
- [xdmp:transaction-mode XQuery Prolog Option Deprecated](#)
- [Deprecation of transaction-mode Option to xdmp:eval](#)
- [XCC Session.setTransactionMode is Deprecated](#)
- [Java Client API 4.0.2 Deprecations](#)
- [Java Client API 4.0.4 Deprecations](#)
- [REST Client API Namespace Configuration Deprecation](#)
- [Configuration Packaging XQuery Library Deprecated](#)
- [Configuration Manager Deprecated](#)

### 5.1 Hadoop Connector Deprecated

The Hadoop Connector is deprecated starting with MarkLogic release 9.0-12 and is removed from the product in a 9.0-13.

### 5.2 HDFS Support Deprecated

HDFS support was deprecated in MarkLogic 9.0-12.

### 5.3 Packaging API Deprecated

In MarkLogic 9, the Packaging API has been deprecated. It will be removed from the product in MarkLogic 10.

The following table lists the deprecated endpoints and the new alternative.

Deprecated Endpoints	Alternative
<p><b>Base path</b> /manage/v2/packages:</p> <ul style="list-style-type: none"> <li>• GET /manage/v2/packages</li> <li>• POST /manage/v2/packages</li> <li>• GET /manage/v2/packages/{pkgname}</li> <li>• POST /manage/v2/packages/{pkgname}</li> <li>• HEAD /manage/v2/packages/{pkgname}</li> <li>• DELETE /manage/v2/packages/{pkgname}</li> <li>• GET /manage/v2/packages/{pkgname}/databases</li> <li>• GET /manage/v2/packages/{pkgname}/databases/{name}</li> <li>• POST /manage/v2/packages/{pkgname}/databases/{name}</li> <li>• HEAD /manage/v2/packages/{pkgname}/databases/{name}</li> <li>• DELETE /manage/v2/packages/{pkgname}/databases/{name}</li> <li>• POST /manage/v2/packages/{pkgname}/install</li> <li>• GET /manage/v2/packages/{pkgname}/servers</li> <li>• GET /manage/v2/packages/{pkgname}/servers/{name}</li> <li>• POST /manage/v2/packages/{pkgname}/servers/{name}</li> <li>• HEAD /manage/v2/packages/{pkgname}/servers/{name}</li> <li>• DELETE /manage/v2/packages/{pkgname}/servers/{name}</li> </ul>	<p>Use the following endpoints of CMA REST API instead:</p> <ul style="list-style-type: none"> <li>• GET /manage/v3</li> <li>• POST /manage/v3</li> </ul> <p>For more details, see <a href="#">“Configuration Management API (CMA) REST Endpoints”</a> on page 43.</p>
<p><b>Base path</b> /manage/v2/tickets:</p> <ul style="list-style-type: none"> <li>• POST /manage/v2/tickets/{ticketnumber}/revert</li> </ul>	

## 5.4 info and infodev APIs Deprecated

The XQuery library modules in the `info` and `infodev` namespaces are deprecated as of MarkLogic 9 and will be removed from the product in a future release. For example, the functions `info:ticket` and `infodev:ticket-create` are deprecated.

## 5.5 Annotated Query Output from `search:parse` Deprecated

The `search:parse` XQuery function and `search.parse` Server-Side JavaScript functions can return an annotated `cts:query` if you pass in `"cts:annotated-query"` as the output format parameter value. As of MarkLogic 9, use of `"cts:annotated-query"` is deprecated. Support for this format will be removed in a future release.

If you currently use the annotated query output as an intermediate step in a transformation, you should use the structured query (`"search:query"`) output format instead. Runtime modification of queries is a primary use case for structured query. For more details, see [Searching Using Structured Queries](#) in the *Search Developer's Guide*.

If you currently use the annotated query output format to recover the original query text using `search:unparse`, you should cache the original query text yourself.

## 5.6 Search API Grammar Customization Deprecated

Customizing the string query grammar through the Search API `grammar` query option is now deprecated. Support for this feature will be removed in a future release.

