



Administration Documentation

Nuxeo Platform 5.5

Table of Contents

1. Installation and Administration	4
1.1 Hardware and Software requirements	4
1.1.1 Supported application servers	7
1.1.2 Supported databases	7
1.2 Installation	7
1.2.1 Installing the Nuxeo Platform on Mac OS	7
1.2.2 Installing the Nuxeo Platform on Windows	8
1.2.2.1 Installing the Nuxeo Platform as a Windows service	11
1.2.2.2 Running multiple server instances in Windows	16
1.2.3 Installing the Nuxeo Platform on Linux	16
1.2.3.1 Configuring the Nuxeo Platform as a daemon on Debian	18
1.2.4 Deploying Nuxeo on your (Amazon AWS) cloud	20
1.2.5 Installing and setting up related software	22
1.3 Setup	24
1.3.1 Recommended configurations	31
1.3.2 Configuration examples	33
1.3.3 Configuration templates	34
1.3.4 Case Management recommended configurations	35
1.3.5 Configuration parameters index (nuxeo.conf)	37
1.3.6 Making CM and DM work with the same repository	43
1.4 Database	43
1.4.1 Configuring PostgreSQL	43
1.4.2 Configuring Oracle	50
1.4.3 Configuring MS SQL Server	53
1.4.4 Connecting Nuxeo to the database	55
1.5 Authentication, users and groups	58
1.5.1 Using a LDAP directory	60
1.5.2 Using CAS2 authentication	65
1.5.3 Using OAuth	67
1.5.4 Using Shibboleth	70
1.6 HTTP and HTTPS reverse-proxy configuration	75
1.7 Advanced configuration	78
1.7.1 Adding custom templates	78
1.7.2 Changing context path	78
1.7.3 Configure User & Group storage and Authentication	79
1.7.4 Firewall consideration	79
1.7.5 How to create a Nuxeo static WAR?	79
1.7.6 JDBC	80
1.7.7 Nuxeo clustering configuration	80
1.7.8 VCS Configuration	82
1.8 Server start and stop	87
1.8.1 nuxeoctl and Control Panel usage	89
1.9 Monitoring	91
1.9.1 Nuxeo Management	93
1.9.2 Remote monitoring through a SSH tunnel	97
1.9.3 Transactions and connections	98
1.10 Nuxeo Shell (admin's best friend)	99
1.10.1 Shell Command Index	102
1.10.1.1 Built-in Commands	102
1.10.1.2 File System Commands	105
1.10.1.3 Nuxeo Server Commands	107
1.10.1.4 Nuxeo Automation Commands	118
1.10.1.5 Configuration Commands	141
1.10.2 Shell Batch Mode	142
1.11 Backup and restore	144
1.12 Import a file system	144
1.13 Upgrading your Nuxeo Version	145
1.13.1 Upgrade from 5.4.2 to 5.5	145
1.13.2 Upgrade to 5.4.2 with Oracle	149
1.13.3 Upgrade from 5.3.2 to 5.4.0	150
1.13.3.1 From the old workflow system to the new 5.4 workflow system	150
1.13.3.2 Upgrade to 5.4 and JBoss 5	150
1.13.4 Upgrade from 5.3.1 to 5.3.2	156
1.13.5 Upgrade from 5.3.1 with MySQL to 5.3.2	157
1.13.6 Upgrade from 5.3.0 to 5.3.1	159
1.13.7 Upgrade from 5.1.6 with JCR + PostgreSQL to 5.2.0	161
1.13.8 Upgrade from 5.1.2 to 5.1.3	162

1.14 Marketplace addons	162
1.14.1 Amazon S3 Online Storage	163
1.14.2 Nuxeo GSA Connector	165
1.15 Purge audit logs	170

Installation and Administration

Installation and Administration Guide for Nuxeo Platform 5.5

Welcome to the Nuxeo Platform installation and administration guide.

In this guide, you will find all the information to install and manage the Nuxeo Platform: how to install it, how to configure a database, install new packages from the Nuxeo Admin Center etc. The installation and administration principles described in this guide apply to all modules of the Platform: document management, digital asset management, case management, etc.

[Download](#)



[Download this documentation in PDF.](#)

Quick start – For evaluation purpose

You want to evaluate or test the platform? Here the main steps you should follow to quickly install your Nuxeo application and get ready to use it.

1. Check out the requirements
2. Install
3. Start
4. Setup the platform with a preset module



License

This documentation is copyrighted by Nuxeo and published under the Creative Common BY-SA license. More details on the [Nuxeo documentation license](#) page.

Full installation – For advanced testing and production purpose

You want to install the application to use it or test it in a production environment? Follow the steps below:

- Hardware and Software requirements
- Installation
- Setup
- Database
- Authentication, users and groups
- HTTP and HTTPS reverse-proxy configuration
- Advanced configuration
- Server start and stop
- Monitoring
- Nuxeo Shell (admin's best friend)
- Backup and restore
- Import a file system
- Upgrading your Nuxeo Version
- Marketplace addons
- Purge audit logs

Hardware and Software requirements

This section presents information about the running environment for a Nuxeo server. Listing all required software, giving a recommended configuration and listing some others, known as operational, this sections aims at helping you to validate or define your production environment. However the list is not exhaustive and needs to be completed with the users' experience.

On this page

- Requirements
 - Checking your Java version
 - Installing Java
 - For Windows users
 - For Linux users
 - For Mac OS X users
- Recommendations
 - Hardware configuration
 - Default configuration
 - For optimal performances
 - Known working configurations
 - OS
 - JVM
 - Storage backends
 - LDAP
 - Browsers

Requirements

The Nuxeo Platform can run on Linux, Mac OS X and Windows operating systems.

All you need for a quick start is a **Sun Java Development Kit (JDK)** (*a JRE is not enough*). **Java 6** (also called Java 1.6) is required.



We currently support Sun's JDK, don't hesitate to contact us if you need us to support a JDK from another vendor.

Checking your Java version

To check that you have the right version of Java:

1. Open a terminal:
 - on Linux or Mac OS X: open a terminal.
 - on Windows: press "windows" key + r, type cmd (or command) in the Run window and press "OK" or open the "Prompt command" from "Start > Programs > Accessories" menu.
2. Type java -version and press **Enter**.
If Java is correctly installed on your computer, the name and version of your Java virtual machine is displayed:

```
$ java -version
java version "1.6.0_20"
Java(TM) SE Runtime Environment (build 1.6.0_20-b02-279-10M3065)
Java HotSpot(TM) 64-Bit Server VM (build 16.3-b01-279, mixed mode)
```

If Java is not installed on your computer, it fails to display the Java version. Then, you need to install Java (see below).

If Java is installed but not included in the PATH, it fails to find the Java command. Then, you need to add \$JAVA_HOME/bin/ in your PATH (see [How do I set or change the PATH system variable?](#)).

Installing Java

For Windows users

If the required version of Java is not installed on your computer:

1. Download it from [Sun website](#) and choose the appropriate platform for your hardware and Windows version.
2. Run the downloaded .exe file and follow the instructions displayed.

For Linux users

On a modern Linux distribution, you can now install Sun Java 6 with the packaging system of your distribution (using rpm, apt, yum, etc.).

If the JDK from Sun is not available for your OS, you can still [download it from Sun website](#).



 Since Ubuntu 10.4 (Lucid Lynx), sun-java-6 is available from partners software sources (apt repositories).

