

---

# **MarkLogic Server**

---

## **Installation Guide for All Platforms**

Release 4.1  
June, 2009

Last Revised: 4.1-4, December, 2009

## Copyright

© Copyright 2002-2009 by Mark Logic Corporation. All rights reserved worldwide.

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

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

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

Antenna House OfficeHTML Copyright © 2000-2008 Antenna House, Inc. All rights reserved.

Argus Copyright ©1999-2008 Icenit Technology Ltd. All rights reserved.

Contains Rosette Linguistics Platform 6.0 from Basis Technology Corporation, Copyright © 2004-2008 Basis Technology Corporation. All rights reserved.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>) Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved. Copyright © 1998-2001 The OpenSSL Project. All rights reserved. Contains software derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm. Copyright © 1991-1992, RSA Data Security, Inc. Created 1991. All rights reserved.

---



---

## Table of Contents

---

### Installation Guide for All Platforms

Copyright .....	2
1.0 Requirements and Database Compatibility .....	4
1.1 Introduction .....	4
1.2 Memory, Disk Space, and Swap Space Requirements .....	5
1.3 Supported Platforms .....	6
1.4 Upgrades and Database Compatibility .....	7
1.4.1 4.1 Database Compatibility With 4.0 .....	7
1.4.2 4.0 Database Compatibility With 3.2 .....	8
1.4.3 3.2 Database Compatibility With 3.1 .....	8
1.4.4 3.1/3.2 Database Compatibility With 3.0 .....	9
2.0 Procedures .....	11
2.1 Upgrading from Previous Releases .....	11
2.1.1 Note That Reindexing is On By Default .....	12
2.1.2 Upgrading from Release 4.1-1 or Later .....	12
2.1.3 Upgrading from Release 4.0 .....	13
2.1.4 Upgrading from Release 3.2 or earlier .....	14
2.2 Turning Off Internet Explorer User-Friendly HTTP Messages .....	14
2.3 Installing MarkLogic Server .....	16
2.4 Starting MarkLogic Server .....	19
2.5 Entering the License Key and Accepting the License Agreement .....	19
2.6 Configuring a New Installation .....	23
2.7 Enterprise Edition: Configuring the First and Subsequent Hosts .....	25
2.7.1 Enterprise Edition: Configuring a Single Host or the First Host in a Cluster 25	25
2.7.2 Enterprise Edition: Configuring an Additional Host in a Cluster .....	27
2.7.3 Enterprise Edition: Leaving a Cluster and Becoming a Single Host .....	29
2.8 Checking for the Correct Software Version .....	30
2.9 Configuring MarkLogic Server on UNIX Systems to Run as a Non-daemon User . 31	31
2.10 Removing MarkLogic Server .....	33
3.0 Technical Support .....	34

## 1.0 Requirements and Database Compatibility

This chapter introduces MarkLogic Server, lists the product requirements and supported platforms, and describes the database compatibility with previous releases. It includes the following sections:

- [Introduction](#)
- [Memory, Disk Space, and Swap Space Requirements](#)
- [Supported Platforms](#)
- [Upgrades and Database Compatibility](#)

### 1.1 Introduction

MarkLogic Server is a powerful software solution for harnessing your digital content base. MarkLogic Server enables you to build complex applications that interact with large volumes of content in XML, SGML, HTML and other popular content formats. The unique architecture of MarkLogic Server ensures that your applications are both scalable and high-performance, delivering query results at search-engine speeds while providing transactional integrity over the underlying content repository. MarkLogic Server Enterprise Edition can be configured for a distributed environment, enabling you to scale your infrastructure through hardware expansion.

This document describes how to install MarkLogic Server on your system. It covers both Enterprise and Standard Edition on all supported platforms. It is intended for a technical audience, specifically an IT staff with experience in XML.

This installation guide explains the procedures needed to install MarkLogic Server on your system. This document only explains how to install the software, not how to use the software. To learn how to get started using the software, see the following documents (available on [developer.marklogic.com](http://developer.marklogic.com)):

- *Getting Started With MarkLogic Server*
- *MarkLogic Server Administrator's Guide*
- *MarkLogic Server Application Developer's Guide*
- *MarkLogic Server Search Developer's Guide*
- *Mark Logic Built-In and Module Functions Reference*
- *MarkLogic Server Release Notes*

## 1.2 Memory, Disk Space, and Swap Space Requirements

Before installing the software, be sure that your system meets the following requirements:

- 512 MB of system memory, minimum. 2 GB or more recommended, depending on database size.

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.

- 3 times the disk space of the total forest size. Specifically, each forest on a filesystem requires its filesystem to have at least 3 times the forest size in disk space. This translates to three times the disk space of the source content after it is loaded.

For example, if you plan on loading content that will result in a 10 GB database, reserve at least 30 GB of disk space. The disk space reserve is required for merges.

- On UNIX systems, swap space at least equal to the amount of physical memory on the machine. Swap space equal to twice the amount of physical memory is highly recommended.

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

- On Linux systems, Mark Logic recommends setting Linux Huge Pages to 3/8 the size of your physical memory. For details on setting up Huge Pages, see <http://magazine.redhat.com/2007/05/29/how-do-i-set-up-hugepages-in-red-hat-enterprise-linux-4/>.

If you have Huge Pages set up on a Linux system, your swap space on that machine should be twice the size of your physical memory minus the size of your Huge Page (because Linux Huge Pages are not swapped). For example, if you have 64GB of physical memory, and if you have Huge Pages set to 24 GB, then you need swap space of 80 GB  $[(64 - 24) \times 2]$ .

At system startup on Linux machines, MarkLogic Server logs a message to the `ErrorLog.txt` file showing the Huge Page size, and the message indicates if the size is below the recommended level.

### 1.3 Supported Platforms

MarkLogic Server is supported on the following platforms:

- Microsoft Windows Server 2008 (x86), Microsoft Windows 2003 Server (x86), Microsoft Windows XP SP2, Microsoft Windows Vista 32-bit Edition (x86)\*
- Microsoft Windows Server 2008 (x64), Windows 2003 Server 64-bit Edition (x64), Windows Vista Server 64-bit Edition (x64)\*
- Sun Solaris 10 (64-bit SPARC)
- Sun Solaris 10 (x64)
- Red Hat Enterprise Linux 4.0 and 5.0 (x86)\*\* \*\*\*
- Red Hat Enterprise Linux 4.0 and 5.0 (x64)\*\* \*\*\* \*\*\*\*
- CentOS 5 (x64)\*\* \*\*\*

\* Microsoft Windows Vista is supported for development only.

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

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

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

64-bit versions of Windows require the following DLL:

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

\*\* The deadline I/O scheduler is required on Red Hat Linux platforms. The deadline scheduler is optimized to ensure efficient disk I/O for multi-threaded processes, and MarkLogic Server can have many simultaneous threads. For information on the deadline scheduler, see the Red Hat documentation (for example,

<http://www.redhat.com/docs/wp/performance/iotuning/iosubsystem-scheduler-deadline.html>,  
<http://www.redhat.com/docs/wp/performance/iotuning/iosubsystem-scheduler-selection.html>, and  
<http://www.redhat.com/magazine/008jun05/features/schedulers/>).