If your application currently relies on a Search API grammar customization, you should consider alternatives such as the following:

- `xqysp` (<http://github.com/mblakele/xqysp>) for XQuery.
- `PEG.js` (<http://pegjs.org/>) or `Jison` (<http://github.com/zaach/jison>) for Server-Side JavaScript.

## 5.7 Search API searchable-expression Deprecated

Due to security and performance considerations, beginning in MarkLogic 9.0-10, the `searchable-expression` property/element in query options is deprecated.

In addition, in 9.0-10 and moving forward, the `searchable-expression` requires the `eval-search-string` privilege.

Before MarkLogic 11, Search API users should modify queries that use query options with searchable expressions to remove the `searchable-expression`. The replacement is as follows:

- For the parts of the searchable expression resolvable with indexes, use query clauses instead
- For the parts of the searchable expression resolvable only by scanning (such as regex predicates), the preferred approach is to remove them in favor of using indexes. Where necessary, filter the results in post-processing

## 5.8 The `mlcp` Option `-tolerate_errors` Deprecated

The `-tolerate_errors` option of the `mlcp import` command is deprecated (and ignored) as of MarkLogic 9.0-2. The option will be removed in a future release. `mlcp` now always tolerates errors.

## 5.9 `xdmp:transaction-mode` XQuery Prolog Option Deprecated

The XQuery prolog option `xdmp:transaction-mode` is deprecated as of MarkLogic 9.0-2. Use the `xdmp:commit` and `xdmp:update` prolog options instead.

Note that the new prolog options differ from `xdmp:transaction-mode` in that they affect only the transaction create after their declaration, where as `xdmp:transaction-mode` settings persist across an entire session.

The following table illustrates the correspondence between the old and new options settings.

<code>xdmp:transaction-mode</code> Value	Equivalent <code>xdmp:commit</code> and <code>xdmp:update</code> Option Settings
"auto"	<code>declare option xdmp:commit "auto";</code> <code>declare option xdmp:update "auto";</code>
"update-auto-commit"	<code>declare option xdmp:commit "auto";</code> <code>declare option xdmp:update "true";</code>
"update"	<code>declare option xdmp:commit "explicit";</code> <code>declare option xdmp:update "true";</code>
"query"	<code>declare option xdmp:commit "explicit";</code> <code>declare option xdmp:update "false";</code>

Note that the default values for `xdmp:commit` and `xdmp:update` are both “auto”, so you do not need to set this value explicitly in most cases.

For more details, see [xdmp:update](#) and [xdmp:commit](#) in the *XQuery and XSLT Reference Guide*.



## 5.10 Deprecation of transaction-mode Option to xdmp:eval

The `transaction-mode` option of the `xdmp:eval` XQuery function and the `xdmp.eval` JavaScript function is deprecated as of MarkLogic 9.0-2. Use the `commit` and `update` options instead. For more details, see the function reference documentation for `xdmp:eval` (XQuery) and `xdmp.eval` (JavaScript).

This option deprecation (and alternative option settings apply to the following functions:

XQuery	JavaScript
<code>xdmp:eval</code>	<code>xdmp.eval</code>
<code>xdmp:javascript-eval</code>	<code>xdmp.xqueryEval</code>
<code>xdmp:invoke</code>	<code>xdmp.invoke</code>
<code>xdmp:invoke-function</code>	<code>xdmp.invokeFunction</code>
<code>xdmp:spawn</code>	<code>xdmp.spawn</code>
<code>xdmp:spawn-function</code>	

The following table illustrates the correspondence between the old and new option settings:

<code>transaction-mode</code> Option Value	Equivalent <code>commit</code> and <code>update</code> Option Values
<code>auto</code>	<code>commit: "auto"</code> <code>update: "auto"</code>
<code>update-auto-commit</code>	<code>commit: "auto"</code> <code>update: "true"</code>
<code>update</code>	<code>commit: "explicit"</code> <code>update: "true"</code>
<code>query</code>	<code>commit "explicit"</code> <code>update "false"</code>

## 5.11 XCC Session.setTransactionMode is Deprecated

Use of `Session.setTransactionMode` to specify commit semantics and transaction type is deprecated as MarkLogic 9.0-2. This function will be removed in a future version. Use the new `Session.setAutoCommit` and `Session.setUpdate` methods instead.