For Mac OS X users

An appropriate version of Java should already be installed on your computer if you are using Snow Leopard (Java 6).

Recommendations

Hardware configuration

The Nuxeo Platform is designed to be scalable and can thus be deployed on many servers. It can be installed on only one server for a start, and can also easily be installed on many servers. The constant is that there is the need to have one modern server with good performances. Then the other servers can be more lower-end.

The numbers below are given for the one needed high-end server.

- RAM: 2 GB is the minimum requirement for using Nuxeo,
- CPU: Intel Core 2 or equivalent and better.
You might want to avoid machines from the Intel Pentium 4 Xeon series since some models have a too small amount of cache. This impairs performance greatly compared to other CPU architecture of the same generation. (Intel Pentium 4 servers are quite widespread because of an attractive price policy.)
- Storage (disk) space: the minimum Nuxeo installation, along with the needed server and libs, takes something between 200 MB and 280 MB on a filesystem. Then, the final size will of course depend on the amount of data that will be stored in Nuxeo. A safe bet (until we provide better numbers) is to consider data space ratio of 1.5 to 2.

Default configuration

The default persistence configuration is lightweight and easy to use, but it is not made for performance.

The Nuxeo Platform uses:

- H2 for SQL Data (directories, JBPM, Relations ...),
- Filesystem persistence with [VCS](#) for the Document repository.

For optimal performances

- Linux 64 bits,
- PostgreSQL 8.4 or 9.0,
Use PostgreSQL for document repository and all other services.



With Nuxeo EP 5.1 or 5.2, configure the document repository to externalize the blobs to filesystem.

- Have plenty of RAM (>= 4 GB).

Known working configurations

OS

- Debian GNU/Linux 5.0 Lenny or more recent
- Linux Ubuntu 32 and 64 bits: 10.10 or more recent
- Linux Mandriva 2008.1
- Red Hat Linux RHEL 5 and 6
- CentOS 5
- OpenSUSE
- Other Unix variants, as long as there is an implementation of Java 6 (such as Solaris)
- Mac OS X Leopard (10.5), Snow Leopard (10.6)
- Microsoft Windows 2000, Windows 2003 server 32 and 64 bits, Windows XP, Windows 7

JVM

- Sun JDK 6, 64 bits recommended especially on Windows environment.

Storage backends

Different backends may be set as well for Nuxeo Core repository as for all other Nuxeo services that persist data. Please see the [list of supported databases](#) for each version of Nuxeo.

LDAP

- OpenLDAP
- OpenDS
- Microsoft Active Directory

Browsers

Nuxeo applications can be used with the browsers below.

- IE 7 and greater
- Firefox 3.5 and greater
- Google Chrome 3 and greater
- Safari 4 and greater



Browser extensions for Drag & Drop and Live Edit are available for Internet Explorer and Firefox only.

Supported application servers

The page Supported application servers does not exist.

Supported databases

The page Supported databases does not exist.

Installation

The Nuxeo Platform comes in different packages and can be installed on any operating system. You may have to install:

- a zip archive (works on any operating system),
- a Windows installer (.exe),
- a virtual machine image (works on any operating system),
- a .deb package (works on Linux Debian and Ubuntu).

Our installation recipies:

- [Installing the Nuxeo Platform on Mac OS](#)
- [Installing the Nuxeo Platform on Windows](#)
 - [Installing the Nuxeo Platform as a Windows service](#)
 - [Running multiple server instances in Windows](#)
- [Installing the Nuxeo Platform on Linux](#)
 - [Configuring the Nuxeo Platform as a daemon on Debian](#)
- [Deploying Nuxeo on your \(Amazon AWS\) cloud](#)
- [Installing and setting up related software](#)

Installing the Nuxeo Platform on Mac OS

On Mac OS, you can install the Nuxeo Platform using two different packages:

- the .zip archive,
- the virtual machine image.

How to install the Nuxeo Platform from the .zip archive

Installing the Nuxeo Platform using the .zip package installs the Nuxeo Platform only. External dependencies must be installed separately.

To install the Nuxeo Platform zip archive:

Unzip the .zip archive using your favorite tool.

What's next?

You want to evaluate the application? You can now [start the server](#).

You want to do a complete installation, compatible for a production environment? You should now [prepare your environment](#).

How to install a Nuxeo Virtual machine image

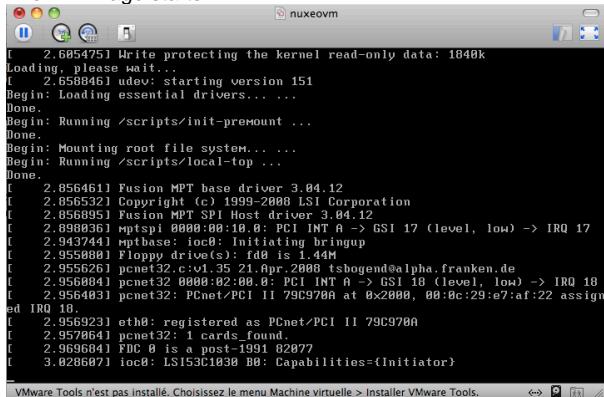
The Nuxeo Platform is available as ready-to-use virtual machine images from [nuxeo.com](#). VM images are available for VMWare and Virtual Box. They provide a full environment (OS, database...) and all required dependencies to make the Nuxeo Platform work.

To install the Nuxeo virtual machine image and start Nuxeo:

1. Unzip the downloaded package.
You get a folder with the required file for the virtual machine image to run.
2. Start the virtual machine image in your virtual machine application by double-clicking on it.

- For the VMWare package, double-click on the file "nuxeo.vmx".
- For the OVF package, double-click on the .ovf file that corresponds to the supported standard: "nuxeo_OVF10.ovf" for Open Virtualization Format 1.0, supported by Virtual Box for instance, or "nuxeo_OVF09.ovf" for Open Virtualization Format 0.9. Then start the imported virtual machine.

The VM image starts.