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

\*\*\*\*Red Hat Linux 5.0 (x64) is also supported in a VMWare ESX 3.0.2 (installed on bare metal) environment.

## 1.4 Upgrades and Database Compatibility

MarkLogic Server 4.1 supports upgrades from MarkLogic Server 4.1 or from MarkLogic Server 4.0-1 or later databases. If you are upgrading from an earlier version of MarkLogic Server, you must first upgrade to 4.0 before moving to 4.1. For the procedure for upgrading to MarkLogic Server 4.1 see “Upgrading from Previous Releases” on page 11.

During the upgrade, the security database and the schemas databases are automatically upgraded. The security database is upgraded with the latest execute privileges and the schemas database is upgraded with the latest version of the schemas used by MarkLogic Server. This upgrades occur as part of the installation procedure.

Databases that contain your own content are also upgraded to work with MarkLogic Server 4.1; once you upgrade to 4.1, you will no longer be able to use that database with MarkLogic Server 4.0. Mark Logic Corporation strongly recommends performing a backup of your databases before upgrading to MarkLogic Server 4.1.

For the procedure for upgrading to MarkLogic Server 4.1, see “Upgrading from Previous Releases” on page 11. For details about known incompatibilities between MarkLogic Server 4.0 and 4.1, see [Known Incompatibilities with Previous Releases](#) in the *Release Notes*.

This section contains database compatibility information between various releases, and includes the following sections.

- [4.1 Database Compatibility With 4.0](#)
- [4.0 Database Compatibility With 3.2](#)
- [3.2 Database Compatibility With 3.1](#)
- [3.1/3.2 Database Compatibility With 3.0](#)

### 1.4.1 4.1 Database Compatibility With 4.0

MarkLogic Server 4.1 databases include improvements compared to 4.0 databases. Therefore, to take advantage of these improvements, you need to reindex your 4.0 databases in 4.1. When you upgrade to 4.1, all databases with `reindex enable` set to `true` will automatically begin reindexing immediately. If you do not want the databases to reindex, you must set `reindex enable` to `false` before upgrading (that is, you must set `reindex enable` to `false` in 4.0). You can always reindex your content later by changing this setting back to `true` after installing MarkLogic Server 4.1.

Also, you can reindex your content a little bit at a time by enabling reindexing for a while, then disabling it for a while, then enabling it, and so on. You might want to use this technique to reindex your database during non-peak hours, for example, over a period of hours or days or weeks, depending on how large your database is. Always ensure that you have the proper system requirements, as defined in “Memory, Disk Space, and Swap Space Requirements” on page 5.

**Note:** If you are upgrading from 4.0 and you want to install Application Builder, after you upgrade and start MarkLogic Server, you can add the App Server and database needed for the MarkLogic Server as described in [Setting Up and Starting Application Builder](#) in the *Application Builder Developer’s Guide*.

### 1.4.2 4.0 Database Compatibility With 3.2

**Note:** This section is here for reference only about the compatibility between MarkLogic Server 3.2 and 3.1.

MarkLogic Server 4.0 databases can have several new indexes compared to 3.2 databases. Therefore, to take advantage of these new features, you need to reindex your 3.2 databases in 4.0. When you upgrade to 4.0, all databases with `reindex enable` set to `true` will automatically begin reindexing immediately. If you do not want the databases to reindex, you must set `reindex enable` to `false` *before* upgrading (that is, you must set `reindex enable` to `false` in 3.2). You can always reindex your content later by changing this setting back to `true` after installing MarkLogic Server 4.0.

Also, you can reindex your content a little bit at a time by enabling reindexing for a while, then disabling it for a while, then enabling it, and so on. You might want to use this technique to reindex your database during non-peak hours, for example, over a period of hours or days or weeks, depending on how large your database is. Always ensure that you have the proper system requirements, as defined in “Memory, Disk Space, and Swap Space Requirements” on page 5.

### 1.4.3 3.2 Database Compatibility With 3.1

**Note:** This section is here for reference only about the compatibility between MarkLogic Server 3.2 and 3.1.

MarkLogic Server 3.2 and newer databases are language-aware. Details such as stemming and tokenization are different for different languages. For an overview of how languages work in MarkLogic Server, see the “Languages, Encodings, and Collations” chapter of the *Application Developer’s Guide*.

Because of the differences in stemming and tokenization between languages, and because MarkLogic Server allows you to specify languages at the element level, the index data in MarkLogic Server databases contains language-specific information. Additionally, there are other features in 3.2 that require 3.2 (or newer) indexes.

When you upgrade to 3.2, all databases with `reindex enable` set to `true` will begin reindexing immediately. If you do not want the databases to reindex, you must set `reindex enable` to `false` *before* upgrading (that is, you must set `reindex enable` to `false` in 3.1). You can always reindex your content later by changing this setting back to `true` after installing MarkLogic Server 3.2.

MarkLogic Server 3.2 automatically detects the format of the content in the database. If you do not reindex the content when upgrading to 3.2 (that is, if you set `reindex enable` to `false` in your database configuration *before* you install 3.2), then the database will run in compatibility mode, and certain 3.2 features will be disabled (for example, non-English languages and `cts:frequency` for lexicon values).

#### 1.4.4 3.1/3.2 Database Compatibility With 3.0

**Note:** This section is here for reference only about the compatibility between MarkLogic Server 3.1 and 3.0.

You can have this same level of compatibility as other versions and 4.0 between 4.0 and 3.0 as well. You must upgrade to MarkLogic Server 3.1 before upgrading to MarkLogic Server 3.2 and 4.0. For details on upgraded from 3.0 to 3.1, see the MarkLogic Server 3.1 *Installation Guide*.

MarkLogic Server 3.1 and newer has diacritic-aware indexes. MarkLogic Server 3.0 does not have diacritic-aware indexes. Therefore, databases created with 3.0 will have different default behavior in 3.1 in the handling of searches that involve diacritic characters. Additionally, databases created in 3.0 that are not reindexed might miss diacritic-insensitive matches on `cts:search` operations. Note the following about user-created databases that are upgraded to MarkLogic Server 3.1:

- If reindexing is enabled, MarkLogic Server 3.1 will reindex any 3.0 database in order to rebuild the default text indexes as diacritic-insensitive indexes. After the reindex operation is completed, all diacritic searches will work correctly.
- If reindexing is not enabled, then diacritic searches using MarkLogic Server 3.1 will behave differently against 3.0 databases than 3.1 databases. Any `cts:search` operations on upgraded, un-reindexed 3.0 databases, with terms that incorporate diacritics (in the search specification or in the content loaded in MarkLogic Server) have the following characteristics:
  - If the term passed to `cts:search` incorporates one or more diacritic characters, and the search is either explicitly or implicitly `diacritic-sensitive`, then the search will work correctly.
  - If the term passed to `cts:search` incorporates one or more diacritic characters, and the search is explicitly `diacritic-insensitive`, then the search will not match terms in the source content that had diacritics.
  - If the term passed to `cts:search` does not incorporate any diacritic characters, and the search is explicitly `diacritic-sensitive`, then the search will work correctly.