The following table illustrates how to replace calls to `setTransactionMode` with equivalent calls to `setAutoCommit` and `setUpdate`.

If you call <code>setTransactionMode</code> with this value:	Then replace it with the following calls on the same Session object
AUTO	<pre>setAutoCommit(true); setUpdate(Session.Update.AUTO);</pre>
MULTI_AUTO	<pre>setAutoCommit(false); setUpdate(Session.Update.AUTO);</pre>
UPDATE	<pre>setAutoCommit(false); setUpdate(Session.Update.TRUE);</pre>
QUERY	<pre>setAutoCommit(false); setUpdate(Session.Update.FALSE);</pre>
UPDATE_AUTO_COMMIT	<pre>setAutoCommit(true); setUpdate(Session.Update.TRUE);</pre>
QUERY_SINGLE_STATEMENT	<pre>setAutoCommit(true); setUpdate(Session.Update.FALSE);</pre>

## 5.12 Java Client API 4.0.2 Deprecations

Java Client API 4.0.2 introduces the following deprecations.

The `com.marklogic.client.extra.httpclient.HttpClientConfigurator` interface is deprecated and will be removed in a future release. Use the new

`com.marklogic.client.extra.okhttpclient.OkHttpClientConfigurator` interface instead.

Attaching a configurator based on `HttpClientConfigurator` to a `DatabaseClientFactory` object no longer has any effect on the HTTP configuration.

The single parameter version of `DatabaseClientFactory.SecurityContext.withSSLContext` has been deprecated and will be removed in a future release. Instead, use the new version that requires an `X509TrustManager` as its second parameter. This change affects all classes that implement `DatabaseClientFactory.SecurityContext`, such as `DatabaseClientFactory.CertificateAuthContext`, `DatabaseClientFactory.KerberosAuthContext`, and `DatabaseClientFactory.DigestAuthContext`.

## 5.13 Java Client API 4.0.4 Deprecations

Java Client API 4.0.4 deprecates the following interfaces:

- [NamespacesManager Interface Deprecated](#)
- [QueryBatcher.getQuerySuccessListeners Deprecated](#)

### 5.13.1 NamespacesManager Interface Deprecated

The Java Client API interfaces for configuration namespace bindings are deprecated as of version 4.0.4 and will be removed in a future release.

Instead, you should use the REST Management API to define namespace bindings for your App Server. For details, see the `namespaces` property of the payload for `PUT:/manage/v2/servers/{id|name}/properties`.

This deprecation notice affects the following components of the Java Client API:

- The `com.marklogic.client.admin.NamespacesManager` interface.
- The `com.marklogic.client.admin.ServerConfigurationManager.newNamespacesManager` method.

### 5.13.2 QueryBatcher.getQuerySuccessListeners Deprecated

The method `QueryBatcher.getQuerySuccessListeners` is deprecated as of Java Client API 4.0.0 and will be removed in a future release. Use `QueryBatcher.getUriReaderListeners` instead.

Both methods do the same thing in the same way. You do not need to change anything but the method name.

## 5.14 REST Client API Namespace Configuration Deprecation

The REST Client API methods for configuring namespace bindings are deprecated as of MarkLogic 9.0-5 and will be removed in a future release.

Instead, you should use the REST Management API to define namespace bindings for your App Server. For details, see the `namespaces` property of the payload for `PUT:/manage/v2/servers/{id|name}/properties`.

This deprecation notice affects the following methods of the REST Client API:

- GET, POST, PUT and DELETE on `/v1/config/namespaces`
- GET, PUT, and DELETE on `/v1/config/namespaces/{prefix}`

## 5.15 Configuration Packaging XQuery Library Deprecated

In MarkLogic 9.0-5, the Configuration Packaging XQuery API library has been deprecated. It will be removed from the product in a future release.