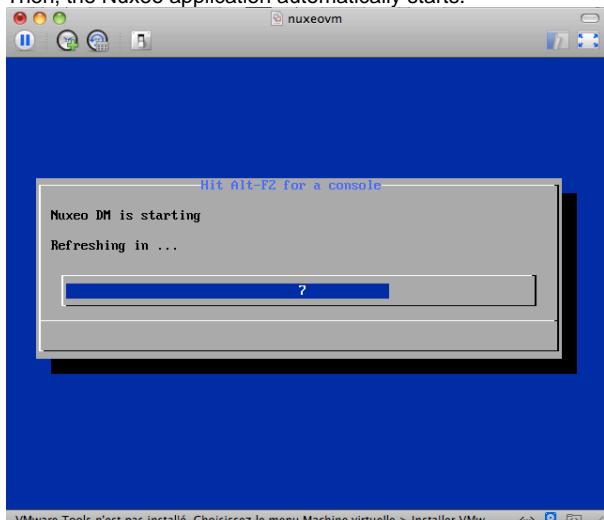
```

[ 2.685475] Write protecting the kernel read-only data: 1840k
Loading, please wait...
[ 2.658846] udev: starting version 151
Begin: Loading essential drivers... ...
Done.
Begin: Running /scripts/init-premount ...
Done.
Begin: Mounting root file system... ...
Begin: Running /scripts/local-top ...
Done.
[ 2.856461] Fusion MPT base driver 3.04.12
[ 2.856523] Copyright (c) 1999-2009 LSI Corporation
[ 2.856095] Fusion MPT SPI Host Driver 3.04.12
[ 2.889836] Mptbase 0000:00:10.0: PCI INT A -> GSI 17 (level, low) -> IRQ 17
[ 2.947744] mphase: ioc8: Initializing bringup
[ 2.955888] Floppy drive(s): fd0 is 1.44M
[ 2.956261] pcnet32.c:wl.35 21 Apr 2008 tsbogend@alpha.franken.de
[ 2.956884] pcnet32 0000:02:00.0: PCI INT A -> GSI 18 (level, low) -> IRQ 18
[ 2.956403] pcnet32: PCnet/PCI II 79C970A at 0x2000, 00:0c:29:e7:af:22 assigned IRQ 18.
[ 2.956923] eth0: registered as PCnet/PCI II 79C970A
[ 2.957864] pcnet32: 1 cards found.
[ 2.969684] FDC 0 is a post-1991 82077
[ 3.028687] ioc8: LSI153C103B BB: Capabilities={Initiator}

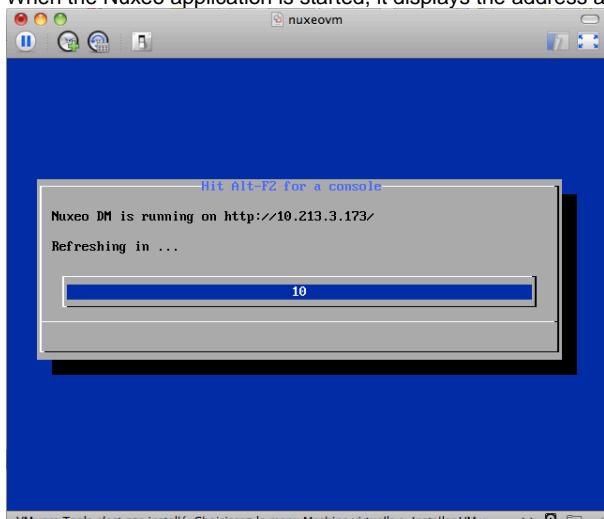
```

VMware Tools n'est pas installé. Choisissez le menu Machine virtuelle > Installer VMware Tools.

Then, the Nuxeo application automatically starts.



When the Nuxeo application is started, it displays the address at which it is available.



3. In your browser, type the indicated address.

The startup wizard is displayed to help you configure your application.



Shell root access

The password for the root and nuxeo users are generated the first time you start the virtual machine and are displayed on the console.

Installing the Nuxeo Platform on Windows

You can install the Nuxeo Platform on Windows using several packages:

- the Windows installer (.exe),
- the .zip archive,
- the virtual machine image.

On this page

- [How to install the Nuxeo Platform from the Windows installer](#)
- [How to install the Nuxeo Platform from the .zip archive](#)
- [How to install a Nuxeo Virtual machine image](#)
- [Related pages:](#)

How to install the Nuxeo Platform from the Windows installer

The Nuxeo Platform is available with a Windows installer that guides you in the install process.

To install the Nuxeo Platform using the Windows installer:

To install the application using the Windows installer (.exe), double-click on the .exe installer you downloaded and follow the instructions displayed.



On Windows 7, because of rights issues, it is highly recommended to install your Nuxeo application at the root of C: in order for your application to restart correctly at the end of the [startup wizard steps](#).



What's next?

You want to evaluate the application? You can now [start the server](#).

You want to do a complete installation, compatible for a production environment? You should now [prepare your environment](#).

How to install the Nuxeo Platform from the .zip archive

Installing the Nuxeo Platform using the .zip package installs the Nuxeo Platform only. External dependencies must be installed separately.

To install the Nuxeo Platform zip archive:

Unzip the .zip archive using your favorite tool.



Because of the limitation on paths length, it is recommended to extract the content of the archive at root of C:.

What's next?

You want to evaluate the application? You can now [start the server](#).

You want to do a complete installation, compatible for a production environment? You should now [prepare your environment](#).

How to install a Nuxeo Virtual machine image

The Nuxeo Platform is available as ready-to-use virtual machine images from [nuxeo.com](#). VM images are available for VMWare and Virtual Box. They provide a full environment (OS, database...) and all required dependencies to make the Nuxeo Platform work.

To install the Nuxeo virtual machine image and start Nuxeo:

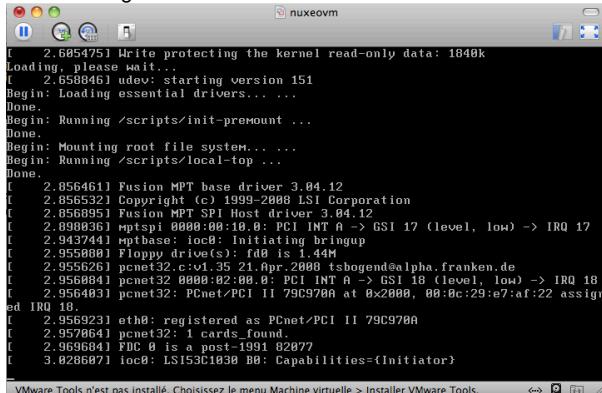
1. Unzip the downloaded package.

You get a folder with the required file for the virtual machine image to run.

2. Start the virtual machine image in your virtual machine application by double-clicking on it.

- For the VMWare package, double-click on the file "nuxeo.vmx".
- For the OVF package, double-click on the .ovf file that corresponds to the supported standard: "nuxeo_OVF10.ovf" for Open Virtualization Format 1.0, supported by Virtual Box for instance, or "nuxeo_OVF09.ovf" for Open Virtualization Format 0.9. Then start the imported virtual machine.

The VM image starts.

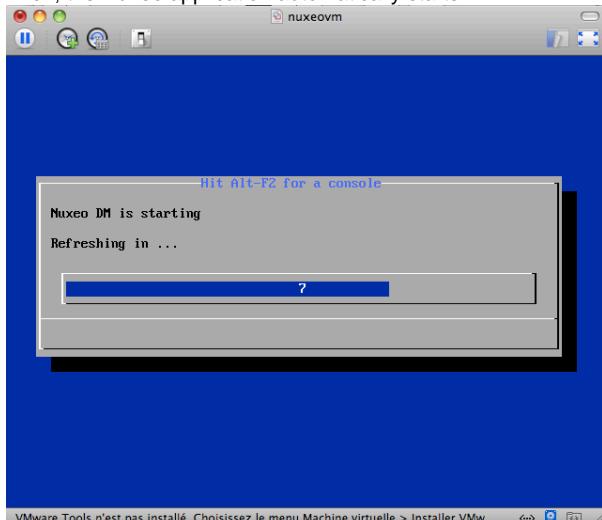


```

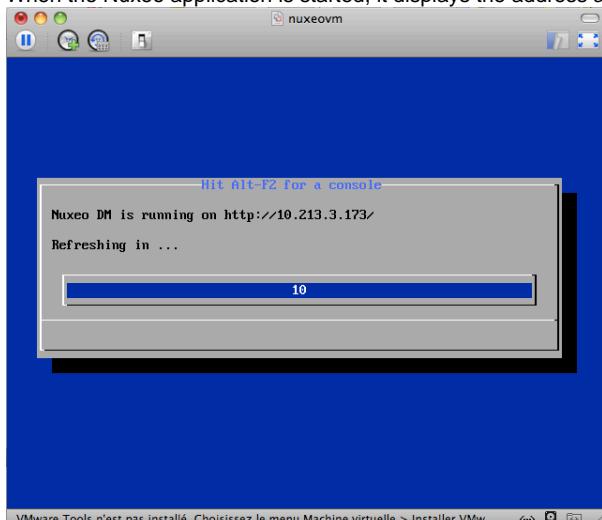
[ 2.685475] Write protecting the kernel read-only data: 1840k
[    Loading, please wait...
[    2.658846] udev: starting version 151
[    Begin: Loading essential drivers... ...
[    Done.
[    Begin: Running /scripts/init-premount ...
[    Done.
[    Begin: Mounting root file system... ...
[    Begin: Running /scripts/local-top ...
[    Done.
[    2.856461] Fusion MPT base driver 3.04.12
[    2.856532] Copyright (c) 1999-2008 LSI Corporation
[    2.856895] Fusion MPT SPI Host driver 3.04.12
[    2.898836] mptscsi 0000:00:10.0: PCI INT A -> GSI 17 (level, low) -> IRQ 17
[    2.943744] mptbase: ioc0: Initiating bringup
[    2.955888] Floppy drive(s): fd0 is 1.44M
[    2.955626] pcnet32.c:v1.30 21.Apr.2008 tsbogend@alpha.franken.de
[    2.956084] pcnet32 0000:02:00.0: PCI INT A -> GSI 18 (level, low) -> IRQ 18
[    2.956403] pcnet32: PCnet/PCI II 79C970A at 0x2000, 00:0c:29:e7:af:22 assigned IRQ 18
[    2.956923] eth0: registered as PCnet/PCI II 79C970A
[    2.957864] pcnet32: 1 cards found.
[    2.969684] FDC 0 is a post-1991 E2077
[    3.028687] ioc0: LS153C1030 00: Capabilities={Initiator}