- If the term passed to `cts:search` does not incorporate any diacritic characters, and if the search is implicitly or explicitly `diacritic-insensitive`, then the search will not match terms in the source content that incorporated diacritics (even though you specified `diacritic-insensitive matches`).
- If the source content does not incorporate diacritic characters, or if diacritic-related recall is not important to your application, these incomplete search issues may not pose a problem for you. However, you should understand that at any point in the future, if you enable reindexing, the database will undertake this reindexing task concurrently with any other reindexing tasks that may be required.
- Note that the database has no way to know whether or not your source content incorporates diacritics, so if reindexing is enabled, it will reindex all the documents in order to ensure that text indexes are diacritic-aware.
- Any new or changed content will be added diacritic-aware, potentially creating different diacritic behavior for different parts of your content set.
- If you choose not to reindex your 3.0 database, and if you always specify the `diacritic-sensitive` option to `cts:search`, then your results will be the same as they were in MarkLogic Server 3.0.

## 2.0 Procedures

This section describes the following procedures to install MarkLogic Server on your system.

- [Upgrading from Previous Releases](#)
- [Turning Off Internet Explorer User-Friendly HTTP Messages](#)
- [Installing MarkLogic Server](#)
- [Starting MarkLogic Server](#)
- [Entering the License Key and Accepting the License Agreement](#)
- [Configuring a New Installation](#)
- [Enterprise Edition: Configuring the First and Subsequent Hosts](#)
- [Checking for the Correct Software Version](#)
- [Configuring MarkLogic Server on UNIX Systems to Run as a Non-daemon User](#)
- [Removing MarkLogic Server](#)

Be sure to complete each procedure in the order presented.

### 2.1 Upgrading from Previous Releases

If you have previously installed MarkLogic Server on this machine, you must uninstall the old release before proceeding with the new installation. For information on removing the software, see “Removing MarkLogic Server” on page 33 or the *Installation Guide* from the previous release. This section describes the following information and upgrade paths:

- [Note That Reindexing is On By Default](#)
- [Upgrading from Release 4.1-1 or Later](#)
- [Upgrading from Release 4.0](#)
- [Upgrading from Release 3.2 or earlier](#)

If you are upgrading a cluster to a new release, see [Upgrading a Cluster to a New Maintenance Release of MarkLogic Server](#) in the *Scalability, Clustering, and Forest-Level Failover*.

**Note:** If you are upgrading from 4.0 and you want to install Application Builder, after you upgrade and start MarkLogic Server, you can add the App Server and database needed for the MarkLogic Server as described in [Setting Up and Starting Application Builder](#) in the *Application Builder Developer’s Guide*.

### 2.1.1 Note That Reindexing is On By Default

Before upgrading from 3.2 to 4.0, be aware that database reindexing is enabled by default. If you have a very large database, or if you have many databases configured, they will start reindexing after you install MarkLogic Server as soon as you accept the license key. Reindexing is a CPU and disk-intensive process, and it can take some time. Consequently, it can slow down the machine, particularly if the machine has a slow disk system (for example, a development laptop).

If you want to delay reindexing, set `reindex enable` to `false` in your database configurations *before* you install MarkLogic Server 4.0 (that is, you must set `reindex enable` to `false` in 3.2). MarkLogic Server 4.0 will then run in 3.2 compatibility mode until reindexing is enabled and completes. To re-enable indexing, set `reindex enable` to `true` in your database configurations after completing the installation (that is, after accepting the license key, and performing the configuration and security database upgrades). After reindexing has been enabled, it will commence immediately and continue until it is done, or until reindexing is disabled. For details about database compatibility, see “Upgrades and Database Compatibility” on page 7.

### 2.1.2 Upgrading from Release 4.1-1 or Later

To upgrade from release 4.1-1 or later to the current 4.1 release, perform the following basic steps:

1. Shut down MarkLogic Server 4.1 (as described in step 1 of “Removing MarkLogic Server” on page 33).
2. Uninstall the old 4.1 release (as described in “Removing MarkLogic Server” on page 33).
3. Install the new 4.1 release (as described in “Installing MarkLogic Server” on page 16).

If you are upgrading a cluster to a new release, see [Upgrading a Cluster to a New Maintenance Release of MarkLogic Server](#) in the *Scalability, Clustering, and Forest-Level Failover*.

**Warning:** There is no direct upgrade path from Early Access releases of 4.1 to this release. If you need to move any data from an Early Access release to this release, you must re-create the data in the current release.

### 2.1.3 Upgrading from Release 4.0

MarkLogic Server 4.1 installs in the same default directory as Release 4.0, so there is no need to move any old files around. The upgrade to Release 4.1 does incorporate an automatic update to the Security database and to the configuration files. The Security database upgrade occurs when you first access the Admin Interface after the Release 4.1 installation (and after accepting the license agreement).

Uninstalling a previous release of MarkLogic Server does not remove or delete the user data files (the forests and configuration information). When upgrading to MarkLogic Server 4.1, you must first uninstall previous releases of MarkLogic Server.

**Note:** For Enterprise Edition installations, if you are upgrading a cluster of two or more servers, you must first upgrade the server in which the Security database is mounted. The Security database must be available before subsequent hosts can join the cluster.

The basic steps to upgrade from Release 4.0 to Release 4.1 are as follows:

1. As a precaution, perform database backups on your MarkLogic Server 4.0 databases.
2. If you do not want to reindex a database, disable reindexing for that database before installing MarkLogic Server 4.1 (set `reindexer enable` to `false` on the Database Configuration page of the Admin Interface). The entire database will be reindexed after upgrading to 4.1 unless you disable reindexing. For details, see “Upgrades and Database Compatibility” on page 7. You can always reindex the database later by setting `reindexer enable` database configuration option to `true`.
3. Shut down MarkLogic Server 4.0 (as described in step 1 of “Removing MarkLogic Server” on page 33).
4. Uninstall MarkLogic Server 4.0 (as described in “Removing MarkLogic Server” on page 33).
5. Install MarkLogic Server 4.1 (as described in “Installing MarkLogic Server” on page 16).
6. Start MarkLogic Server 4.1 (as described in “Starting MarkLogic Server” on page 19).
7. Open the Admin Interface in a browser (<http://localhost:8001/>).
8. Accept the license agreement.

**Note:** Accepting this license agreement is optional. However, to begin using MarkLogic Server, you must accept the terms and conditions outlined in the license agreement. If you have executed a written software license with Mark Logic Corporation, the agreement displayed references that written agreement.