The following table lists the deprecated methods and the new alternative.

Deprecated Methods	Alternative
<p>Configuration Packaging (pkg:) library:</p> <ul style="list-style-type: none"> <li>• pkg:create</li> <li>• pkg:database-configuration</li> <li>• pkg:delete</li> <li>• pkg:differences</li> <li>• pkg:errors</li> <li>• pkg:exists</li> <li>• pkg:get-database</li> <li>• pkg:get-database-list</li> <li>• pkg:get-modules</li> <li>• pkg:get-package</li> <li>• pkg:get-package-list</li> <li>• pkg:get-server</li> <li>• pkg:get-server-list</li> <li>• pkg:install</li> <li>• pkg:installable</li> <li>• pkg:put-database</li> <li>• pkg:put-modules</li> <li>• pkg:put-server</li> <li>• pkg:remove-database</li> <li>• pkg:remove-modules</li> <li>• pkg:remove-server</li> <li>• pkg:revert</li> <li>• pkg:server-configuration</li> <li>• pkg:valid</li> </ul>	<p>Use the following methods of Configuration Management (cma:) library instead:</p> <ul style="list-style-type: none"> <li>• cma:generate-config</li> <li>• cma:apply-config</li> </ul> <p>For more details, see “Configuration Management API (CMA) XQuery and JavaScript Libraries” on page 43.</p>

## 5.16 Configuration Manager Deprecated

The Configuration Manager tool is deprecated starting with MarkLogic 9.0-5 and removed from MarkLogic Server in version 10.0-3.

## 6.0 Other Notes

This section provides the following information about MarkLogic Server:

- [Memory and Disk Space Requirements](#)
- [Compatibility with XQuery Specifications](#)
- [XQuery Extensions](#)
- [SQL Queries](#)
- [Documentation](#)
- [Browser Requirements](#)
- [Security: Prevent Abuse of System Entity Expansion](#)

### 6.1 Memory and Disk Space Requirements

MarkLogic Server requires at least 2 GB 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 1.5 times the disk space of the total forest size. Specifically, each forest on a filesystem requires its filesystem to have at least 1.5 times the forest size in disk space (or, for each forest less than 48 GB, 3 times the forest size) when the `merge_max_size` database merge setting is set to the default of 48 GB. This translates to approximately 1.5 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 200 GB database, reserve at least 300 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*.

\* You need at least 2 times the `merge_max_size` of free space per forest, regardless of the forest size. Therefore, with the default `merge_max_size` of 48 GB, you need at least 96 GB of free space. Additionally, if your journals are not yet created, you need 2 times the journal size of free disk space (if the journal space is not yet allocated). Therefore, to be safe, you need (with the default `merge_max_size` and a 2G journal size) at least 100 GB of free space for each forest, no matter what size the forest is.

## 6.2 Compatibility with XQuery Specifications

MarkLogic 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*.

## 6.3 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>.

## 6.4 SQL Queries

This section lists the SQL queries that are known not to work in 9.0-13 and those that are not yet fully optimized.

- Sub-select / sub-query (with any operator - `IN`, `NOT IN`, `=`, `>`, etc) is not optimized.
- Queries that return tens of thousands of rows are not optimized. Use `LIMIT` to restrict the number of rows returned, or focus your query with more filters in the `WHERE` clause or a `cts:query` to restrict by collection, date range, etc. If you are using a BI Tool, configure it to avoid an unrestricted `SELECT * FROM table`.
- Special use of the `MATCH` keyword is retained in MarkLogic 9.0-13, but `MATCH` queries run filtered (`MATCH` does not use the Universal Index).
- `FULL OUTER JOIN` is not supported in 9.0-13

- SQL is read-only in 9.0-13. The only way to create a persistent view is by means of a template view or a range view. You can't create a view over views and save it. And you can't create a `TEMP` table. Applications, such as BI Tools, that rely on creating `TEMP` tables to cache results need to be configured to use some other method.