```

Then, the Nuxeo application automatically starts.



When the Nuxeo application is started, it displays the address at which it is available.



3. In your browser, type the indicated address.

The startup wizard is displayed to help you configure your application.

**Shell root access**

The password for the root and nuxeo users are generated the first time you start the virtual machine and are displayed on the console.

Related pages:

- [Installing the Nuxeo Platform on Windows](#)
- [Running multiple server instances in Windows](#)
- [Installing the Nuxeo Platform as a Windows service](#)

Installing the Nuxeo Platform as a Windows service

Installing Nuxeo as a Windows service is independent of Nuxeo. So, this is no longer in our development scope since Nuxeo 5.4.

Multiple solutions are available, here are some of them, given without any warranty.

Once a batch is installed as a service, it cannot be changed: you must first uninstall it, then edit and reinstall in order to change its content.
So, it's generally a good idea to write a batch file wrapping calls to `nuxeoctl.bat` and install that `NuxeoWrapper.bat` as a service, which will be responsible of starting Nuxeo with the wanted user and environment parameters.

On this page

- Prerequisites
- Available solutions
 - Yet Another Java Service Wrapper (recommended)
 - Installing Nuxeo as a Windows service using YAJSW
 - JBoss Native Windows (aka JBossSVC, JBossService and JavaService)
 - Tomcat Service Install/Uninstall script
 - JavaServiceWrapper by Tanuki
 - .NET InstallUtil
 - Related pages

Prerequisites

In order to run as a service, you have to manage the directory rights for the super-user running the service. There are behavior changes depending on the Windows version.

Also, take care that network directories are usually not available when a service is executing. So, if you need to use some, you will have to mount them in the batch script before starting Nuxeo.

The database used by Nuxeo has to be installed as a service and started before the Nuxeo service.

Available solutions

Yet Another Java Service Wrapper (recommended)

YAJSW is a Java centric implementation of the Java Service Wrapper by tanuki (JSW). It aims at being mostly configuration compliant with the original. YAJSW is LGPL licensed. That solution seems to be the more flexible, robust and multi-OS compliant.

Installing Nuxeo as a Windows service using YAJSW

1. Download YAJSW and unzip the archive.
2. Set the system environment variable `NUXEON_CONF` to the location of your `nuxeo.conf` file, something like `%NUXEON_HOME%\bin\nuxeo.conf`.
3. Start Nuxeo DM from the command line:

```
nuxeoctl.bat nogui start
```

4. Once the server is started, you'll get a message like below where XXXX is the process ID of the running Nuxeo application:

```
Server started with process ID XXXX.
```

5. Start a Command Prompt as an Administrator.
6. Go to the %YAJSW_HOME%\bat folder.
7. Execute the genConfig command with the process ID as parameter:

```
genConfig.bat XXXX
```

The configuration is written in the file %YAJSW_HOME%\conf\wrapper.conf.

8. Stop Nuxeo DM:

```
nuxeoctl.bat nogui stop
```

9. Execute your wrapped application as console application by calling this command and check your application is accessible:

```
runConsole.bat
```

10. Edit the file %YAJSW_HOME%\conf\wrapper.conf and set your custom values for these parameters:

```
# Name of the service
wrapper.ntservice.name=NuxeoDM
# Display name of the service
wrapper.ntservice.displayname=Nuxeo DM
# Description of the service
wrapper.ntservice.description=Service to manage Nuxeo DM
```

11. To install the application as service call, execute:

```
installService.bat
```

Your service is installed and you can run Nuxeo DM from its service ("Windows Computer Management > Services" on Windows 7).

JBoss Native Windows (aka JBossSVC, JBossService and JavaService)

Deprecated Nuxeo scripts managing install as a Windows service were previously used. They were based on [JBoss Native Windows](#) which is now not recommended because of a number of defects. However, it was relatively easy to use and provides a quick solution.

As an example, here is the content of jboss-native-2.0.4/bin/service.bat:

Nuxeo JBoss Service Script for Windows

```
@echo off
REM JBoss, the OpenSource webOS
REM
REM Distributable under LGPL license.
REM See terms of license at gnu.org.
REM
REM -----
REM JBoss Service Script for Windows
REM -----
```

```
@if not "%ECHO%" == "" echo %ECHO%
@if "%OS%" == "Windows_NT" setlocal
set DIRNAME=%CD%

REM
REM VERSION, VERSION_MAJOR and VERSION_MINOR are populated
REM during the build with ant filter.
REM
set SVCNAME=NuxeoEP
set SVCDISP=NuxeoEP
set SVCDESC=Nuxeo 5.3.0-GA / JBoss Application Server 4.2.3 GA / Platform: Windows 64
set NOPAUSE=Y

REM Suppress killing service on logoff event
set JAVA_OPTS=-Xrs

REM Figure out the running mode

if /I "%1" == "install" goto cmdInstall
if /I "%1" == "uninstall" goto cmdUninstall
if /I "%1" == "start" goto cmdStart
if /I "%1" == "stop" goto cmdStop
if /I "%1" == "restart" goto cmdRestart
if /I "%1" == "signal" goto cmdSignal
echo Usage: service install^|uninstall^|start^|stop^|restart^|signal
goto cmdEnd

REM jbosssvc return values
REM ERR_RET_USAGE      1
REM ERR_RET_VERSION    2
REM ERR_RET_INSTALL     3
REM ERR_RET_REMOVE      4
REM ERR_RET_PARAMS      5
REM ERR_RET_MODE        6

:errExplain
if errorlevel 1 echo Invalid command line parameters
if errorlevel 2 echo Failed installing %SVCDISP%
if errorlevel 4 echo Failed removing %SVCDISP%
if errorlevel 6 echo Unknown service mode for %SVCDISP%
goto cmdEnd

:cmdInstall
jbosssvc.exe -imwdc %SVCNAME% "%DIRNAME%" "%SVCDISP%" "%SVCDESC%" service.bat
if not errorlevel 0 goto errExplain
echo Service %SVCDISP% installed
goto cmdEnd

:cmdUninstall
jbosssvc.exe -u %SVCNAME%
if not errorlevel 0 goto errExplain
echo Service %SVCDISP% removed
goto cmdEnd

:cmdStart
REM Executed on service start
del .r.lock 2>&1 | findstr /C:"being used" > nul
if not errorlevel 1 (
```

```
echo Could not continue. Locking file already in use.
goto cmdEnd
)
echo Y > .r.lock
jbosssvc.exe -p 1 "Starting %SVCDISP%" > run.log
call run.bat -b 0.0.0.0 < .r.lock >> run.log 2>&1
jbosssvc.exe -p 1 "Shutdown %SVCDISP% service" >> run.log
del .r.lock
goto cmdEnd

:cmdStop
REM Executed on service stop
echo Y > .s.lock
jbosssvc.exe -p 1 "Shutting down %SVCDISP%" > shutdown.log
call shutdown -S < .s.lock >> shutdown.log 2>&1
jbosssvc.exe -p 1 "Shutdown %SVCDISP% service" >> shutdown.log
del .s.lock
goto cmdEnd

:cmdRestart
REM Executed manually from command line
REM Note: We can only stop and start
echo Y > .s.lock
jbosssvc.exe -p 1 "Shutting down %SVCDISP%" >> shutdown.log
call shutdown -S < .s.lock >> shutdown.log 2>&1
del .s.lock
:waitRun
REM Delete lock file
del .r.lock > nul 2>&1
REM Wait one second if lock file exist
jbosssvc.exe -s 1
if exist ".r.lock" goto waitRun
echo Y > .r.lock
jbosssvc.exe -p 1 "Restarting %SVCDISP%" >> run.log
call run.bat < .r.lock >> run.log 2>&1
jbosssvc.exe -p 1 "Shutdown %SVCDISP% service" >> run.log
del .r.lock
goto cmdEnd

:cmdSignal
REM Send signal to the service.
REM Requires jbosssch.dll to be loaded in JVM
@if not "%2" == "" goto execSignal
echo Missing signal parameter.
echo Usage: service signal [0...9]
goto cmdEnd
:execSignal
jbosssvc.exe -k%2 %SVCNAME%
goto cmdEnd
```

```
:cmdEnd
```

Other implementations were available from JBoss.



They were licensed under LGPL and so redistributable but there are not fully satisfying.

Tomcat Service Install/Uninstall script

Using the Tomcat distribution of Nuxeo, you will find a `service.bat` script in the `bin` directory that could be adapted to install Nuxeo as a Windows service.

JavaServiceWrapper by Tanuki

Tanuki's library provides [multiple methods](#) for integrating a software as a service on various OS, the easier is to use the `WrapperSimpleApp` helper class to launch the application: see [the example of JBoss installed as a Windows service](#). It requires to unzip the downloaded wrapper file, configure a `wrapper.conf` file pointing to `%NUXEHOME%\bin\nuxeoctl.bat`, then write a `wrapper.bat` file for managing test/install/uninstall:

JavaServiceWrapper usage

```
REM Test:  
wrapper.exe -c %NUXEHOME%\wrapper\wrapper.conf  
  
REM Install:  
wrapper.exe -i %NUXEHOME%\wrapper\wrapper.conf  
  
REM Uninstall:  
wrapper.exe -r %NUXEHOME%\wrapper\wrapper.conf
```

This solution is known to work well but is sadly not redistributable for us because of its GPL/Commercial license.

.NET InstallUtil

.NET framework provides an `InstallUtil.exe` tool for installing/uninstalling services.

InstallUtil usage

```
REM Install  
InstallUtil /i %NUXEHOME%\bin\service.bat  
  
REM Uninstall  
InstallUtil /u %NUXEHOME%\bin\service.bat
```



There are some disadvantages such as failures in case of multiple frameworks installed and forward/backward incompatibilities.

You may want to have a look at <http://msdn2.microsoft.com/en-US/library/system.configuration.install.managedinstallerclass.aspx> for managing that programmatically.

Related pages

- [Installing the Nuxeo Platform as a Windows service](#)
- [Installing the Nuxeo Platform on Windows](#)
- [Running multiple server instances in Windows](#)

Running multiple server instances in Windows



This applies for all Nuxeo products since version 5.3.2.

The location of the `nuxeo.conf` is defined by that order of priority (i.e. first one of those found is used):

- Registry key `HKEY_LOCAL_MACHINE\SOFTWARE\PRODNAME\ConfigFile` with PRODNAME equals "Nuxeo CAP", "Nuxeo DM", "Nuxeo DAM", ...
- Environment variable `NUXEO_CONF`
- "`nuxeo.conf`" file in the working directory
- "`nuxeo.conf`" file on the Desktop
- "`nuxeo.conf`" file in the same location as the (real) `NuxeoCtl.exe` (for versions < 5.4.1) or `nuxeoctl.bat` (for versions 5.4.1).

To launch multiple instances of Nuxeo you'd need to remove the registry key (set up by the Windows installer) and have wrappers around `NuxeoCtl.exe/nuxeoctl.bat` that define different `NUXEO_CONF` environment variables.

Note that you'd also want to have different `nuxeo.data.dir`, `nuxeo.log.dir`, `nuxeo.tmp.dir`, `nuxeo.server.http.port` and `nuxeo.server.tomcat-admin.port` in the two `nuxeo.conf` files (you can set `nuxeo.server.ajp.port` to 0 to disable AJP if you don't use it).

Installing the Nuxeo Platform on Linux

On Linux, you can install the Nuxeo Platform using the packages below:

- the .zip archive,
- the virtual machine image,
- from the APT repository for Debian and Ubuntu.

On this page

- [How to install the Nuxeo Platform from the .zip archive](#)
- [How to install a Nuxeo Virtual machine image](#)
- [How to install the Nuxeo Platform from the APT repository for Debian and Ubuntu](#)
 - [Related pages](#)

How to install the Nuxeo Platform from the .zip archive

Installing the Nuxeo Platform using the .zip package installs the Nuxeo Platform only. External dependencies must be installed separately.

To install the Nuxeo Platform zip archive:

Unzip the .zip archive using your favorite tool.

What's next?

You want to evaluate the platform? You can now [start the server](#).

You want to do a complete installation, compatible for a production environment? You should now [prepare your environment](#).

How to install a Nuxeo Virtual machine image

The Nuxeo Platform is available as ready-to-use virtual machine images from [nuxeo.com](#). VM images are available for VMWare and Virtual Box. They provide a full environment (OS, database...) and all required dependencies to make the Nuxeo Platform work.