9. When the Admin Interface prompts you to upgrade the Security database and the configuration files, click the button to confirm the upgrade.

There are some known application incompatibilities between 4.1 and 4.0, some of which might require minor code changes to your applications. For details on these incompatibilities, see the *Release Notes*.

### 2.1.4 Upgrading from Release 3.2 or earlier


MarkLogic Server 4.1 only supports upgrading from Release 4.0-1 or later; it does not provide a direct upgrade path for previous releases of MarkLogic Server. If you are upgrading from a 3.2 or earlier release of MarkLogic Server, either install this release as a clean installation or upgrade your existing release to Release 4.0 before installing this release. For details on upgrading a MarkLogic Server Release 3.2 or earlier release to 4.0, see the *Installation Guide* for MarkLogic Server 4.0. If you are upgrading from 3.2 and do not want to reindex your content, set the `reindexer enable to false` before upgrading from 4.0 to 4.1. You can run MarkLogic Server 4.1 in either 3.0, 3.1, 3.2, or 4.0 compatibility mode, as described in “Upgrades and Database Compatibility” on page 7.

## 2.2 Turning Off Internet Explorer User-Friendly HTTP Messages

If you are using a Microsoft Internet Explorer browser, turn off the “Show friendly HTTP messages” option. Turning this option off allows any error messages from MarkLogic Server to display in full in your browser. Perform the following steps to turn this option off:

1. Start Internet Explorer.
2. Select Tools > Internet Options from the Internet Explorer menu.  
The Internet Options dialog box appears.
3. Click the Advanced tab.

4. Scroll down and uncheck the “Show friendly HTTP messages” option.



5. Click OK.

## 2.3 Installing MarkLogic Server

This section describes the procedure for installing MarkLogic Server on each platform. Perform the procedure corresponding to the platform to which you are installing.

If you are upgrading a cluster to a new release, see [Upgrading a Cluster to a New Maintenance Release of MarkLogic Server](#) in the *Scalability, Clustering, and Forest-Level Failover*.

Platform	Perform the following:
Windows (x86 and x64)	<ol style="list-style-type: none"> <li data-bbox="428 579 1414 722">1. Shut down and uninstall the previous release of MarkLogic Server (if you are upgrading from 4.0, see “Upgrading from Release 4.0” on page 13, if you are upgrading from 4.0-1 or later, see “Removing MarkLogic Server” on page 33).</li> <li data-bbox="428 743 1414 852">2. Download the MarkLogic Server installation package to your desktop. The latest installation packages are available from the Mark Logic Developer Network (<a href="http://developer.marklogic.com">http://developer.marklogic.com</a>).</li> <li data-bbox="428 873 1414 982">3. Double click the <code>MarkLogic-4.1-1-i686.msi</code> icon to start the installer (<code>MarkLogic-4.1-1-amd64.msi</code> for x64 (64-bit AMD Opteron and 64-bit Intel EM64T)).  <b>Note:</b> If you are installing a release other than 4.1-1, double-click on the appropriately named installer icon.</li> <li data-bbox="428 1129 1045 1163">4. The Welcome page displays. Click Next.</li> <li data-bbox="428 1184 711 1218">5. Select Typical.</li> <li data-bbox="428 1239 688 1272">6. Click Install.</li> <li data-bbox="428 1293 688 1327">7. Click Finish.</li> </ol>

Platform	Perform the following:
Red Hat Linux (x86 and x64)	<ol style="list-style-type: none"> <li data-bbox="428 279 1414 422">1. Shut down and uninstall the previous release of MarkLogic Server (if you are upgrading from 4.0, see “Upgrading from Release 4.0” on page 13, if you are upgrading from 4.0-1 or later, see “Removing MarkLogic Server” on page 33).</li> <li data-bbox="428 443 1414 810">2. Download the package to /tmp using your web browser. The latest installation packages are available from the Mark Logic Developer Network (<a href="http://developer.marklogic.com">http://developer.marklogic.com</a>).  <b>Note:</b> If you are using Firefox or another browser that is configured to associate rpm files, the browser will prompt you for the root password (if you are not already running as root) and you can follow the prompts to complete the installation. When the installation is complete, you can skip the next step. Otherwise, continue to the next step.</li> <li data-bbox="428 852 1414 1136">3. As the root user, install the package with the following command:  <pre>rpm -i /tmp/MarkLogic-4.1-1.i686.rpm</pre> <b>Note:</b> The package name will end in amd64.rpm for x64 (64-bit AMD Opteron and 64-bit Intel EM64T). If you are installing a release other than 4.1-1, replace the characters 4.1-1 in the line above with the appropriate release number.</li> </ol>

Platform	Perform the following:
Sun Solaris (64-bit SPARC and x64)	<ol style="list-style-type: none"> <li>1. Shut down and uninstall the previous release of MarkLogic Server (if you are upgrading from 4.0, see “Upgrading from Release 4.0” on page 13, if you are upgrading from 4.1-1 or later, see “Removing MarkLogic Server” on page 33).</li> <li>2. Download the package to <code>/var/spool/pkg</code> using your web browser. The latest installation packages are available from the Mark Logic Developer Network (<a href="http://developer.marklogic.com">http://developer.marklogic.com</a>).</li> <li>3. Unpack the compressed tar file in <code>/var/spool/pkg</code> with the following shell commands: <pre style="margin-left: 40px;">% cd /var/spool/pkg % uncompress MARKlogic-4.1-1-sparc.tar.Z % tar xf MARKlogic-4.1-1-sparc.tar % rm MARKlogic-4.1-1-sparc.tar</pre> <p><b>Note:</b> The package name will end in <code>amd64.tar.Z</code> for x64 (64-bit AMD Opteron and Intel 64-bit EM64T). If you are installing a release other than 4.1-1, replace the characters 4.1-1 in the line above with the appropriate release number.</p> </li> <li>4. As the <code>root</code> user, install the package with the following command: <pre style="margin-left: 40px;"># pkgadd MARKlogic</pre> </li> </ol>

The following table shows the installation directory and the default data directory for each platform:

Platform	Installation Directory	Default Data Directory (for configuration and log files)
Windows	c:\Program Files\MarkLogic\	c:\Program Files\MarkLogic\Data
Red Hat Linux	/opt/MarkLogic	/var/opt/MarkLogic
Sun Solaris	/opt/MARKlogic	/var/opt/MARKlogic

The default forest directory is the same as the default data directory if the optional data directory is not specified during forest creation. On UNIX platforms, if you want MarkLogic Server to use another location for its default data directory, make your data directory (`/var/opt/MarkLogic` on Linux and `/var/opt/MARKlogic` on Solaris) a soft link to the alternate location.