## 6.5 Documentation

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

Documentation	Description
<i>MarkLogic Server-Side JavaScript Function Reference</i>	API documentation for the Server-Side MarkLogic built-in extensions to the JavaScript standard functions.
<i>MarkLogic XQuery and XSLT Function Reference</i>	API documentation for the MarkLogic built-in and module extensions to the XQuery standard functions, as well as API documentation for the W3C functions implemented in MarkLogic Server.
<i>Getting Started with MarkLogic Server</i>	A quick, step-by-step overview of how to get up and running with MarkLogic Server.
<i>Installation Guide</i>	Provides procedures for installing MarkLogic Server.
<i>Release Notes</i>	Contains a summary of new features and upgrade compatible information.
<i>Concepts Guide</i>	Provides an overview of MarkLogic and conceptual information about the server architecture.
<i>Application Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about general MarkLogic Server application development tasks.
<i>Search Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about search-based application development tasks.
<i>Node.js Application Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about developing MarkLogic Server applications using the Node.js Client API.
<i>Java Application Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about developing MarkLogic Server applications using the Java API.

Documentation	Description
<i>MarkLogic Connector for Hadoop Developer's Guide</i>	Provides information on the MarkLogic Connector for Hadoop, a Java library to help you build applications that combine MarkLogic Server and Hadoop map-reduce jobs.
<i>REST Application Developer's Guide</i>	Provides information on MarkLogic Server administration and application development using the MarkLogic REST API.
<i>Semantics Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about semantic application development tasks and MarkLogic SPARQL and triple support.
<i>Temporal Developer's Guide</i>	Provides information on developing applications using MarkLogic bi-temporal features.
<i>Entity Services Developer's Guide</i>	Provides procedures, methodologies, and conceptual information about developing MarkLogic Applications using entity relationship modeling and the Entity Services API.
<i>Reference Application Architecture Guide</i>	Provides an overview of reference application architectures for multi-tiered applications built using MarkLogic as the database.
<i>Administrator's Guide</i>	Provides procedures for administrative tasks such as creating servers, creating databases, backing up databases, creating users, setting up your security policy, and so on.
<i>Scripting Administrative Tasks Guide</i>	Provides information on writing code to script various administrative tasks such as creating and modifying databases, App Servers, and so on.
<i>Database Replication Guide</i>	Provides information on database replication, useful for disaster recovery scenarios.
<i>Flexible Replication Guide</i>	Provides information on Flexible Replication, useful for information sharing from one database to another.
<i>Ops Director Guide</i>	Provides procedures for setting up and using Ops Director for monitoring MarkLogic clusters.
<i>Monitoring MarkLogic Guide</i>	Provides information on monitoring MarkLogic Server, including using the built-in monitoring tools and integrating with external tools such as HP Operations Manager.



Documentation	Description
<i>MarkLogic Connector for SharePoint® Administrator's Guide</i>	Documentation for the MarkLogic Connector for SharePoint®, which allows you to mirror documents from a Microsoft SharePoint repository in MarkLogic Server.
<i>JavaScript Reference Guide</i>	A language reference for the MarkLogic Server-Side JavaScript language. This book includes MarkLogic-specific Object reference, but is not a comprehensive language reference.
<i>XQuery and XSLT Reference Guide</i>	A condensed overview of the XQuery language, including information on the three XQuery dialects in MarkLogic Server. This book does include some syntax information, although it is primarily intended as an introduction and quick-reference to the language, not as a comprehensive reference.
<i>mlcp User Guide</i>	A procedural guide that explains how to use MarkLogic Content Pump (mlcp) command line tool to load content into MarkLogic, extract content from MarkLogic, or copy content between databases.
<i>Content Processing Framework Guide</i>	Provides an introduction to the Content Processing Framework and procedures for installing the default content processing framework.
<i>Query Performance and Tuning Guide</i>	Provides performance-related information that is useful to application developers and administrators.
<i>Scalability, Availability, and Failover Guide</i>	Provides information on large-scale system architecture, clustering, availability, and details on setting up shared-disk and local-disk failover.
<i>Security Guide</i>	Provides information on the role-based security model in MarkLogic Server.
<i>MarkLogic Server on Amazon Web Services (AWS) Guide</i>	Information about running MarkLogic Server in an EC2 environment.
<i>Query Console User Guide</i>	Provides step-by step information on using Query Console, a tool to create and run arbitrary XQuery code.
<i>Loading Content Into MarkLogic Server Guide</i>	Provides procedures, methodologies, and conceptual information about loading content into MarkLogic Server. Includes an overview of ingestion techniques available using XQuery, Java, REST, and the MarkLogic Content Pump (mlcp).