To install the Nuxeo virtual machine image and start Nuxeo:

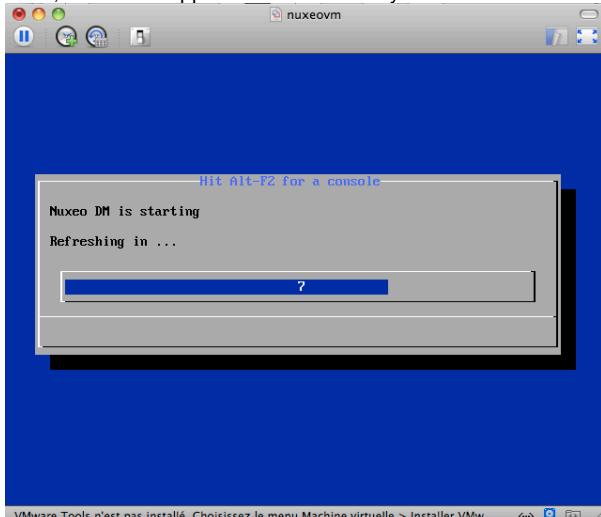
1. Unzip the downloaded package.
You get a folder with the required file for the virtual machine image to run.
2. Start the virtual machine image in your virtual machine application by double-clicking on it.
 - For the VMWare package, double-click on the file "nuxeo.vmx".
 - For the OVF package, double-click on the .ovf file that corresponds to the supported standard: "nuxeo_OVF10.ovf" for Open Virtualization Format 1.0, supported by Virtual Box for instance, or "nuxeo_OVF09.ovf" for Open Virtualization Format 0.9. Then start the imported virtual machine.The VM image starts.

```

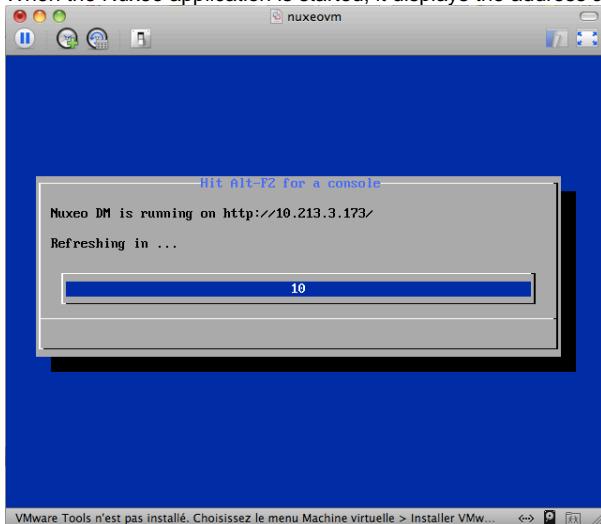
[ 2.685475] Write protecting the kernel read-only data: 1840k
Loading, please wait...
[ 2.658846] udev: starting version 151
Begin: Loading essential drivers... ...
Done.
Begin: Running /scripts/init-premount ...
Done.
Begin: Mounting root file system... ...
Begin: Running /scripts/local-top ...
Done.
[ 2.856461] Fusion MPT base driver 3.04.12
[ 2.856532] Copyright (c) 1999-2009 LSI Corporation
[ 2.856995] Fusion MPT SPI Host driver 3.04.12
[ 2.898836] mpispi 0000:00:10.0: PCI INT A -> GSI 17 (level, low) -> IRQ 17
[ 2.943744] rphbase: ioc0: Initiating bringup
[ 2.955880] Floppy drive(s): fd0 is 1.44M
[ 2.955625] pcnet32.c:v1.35 21-Apr-2008 tsbogend@alpha.franken.de
[ 2.956884] pcnet32 0000:02:00.0: PCI INT A -> GSI 18 (level, low) -> IRQ 18
[ 2.956483] pcnet32: PCnet/PCI II 79C97B at 0x2800, 00:0c:29:e7:af:22 assigned IRQ 18.
[ 2.956923] eth0: registered as PCnet/PCI II 79C97B
[ 2.957864] pcnet32: 1 cards found.
[ 2.969684] FDC 0 is a post-1991 E2807
[ 3.028607] ioc0: LS153C103B B0: Capabilities={Initiator}

```

Then, the Nuxeo application automatically starts.



When the Nuxeo application is started, it displays the address at which it is available.



- In your browser, type the indicated address.

The [startup wizard](#) is displayed to help you configure your application.



Shell root access

The password for the root and nuxeo users are generated the first time you start the virtual machine and are displayed on the console.

How to install the Nuxeo Platform from the APT repository for Debian and Ubuntu

Installing the Nuxeo Platform using the APT sources for Debian and Ubuntu installs and configures the platform, but it also installs all required dependencies for an optimal use of the platform.

You need to use the terminal to install the version 5.5 of the package.

1. Download the package file.

```
wget -q http://apt.nuxeo.org/pool/releases/nuxeo_5.5-01_all.deb
```

2. Install the package

```
sudo dpkg --install nuxeo_5.5-01_all.deb
```

3. Follow the instructions displayed.

If it's your first install, you can configure:

- the bind address,
- the port,
- the database (a preconfigured PostgreSQL database is suggested by default).

The platform is installed as a service. It is automatically started and set to automatically start at boot.

4. Open a browser and type the URL <http://localhost:8080/nuxeo/>.

The [startup wizard](#) is displayed so you can setup your Nuxeo platform and select the module you want to install.

Related pages

- [Installing the Nuxeo Platform on Linux](#)
- [Configuring the Nuxeo Platform as a daemon on Debian](#)
- [Installing the Nuxeo Platform on Linux](#)
- [Configuring the Nuxeo Platform as a Daemon with Systemd](#)
- [Installing the Nuxeo Platform on Linux](#)
- [Configuring the Nuxeo Platform as a Daemon on Debian](#)
- [Configuring the Nuxeo Platform as a daemon on Debian](#)
- [Configuring the Nuxeo Platform as a Daemon on Debian](#)
- [Configuring the Nuxeo Platform as a Daemon with Systemd](#)
- [Configuring the Nuxeo Platform as a Daemon with SysVinit](#)
- [Installing the Nuxeo Platform on Linux](#)
- [Configuring the Nuxeo Platform as a Daemon with SysVinit](#)
- [Installing the Nuxeo Platform on Linux](#)
- [Installing the Nuxeo Platform on Linux](#)

Configuring the Nuxeo Platform as a daemon on Debian

The procedure described here is targeted for the Debian Etch distribution, and should be valid for any Debian-based GNU/Linux distribution such as Ubuntu. In other GNU/Linux distributions some commands may be different.

Here is a sample script based on the one used in the debian package

```
#!/bin/sh
### BEGIN INIT INFO
# Provides:          nuxeo
# Required-Start:    $local_fs $remote_fs $network $syslog
# Required-Stop:     $local_fs $remote_fs $network $syslog
# Default-Start:    2 3 4 5
# Default-Stop:      0 1 6
# Short-Description: Start/stop Nuxeo
### END INIT INFO

DESC="Nuxeo"
```

```
NUXEO_USER=nuxeo
NUXEOCTL="/var/lib/nuxeo/server/bin/nuxeoctl"
NUXEO_CONF="/etc/nuxeo/nuxeo.conf"
export NUXEO_CONF

. /lib/init/vars.sh
. /lib/lsb/init-functions

create_pid_dir() {
    mkdir -p /var/run/nuxeo
    chown $NUXEO_USER:$NUXEO_USER /var/run/nuxeo
}

# Change ulimit to minimum needed by Nuxeo
ulimit -n 2048

case "$1" in
    start)
        log_daemon_msg "Starting" "$DESC\n"
        create_pid_dir
        su $NUXEO_USER -m -c "$NUXEOCTL --quiet startbg"
        ES=$?
        log_end_msg $ES
        ;;
    stop)
        log_daemon_msg "Stopping" "$DESC\n"
        su $NUXEO_USER -m -c "$NUXEOCTL --quiet stop"
        ES=$?
        log_end_msg $ES
        ;;
    restart)
        create_pid_dir
        su $NUXEO_USER -m -c "$NUXEOCTL --quiet restart"
        ES=$?
        log_end_msg $ES
        ;;
    force-reload)
        create_pid_dir
        su $NUXEO_USER -m -c "$NUXEOCTL --quiet restart"
        ES=$?
        log_end_msg $ES
        ;;
    status)
        su $NUXEO_USER -m -c "$NUXEOCTL --quiet status"
        exit $?
        ;;
    *)
        echo "Usage: $0 {start|stop|restart|force-reload|status}" >&2
        exit 3
    esac
done
```

```
; ;  
esac
```

Copy the shell script to `/etc/init.d/nuxeo`, replacing paths to match your installation.