## 2.4 Starting MarkLogic Server

MarkLogic Server will automatically start when the computer reboots. To start MarkLogic Server without rebooting, perform the following command for the platform on which you are running:

Platform	Perform the following:
Windows	<p>Select Start &gt; Programs &gt; MarkLogic Server &gt; Start MarkLogic Server.</p> <p><b>Note:</b> When you start MarkLogic Server from the Start menu, the Windows service configuration for MarkLogic Server is set to start automatically. Also, if you are using Windows Vista, to start the service you must right-click the Start MarkLogic Server link in the Start menu and choose Run as Administrator, then choose to allow the action.</p>
Red Hat Linux	<p>As the <code>root</code> user, enter the following command:</p> <pre data-bbox="462 940 893 972">/etc/init.d/MarkLogic start</pre>
Sun Solaris	<p>As the <code>root</code> user, enter the following command:</p> <pre data-bbox="462 1073 893 1104">/etc/init.d/MarkLogic start</pre>

This starts a demo HTTP Server running on port 8000, and the Admin Interface HTTP Server running on port 8001.

## 2.5 Entering the License Key and Accepting the License Agreement

If you are upgrading from MarkLogic Server 4.0 or an earlier 4.1 release, there is no need to enter a new license key for 4.1; the old key will continue to work in 4.1 for the features for which you are licensed. In that case, you can skip this section.

You must enter a valid license key before you can execute any queries in MarkLogic Server. If you downloaded MarkLogic Server from [developer.marklogic.com](http://developer.marklogic.com), you can obtain a license key through the License Key Request screen in the Admin Interface. If you are a Mark Logic customer and obtained the release from support, you can obtain your license key from Mark Logic support.

To enter a license key, perform the following steps:

1. If you have not already done so, start MarkLogic Server as described in the previous section.

2. Open the Admin Interface in a browser. For example, if you are running your browser on the computer in which MarkLogic Server is running, open the following URL in a browser:

<http://localhost:8001/>

3. The License Key Entry page appears.

The screenshot shows a web form titled "License Key Entry" with a light beige background. At the top, it says "Please enter the license key for host: dsokolsky-It.marklogic.com". Below this are two input fields: "Licensee" and "License Key". The "Licensee" field has a placeholder text "The licensee." and the "License Key" field has a placeholder text "The license key for this host." and a red error message "You must enter a key". Below the fields is an "ok" button. A horizontal line separates this section from the "Get a License Key" section. This section contains instructions: "You must obtain a valid license key from Mark Logic Corporation to run MarkLogic Server." It then offers two options: "If you have purchased the software but do not know your license key, please contact your sales representative." with a "purchased" button, and "You may request a free, restricted-use license key for personal use, or for evaluation purposes." with a "free" button.

4. If you already have a license key, enter the name of the licensee in the Licensee field and the key in the License Key field. Skip to step [12](#) on page 22.
5. If you have purchased the software but do not have your license key, contact your sales representative or click the Purchased button to send an email.

6. If you want to obtain a free, restricted-use license key for personal use or evaluation, click the Free button.

The free license has capacity and usage restrictions; see the license agreement for details.

After clicking the Free button, the License Key Request screen appears.

**License Key Request**

Hostname: denise.cerisent.com

Please provide the following information to request the free license key.

You may use this form to request a free, restricted-use license key for MarkLogic Server. There are two license options available for you to choose from, Community and Trial. The following fields apply to both license types.

Required fields are denoted with an asterisk (\*).

**Licensee\***   
Your name

**Company\***   
Your company name ('None' if personal)

**Email\***   
Your email address

**Website**   
Your website address

**How did you hear about us?**

**Community License**

The Community License is available free of charge, is not time limited and may be used in production for personal projects. It is limited to a single CPU server and imposes a data set limit of 50 megabytes. The Community License is restricted two copies per company. Support is provided informally through xqzone.marklogic.com.

7. On the License Key Request form, enter your name in the Licensee field, company name in the Company field, and a valid email address in the Email field.

**Note:** The name of the Licensee you enter is associated with the license key that is issued.

8. Decide whether you want a Community license or a Trial license. The Community license does not expire, is limited to 10 gigabytes of content (including index space), and is limited to single- or dual-processor machines. The Trial license allows you to have 40 gigabytes of content, but is limited to 30 days.

9. If you decide on the trial license, fill out the information in the Trial License section.

**Trial License**

For those who need a larger system than the Community License provides but want to get started right away, we provide a 30 day free Trial License. The Trial License is limited to 1 gigabyte of content and can be run on a dual CPU server. The trial license allows Content Interaction Server to be used for development and testing purposes only. Support is provided informally through [xqzone.marklogic.com](http://xqzone.marklogic.com).

Please provide the additional information below about your Company.

**Title\***   
Your title at the company

**Company Address\***   
Street address of your company

**City, State, ZIP\***   
Anytown, CA 98765

**Phone Number\***   
Your contact phone number

**Usage\***   
How do you plan to use the software?

**Employees**   
How many employees in your company?

**Which best describes your organization?**

---

To obtain a full evaluation or production license without the above restrictions, please contact [sales@marklogic.com](mailto:sales@marklogic.com).

Please visit <http://xqzone.marklogic.com> for more information on installing and using MarkLogic Server.

10. Click the Select License button for either the Community or Trial license.
11. The License Key Generation screen appears. Click the OK button.
- The License Key Entry screen appears with your newly-generated key entered in the License Key field. Make sure the Licensee name matches the one you entered in the License Key Request form.
12. Click the OK button.
- If the license key is valid, the server restarts and then the Mark Logic License Agreement screen appears.
13. On the Mark Logic License Agreement screen, read all terms and conditions of the Mark Logic License Agreement.


14. Click Accept to accept the license agreement.

**Note:** Accepting this license agreement is optional. However, to begin using MarkLogic Server, you must accept the terms and conditions outlined in the license agreement. If you have executed a written software license with Mark Logic Corporation, the agreement displayed references that written agreement.

## 2.6 Configuring a New Installation

This section only applies if you are installing MarkLogic Server on this host for the first time. It does not apply if you are upgrading from a previous release of MarkLogic Server.

1. After you accept the license agreement, wait for the server to restart.
2. MarkLogic Server then needs to self-install the initial databases and applications. The Server Install page appears.



3. Click OK to continue.
4. Wait for the server to restart.
5. If you are using Enterprise Edition (in a single-node or a cluster configuration), skip to “Enterprise Edition: Configuring the First and Subsequent Hosts” on page 25.

6. If you are using Standard Edition, the server prompts you to create an admin user. Enter the login name and password for the admin user.

**Security Setup**

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

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

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

**Note:** If you plan on using the digest or digestbasic authentication scheme, decide on the realm you want to use now and enter it in the Realm field. Changing the realm later will invalidate all digest passwords in the security database. For more information on security realms, see *Understanding and Using Security*.

7. Click OK to continue. This creates an admin user.
8. The server prompts you for a username and password; enter the admin username and password you just created.

After logging in with your admin username and password, the browser displays the Admin Interface. MarkLogic Server is now installed. To verify your installation, go to “Checking for the Correct Software Version” on page 30.