Documentation	Description
<i>SQL Data Modeling Guide</i>	Provides information on how to use MarkLogic’s SQL interface, including the creation of relational schemas and views.
<i>Messages and Codes Reference Guide</i>	A reference guide to MarkLogic Server and MarkLogic Application Services error and log messages.
<i>Glossary, Copyright, and Support</i>	Includes a glossary of terms as well as copyright and support information.
<i>MarkLogic REST API Reference</i>	API documentation for the REST API.
<i>Java Client API Documentation</i>	API documentation for the MarkLogic Java Client API.
<i>Node.js Client API Reference</i>	API documentation for the MarkLogic Node.js Client API.
<i>XCC Javadoc API Documentation</i>	API documentation for the MarkLogic XML Contentbase Connector for Java API (XCC/J).
<i>MarkLogic Hadoop MapReduce Connector API</i>	API documentation for the MarkLogic Hadoop MapReduce Connector.
<i>C++ UDF API Reference</i>	API documentation for the C++ User Defined Function (UDF) API.

XQuery language documentation is provided through the W3C working group drafts specified in “Compatibility with XQuery Specifications” on page 142. 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 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.

## 6.6 Browser Requirements

The Admin Interface and many of the other GUI tools (Query Console, Ops Director 1.x, and so on) are certified as follows:

- Internet Explorer 11 on Windows 7 and Windows 10.
- Chrome 58 on Windows and Mac OS.

Ops Director 2.0 is certified as follows:

- Internet Explorer 11 on Windows 10
- Chrome 63 on Windows 10 and Mac OS

Other browser/platform combinations may work but are not as thoroughly tested. Please consult the documentation for your individual GUI tool for exact browser requirements.

## **6.7 Security: Prevent Abuse of System Entity Expansion**

Normal XML processing allows for external entities to be referenced and included in the parsed content of XML files. If you want to disable this processing, set the trace event “Disable XML External Entities”.



## 7.0 Technical Support

MarkLogic provides technical support according to the terms detailed in your Software License Agreement or End User License Agreement.

We invite you to visit our support website at <http://help.marklogic.com> to access information on known and fixed issues, knowledge base articles, and more. For licensed customers with an active maintenance contract, see the [Support Handbook](#) for instructions on registering support contacts and on working with the MarkLogic Technical Support team.

Complete product documentation, the latest product release downloads, and other useful information is available for all developers at <http://developer.marklogic.com>. For technical questions, we encourage you to ask your question on [Stack Overflow](#).



## 8.0 Copyright

MarkLogic Server 9.0 and supporting products.  
Last updated: August 5, 2020

Copyright © 2020 MarkLogic Corporation.

MarkLogic and the MarkLogic logo are trademarks or registered trademarks of MarkLogic Corporation in the United States and other countries.

MarkLogic technology is protected by one or more U.S. Patent Nos. 7,127,469, 7,171,404, 7,756,858, 7,962,474, 8,935,267, 8,892,599, 9,092,507, 10,108,742, 10,114,975, 10,311,088, 10,325,106, 10,339,337, 10,394,889, and 10,503,780.

MarkLogic software incorporates certain third-party software under license. Third-party attributions, copyright notices, and other disclosures required under license are available in the respective notice document for your version of the MarkLogic software.