Then enable the autostart creating the links in the `rcX.d` directories running the command (as root):

```
$ update-rc.d nuxeo defaults
```

Now restart the machine and verify that nuxeo is started automatically looking at the log file.

If you want to remove the automatic startup use the command (as root):

```
$ update-rc.d -f nuxeo remove
```

You can manage the service with the following command:

```
/etc/init.d/nuxeo [status|start|stop|...]
```

Deploying Nuxeo on your (Amazon AWS) cloud

Need a quick Nuxeo instance for your cloud? You can deploy one in just a few minutes with our CloudFormation template, as we provide a template that will automatically install the latest Nuxeo on your Amazon AWS and all the required resources.

Prerequisites

You need an account on Amazon AWS with the CloudFormation service activated.

To sign up for AWS, just go to <http://aws.amazon.com> and click on the "Sign Up Now" link.

To activate the CloudFormation service, sign in to your management console, click on the "CloudFormation" tab and follow the instructions.

If you don't have a keypair, you will also want to create a one so you can connect to your instance later. You can create one in the "EC2" tab in your management console.

You're ready to deploy our template!

Deploying the template

Deploying the Nuxeo template on Amazon AWS installs:

- the latest version of Nuxeo, with a PostgreSQL database and an Apache2 HTTP front-end;
- all the required Amazon resources, which are: an EC2 instance, an elastic IP, an EBS volume.

To deploy the Nuxeo template:

- Sign in to your CloudFormation management console.
- Choose the region you want your stack to be deployed in.

- Start the new stack creation by clicking the "Create New Stack" button.


4. Choose a stack name and fill in the template URL with: <https://nuxeo.s3.amazonaws.com/templates/Nuxeo.template>.

The screenshot shows the 'Create Stack' dialog box with the 'SELECT TEMPLATE' tab selected. It includes fields for 'Stack Name' (set to 'MyVeryOwnNuxeo'), 'Stack Template Source' (radio buttons for 'Use a sample template', 'Upload a Template File', and 'Provide a Template URL' (selected), with the URL 'https://nuxeo.s3.amazonaws.com/templates/NuxeoDM.template' entered), and an unchecked 'Show Advanced Options' checkbox. A 'Continue' button is at the bottom.

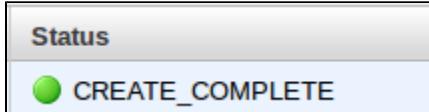
Note: in some regions, AWS will tell you that is not a S3 URL, in that case just download the file locally and use the "Upload a template" option.

5. Fill in your previously created keypair name (KeyName) and the type of amazon instance you want. You can find a list of instance types at <http://aws.amazon.com/ec2/instance-types/>. The default (c1.medium) is suitable for small to medium size installations. If you choose a different instance type, check its "API name" on the instance types page and use that for the "InstanceType" field.

The screenshot shows the 'Create Stack' dialog box with the 'SPECIFY PARAMETERS' tab selected. It displays a 'Template Description' of 'Nuxeo DM installation.' and a 'Specify Parameters' section. Under 'Specify Parameters', there are two entries: 'KeyName' set to 'mykeypairname' (with a note about enabling SSH access) and 'InstanceType' set to 'c1.medium' (with a note about the EC2 instance type). A 'Back' link and a 'Continue' button are at the bottom.

6. Review your settings and click on the "Create Stack" button to start the creation process.

After a few minutes, the instance creation is complete.



Hit the "Refresh" button on the top right corner of the page now and then as it doesn't auto-refresh.

7. Select the line that shows your new CloudFormation stack, then the "Outputs" tab at the bottom. It shows the URL at which you can reach your brand new Nuxeo.



Note that it can take a few more minutes for Nuxeo to be active on that URL as there can still be some installation tasks running at this point.

The template can be used for testing and production purposes. As for every production setup, you will want to check that the configuration suits your needs and tune it as needed: [HTTPS setup](#), disk size, ...

Installing and setting up related software

The Nuxeo Platform modules use external software for some features. They need to be installed on the server in addition to Nuxeo application.

Here is the list of additional resources you may want to install:

- OpenOffice.org and pdftohtml: used for web preview and annotations of office documents in the Document Management module,
- ImageMagick: used for preview and tiling of picture documents in the Document Management and Digital Asset Management module,
- Ffmpeg: for Video features (needed for the Digital Asset Management module),
- Gimp and UFRaw: for RAW format images (needed for the Digital Asset Management module),
- libwpd: used for processing WordPerfect documents.

Under Debian or Ubuntu, all of this can be installed by the following command:

```
sudo apt-get install sun-java6-jdk imagemagick ufraw poppler-utils
openoffice.org ffmpeg libwpd-tools

(apt://sun-java6-jdk,imagemagick,ufraw,poppler-utils,openoffice.org,ffmpeg,libwpd-tools)
```

Windows installer

The following software is already included when using the .exe installer:

- ffmpeg
- ImageMagick
- pdftohtml
- ghostscript

If not already present on the system, you will have the option to automatically install LibreOffice.

On this page

- Setting up OpenOffice/LibreOffice and pdftohtml for preview and annotations on Office documents
- Setting up ImageMagick for picture tiling features
- Setting up ffmpeg
- Setting up Gimp and UFRaw
- Set up libwpd

Setting up OpenOffice/LibreOffice and pdftohtml for preview and annotations on Office documents

Installing OpenOffice/LibreOffice and pdftohtml on the server is only required if you need to use preview (and then possibly annotations) on PDF and office documents.

1. Download and install the following optional components:
 - PDFtoHTML from <http://sourceforge.net/projects/pdftohtml/files/> (necessary for PDF documents preview and annotations)
 - OpenOffice.org 3.x or greater. Available from <http://www.openoffice.org/> (necessary for office documents preview and annotations)
Or LibreOffice 3.x or greater. Available from <http://www.libreoffice.org/> (necessary for office documents preview and annotations)
2. Start the OpenOffice/LibreOffice server (on a single line):

```
soffice.exe -headless -nofirststartwizard  
-accept="socket,host=localhost,port=8100;urp;StarOffice.Service"
```

 If OpenOffice/LibreOffice is already installed on your server, Nuxeo applications come with a daemon that should automatically start it. More information about the daemon configuration below

 If using LibreOffice, you have to add the path to the executable in your nuxeo.conf: jod.office.home=/path/to/libreoffice.

3. Restart the server after launching OpenOffice/LibreOffice server.

More information about the Nuxeo Office daemon

The deprecated OOoDaemonService has been replaced by OOoManagerService. The configuration for the new service can be found in \$NUXEO_HOME/templates/common/config/ooo-manager-config.xml.

Setting up ImageMagick for picture tiling features

The image tiling used in the preview of large images, and so for annotations, needs the installation of the ImageMagick software.