**Note:** The installation process also creates a second user named `nobody`. The `nobody` user that the installation process creates is assigned a randomly-generated password and is given no privileges. If you want to add privileges or change the password of the `nobody` user, use the Admin Interface to modify the user, as described in the *Administrator’s Guide*.

## 2.7 Enterprise Edition: Configuring the First and Subsequent Hosts

The following configuration procedures are for Enterprise Edition only, and are different depending on if you run MarkLogic Server in a cluster configuration or on a single host. The Enterprise Edition only procedures are as follows:

- [Enterprise Edition: Configuring a Single Host or the First Host in a Cluster](#)
- [Enterprise Edition: Configuring an Additional Host in a Cluster](#)
- [Enterprise Edition: Leaving a Cluster and Becoming a Single Host](#)

If you are configuring MarkLogic Server as a standalone host, or if this is the first host in a cluster configuration, follow the installation instructions in Section [2.7.1](#). Otherwise, follow the installation instructions in Section [2.7.2](#).


If you are upgrading a cluster to a new release, see [Upgrading a Cluster to a New Maintenance Release of MarkLogic Server](#) in the *Scalability, Clustering, and Forest-Level Failover*.

### 2.7.1 Enterprise Edition: Configuring a Single Host or the First Host in a Cluster

This section only applies to Enterprise Edition installations.

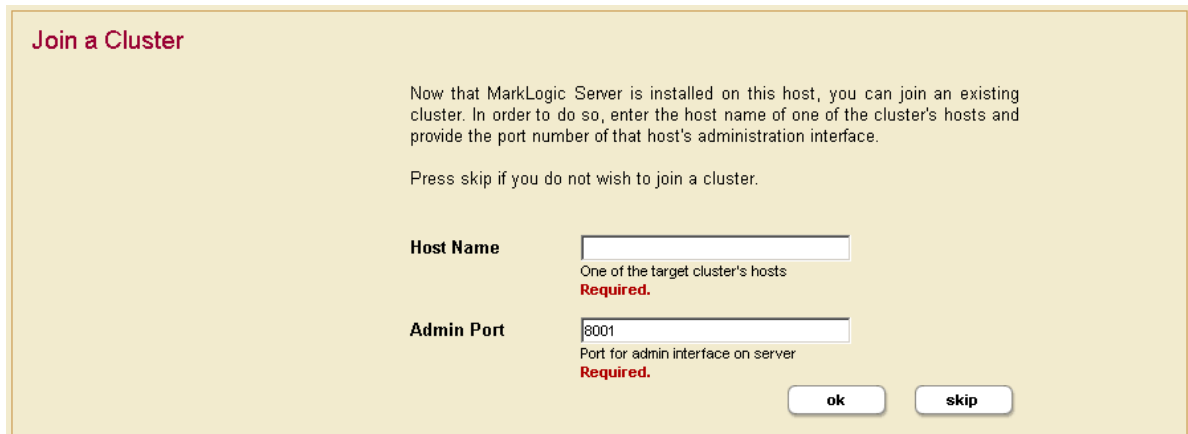
To configure this installation as a single host, or as the first host in a cluster, perform the following steps:

1. Install MarkLogic Server, start MarkLogic Server, and enter your license key, as described in “Installing MarkLogic Server” on page 16, “Starting MarkLogic Server” on page 19, “Entering the License Key and Accepting the License Agreement” on page 19.
2. After accepting the license agreement, MarkLogic Server restarts. Then MarkLogic Server needs to self-install the initial databases and applications. The Server Install page appears.



3. Click OK to continue.
4. Wait for the server to restart.

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



**Join a Cluster**

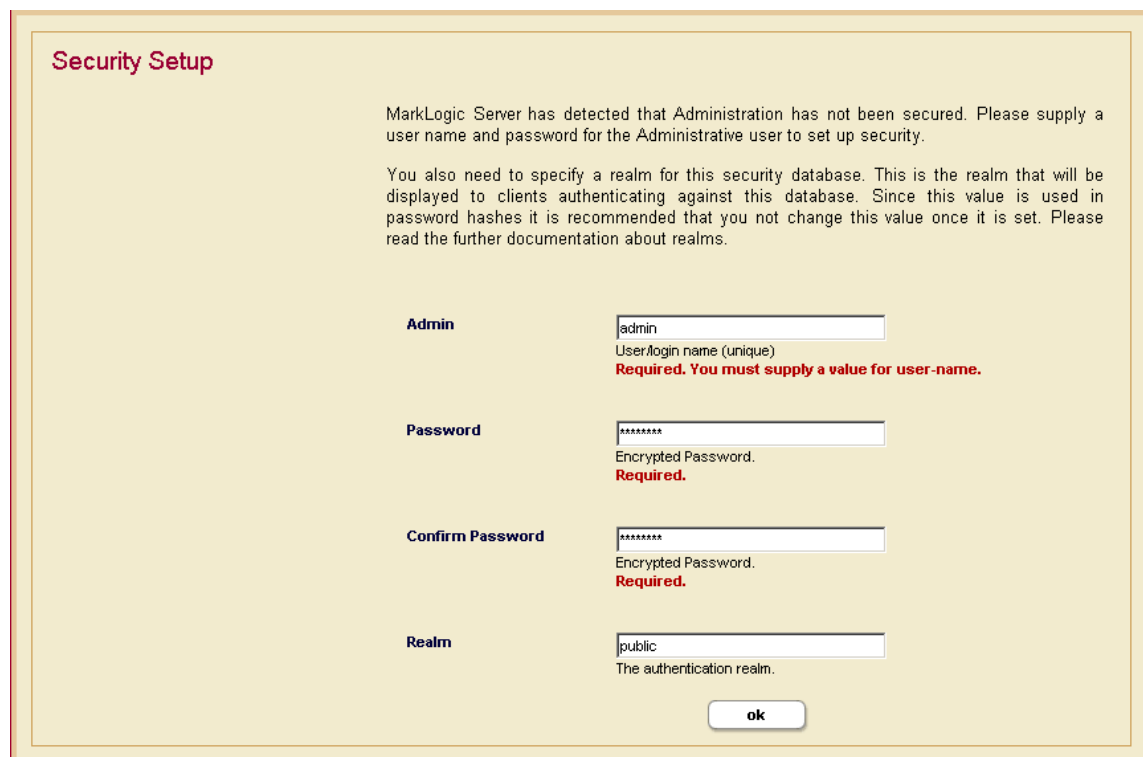
Now that MarkLogic Server is installed on this host, you can join an existing cluster. In order to do so, enter the host name of one of the cluster's hosts and provide the port number of that host's administration interface.

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

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

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

6. Click Skip.
7. You will be prompted to create an admin user. Enter the login name and password for the admin user.



**Security Setup**

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

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

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

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

**Confirm Password**   
Encrypted Password.  
**Required.**

**Realm**   
The authentication realm.

8. Click OK.
9. You will be prompted to log in with your admin username and password.


You will now see the Admin Interface. If you do not need to add any hosts at this time, skip to [Section 2.8](#) on page 30.

## 2.7.2 Enterprise Edition: Configuring an Additional Host in a Cluster

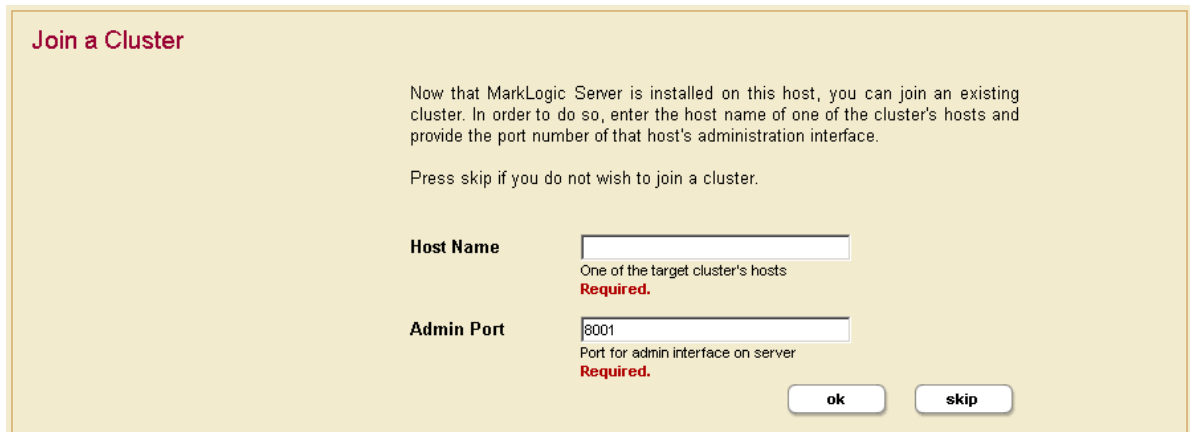
This section only applies to Enterprise Edition installations with two or more hosts.

All hosts in a cluster have to be on the same platform. To configure this installation as an additional host in a cluster of the same platform, perform the following steps:

1. On the node you want to add to an existing cluster, install MarkLogic Server, start MarkLogic Server, and enter your license key, as described in “Installing MarkLogic Server” on page 16, “Starting MarkLogic Server” on page 19, “Entering the License Key and Accepting the License Agreement” on page 19.
2. After accepting the license agreement, MarkLogic Server restarts. Then MarkLogic Server needs to self-install the initial databases and applications. The Server Install page appears.



3. Click OK to continue.
4. Wait for the server to restart.
5. After the server restarts, you will be prompted to join a cluster.



6. Enter the DNS name or the IP address of one of the machines in the cluster. For instance, if this is the second host you are installing, you can enter the DNS name of the first host you installed.
7. Click OK.
8. You will be prompted for an admin username and password. You can use the admin username and password you created when installing the first host. Click OK.

9. Select a Group to assign this host. Click OK.

**Join a Cluster**

In order to join this MarkLogic cluster, the new host must be assigned to one of the cluster's groups and its name confirmed. Please select a group to join and confirm the new host's name:

**Group**  The groups in the cluster

**Host Name**  The name of the new host in the cluster.

10. Click OK to confirm that you are joining the cluster.

**Join a Cluster**

You are about to join this MarkLogic cluster. To proceed click OK.

11. You have now joined the cluster.

**Joined a Cluster**

**Joined a Cluster**  
**MarkLogic Server**

A new host has joined the MarkLogic cluster. Press ok to transfer cluster configuration information to the new host.

12. Click OK to transfer the cluster configuration information.

You have completed the process to join a cluster and will now see the Admin Interface.

### 2.7.3 Enterprise Edition: Leaving a Cluster and Becoming a Single Host

This section only applies to Enterprise Edition installations.

If your host is currently in a cluster of multiple hosts, and you would like to leave the cluster and switch to a single host environment, follow the steps in this section.

When you are switching from a clustered environment to a single-host environment, the forest assignments on the host will be dropped. However, forest data will remain on the file system and you can add the forest back to the host after changing the configuration. For instructions on adding a forest to a host, see the *Administrator's Guide*.

Perform the following steps to leave the cluster to which a host is connected.

1. Run the Admin Interface from the host you want to remove from the cluster.
2. Click the Hosts icon in the left menu tree. The Host Summary page appears.
3. Click the name of the host you want to remove from the cluster, either from the left menu tree or from the Host Summary page. The Host Configuration page appears:

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

**Note:** The Leave button only appears if the Admin Interface is running from this host.

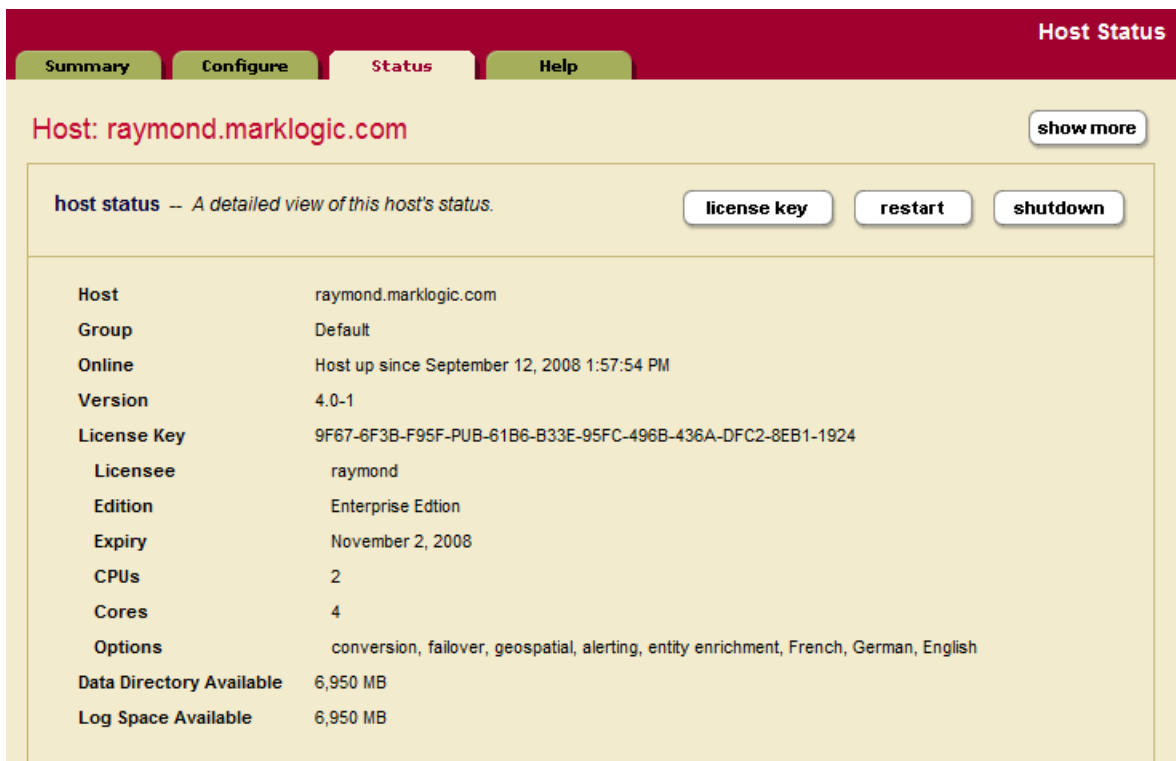
4. Click the Leave button
5. Click OK to confirm leaving the cluster.
6. The host restarts to load the new configuration.
7. Follow the instructions in sections “Configuring a Single Host or the First Host in a Cluster” or “Configuring an Additional Host in a Cluster” as appropriate.

## 2.8 Checking for the Correct Software Version

After logging in with your admin username and password, the Admin Interface appears. In the left corner of the Admin Interface, the version number and product edition are displayed.

To view more details about the release of MarkLogic Server that is installed and licensed, complete the following steps:

1. Click the Hosts icon on the left tree menu.
2. Select the name of the host you just installed, either from the left menu tree or from the Host Summary page.
3. Click the Status tab. The Host Status page appears.
4. Check that `<version>` is correct.



The screenshot shows the 'Host Status' page in the MarkLogic Admin Interface. The page has a dark red header with 'Host Status' on the right and navigation tabs for 'Summary', 'Configure', 'Status', and 'Help'. Below the header, the host name 'raymond.marklogic.com' is displayed, along with a 'show more' button. The main content area is titled 'host status -- A detailed view of this host's status.' and contains three buttons: 'license key', 'restart', and 'shutdown'. Below this is a table of host details:

Host	raymond.marklogic.com
Group	Default
Online	Host up since September 12, 2008 1:57:54 PM
Version	4.0-1
License Key	9F67-6F3B-F95F-PUB-61B6-B33E-95FC-496B-436A-DFC2-8EB1-1924
Licensee	raymond
Edition	Enterprise Edition
Expiry	November 2, 2008
CPUs	2
Cores	4
Options	conversion, failover, geospatial, alerting, entity enrichment, French, German, English
Data Directory Available	6,950 MB
Log Space Available	6,950 MB

To begin using MarkLogic Server, see the following document:

- *Getting Started With MarkLogic Server*

Otherwise, you are finished with the Admin Interface for now. You have successfully installed MarkLogic Server on your system.

## 2.9 Configuring MarkLogic Server on UNIX Systems to Run as a Non-daemon User

On UNIX-based systems (Linux and Solaris), MarkLogic Server runs as the UNIX user named `daemon`. This section describes how to change a configuration to run as a different named UNIX user. This procedure must be run by the `root` user. Additionally, the `root` user is still required for installing and uninstalling MarkLogic Server and for starting and stopping MarkLogic Server from the startup scripts.

To modify an installation to run as a user other than `daemon`, perform the following steps:

1. In a command window on the machine in which you installed MarkLogic Server, log in as the `root` user.
2. Make sure MarkLogic Server is stopped. If it is still running, stop it as follows:

Platform	Perform the following to stop MarkLogic Server:
Red Hat Linux	As the <code>root</code> user, enter the following command:  <code>/etc/init.d/MarkLogic stop</code>
Sun Solaris	As the <code>root</code> user, enter the following command:  <code>/etc/init.d/MarkLogic stop</code>

3. Edit the configuration file for your platform using a text editor such as `vi`.

Platform	Configuration File to Edit
Red Hat Linux	<code>/etc/sysconfig/MarkLogic</code>
Sun Solaris	<code>/etc/MarkLogic.conf</code>

4. In the file, edit the `MARKLOGIC_USER` environment variable to point to the user in which you want MarkLogic Server to run. For example, if you want it to run as a user named `raymond`, change the following line:

```
MARKLOGIC_USER=daemon
```

to the following:

```
MARKLOGIC_USER=raymond
```

5. Save the changes to the `/etc/sysconfig/MarkLogic` or `/etc/MarkLogic.conf` file.

6. If you have not yet started MarkLogic Server after performing a clean installation (that is, after installing into a directory where MarkLogic Server has never been installed), then you are done and you can skip the rest of the steps in this procedure. If you have an existing installation (for example, if you are upgrading to a maintenance release), then continue with the following steps.
7. For all of the MarkLogic Server files owned by `daemon`, you need to change the owner to the new user. This includes all forest data and all of the configuration files. By default, the forest data is in the following directories:

Platform	Default Data Directory (for configuration and log files, and default forest directory)
Red Hat Linux	<code>/var/opt/MarkLogic</code>
Sun Solaris	<code>/var/opt/MARKlogic</code>

For example, on a Linux system, perform a command similar to the following, which changes the owner to the user specified earlier in the `/etc/sysconfig/MarkLogic` file:

```
chown -R raymond /var/opt/MarkLogic
```

8. Make sure to change the owner for all forests in the system, otherwise forests will fail to mount upon startup. Note that the above command only changes the owner for forests installed in the default directory. You need to run a similar command on the data directory for each forest in which a data directory is specified.
9. When you have completed all the file and directory ownership changes, start MarkLogic Server as described in “Starting MarkLogic Server” on page 19.

Once you have performed this procedure, all new files created by MarkLogic Server are created with the new user ownership; there will be no need to change any ownership again. Also, the configuration changes you made in this procedure will not be changed during an upgrade of MarkLogic Server. If you perform a clean installation (not an upgrade installation), however, you will need to run this procedure again.

## 2.10 Removing MarkLogic Server

To remove MarkLogic Server from your system, complete the following steps:

1. Stop MarkLogic Server by performing the following action based on the platform in which you are running:

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

2. Once the server is stopped, you can uninstall MarkLogic Server package by performing the following action based on the platform in which you are running:

Platform	Perform the following:
Windows	Use the Add/Remove Programs Control Panel to uninstall MarkLogic Server.
Red Hat Linux	As the <code>root</code> user, enter the following command:  <code>rpm -e MarkLogic</code>
Sun Solaris	As the <code>root</code> user, enter the following command:  <code>pkgrm MARKlogic</code>

**Note:** Using this procedure to remove MarkLogic Server from your system will not remove user data (configuration information, XQuery files used by HTTP or XDBC servers, or forest content). This data is left in place to simplify the software upgrade process. If you wish to remove the user data, you must do so manually using standard operating system commands.

## 3.0 Technical Support

Mark Logic provides technical support according to the terms detailed in your Software License Agreement. For evaluation licenses, Mark Logic may provide support on an “as possible” basis.

We invite you to visit our support website at <http://support.marklogic.com> to access our full suite of documentation and help materials.

If you have questions or comments, you may contact Mark Logic Technical Support at the following email address:

[support@marklogic.com](mailto:support@marklogic.com)

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