Please see the Nuxeo-Book chapter about "Image tiling". Requirements (ie: ImageMagick 6.3.7 or later) are defined in the installation section.

Linux Debian or Ubuntu:

- `sudo apt-get install imagemagick`

Mac OS X:

- Using Homebrew: `brew install imagemagick`
- Using MacPorts: `sudo port install ImageMagick`

Setting up ffmpeg

To enable video features, you must install ffmpeg on the server.

Windows:

1. Download ffmpeg from <http://ffmpeg.arrozcru.org/autobuilds/>.
2. Extract the ffmpeg archive into a new folder named `C:\ffmpeg\bin` for instance.

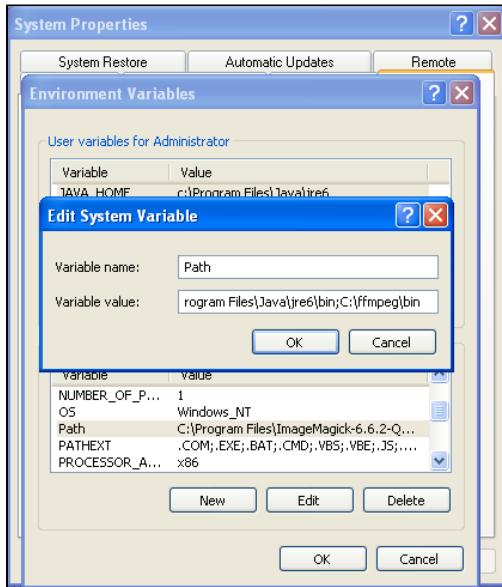


The archives provided by this website should be decompressed with: [7-Zip](#).

3. You have to add the ffmpeg environment variable:
 - Right click on the "My Computer" icon and click on **Properties**.
 - On the "Advanced" tab, edit the `{Path}` system variable and add `;C:\ffmpeg\bin`.



Don't forget the semicolon at the end of existing values.

**Linux Debian or Ubuntu:**

- `sudo apt-get install ffmpeg`

Mac OS X:

- Using [Homebrew](#): `brew install ffmpeg`
- Using [MacPorts](#): `port install ffmpeg +nonfree`

Setting up Gimp and UFRaw

To enable RAW formats in a Nuxeo application, you need to download and install the following optional components:

- Gimp (needed for UFRaw) 2.6.7 or greater from [Gimp Win at SourceForge](#)
- UFRaw from [UFRaw at SourceForge](#)

Linux Debian or Ubuntu:

- `sudo apt-get install ufraw`

Set up libwpd

To enable processing of WordPerfect documents, you need to download and install `libwpd` available at [SourceForge](#)

Linux Debian or Ubuntu:

- `sudo apt-get install libwpd-tools`

Setup

The Nuxeo Platform provides you with easy access to the configuration of your Nuxeo server, thanks to the Admin Center and the Startup Wizard. For advanced configuration or a simple review, manual edition of Nuxeo's configuration file, called `nuxeo.conf`, and a template system is also available.

On this page

- Initial setup of the Nuxeo Platform with the Startup Wizard
 - Server general settings
 - Proxy settings
 - Database settings
 - Mail settings
 - Connect settings
 - Module installation
 - Summary
- Update the application's configuration using the Admin Center
- Manual edition of Nuxeo configuration file nuxeo.conf
 - Related content

Initial setup of the Nuxeo Platform with the Startup Wizard

The first time you start the Nuxeo Platform and go the URL <http://localhost:8080/nuxeo>, a Startup Wizard will guide you to the main configuration steps and enable you to choose which modules you want to enable on the Platform. For each step, a default setting is proposed that enables you to test the application. You can change this default configuration to adapt it to specific environments.

The settings defined during the initial setup can be changed afterward using the Admin Center or by [editing Nuxeo's configuration file](#) manually. Modules can also be added or removed afterwards from the Admin Center.

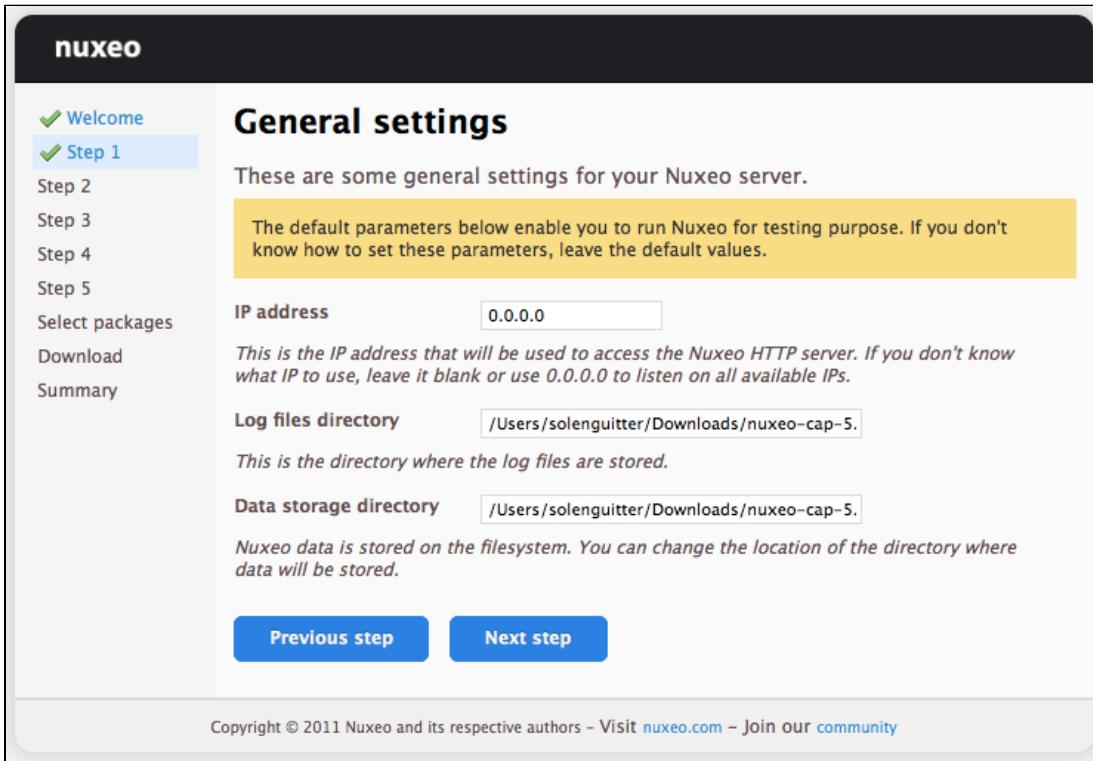
 The Startup wizard will be run only if the configuration sets `nuxeo.wizard.done=false`. You can edit the value in order to replay the wizard (using the Admin Center or editing the `nuxeo.conf` file manually), or simply run `nuxeoctl wizard`

For Internet Explorer 9 users

You need to add the Nuxeo server URL in the trusted sites list to be able to complete the installation and configuration steps. In the Internet Options > Security > Trusted Sites menu, click on the **Sites** button, type the Nuxeo server URL and add it.

Server general settings

This step enables you to change the default IP address of the server and where the logs and data are stored.



nuxeo

General settings

Welcome Step 1 Step 2 Step 3 Step 4 Step 5 Select packages Download Summary

These are some general settings for your Nuxeo server.

The default parameters below enable you to run Nuxeo for testing purpose. If you don't know how to set these parameters, leave the default values.

IP address 0.0.0.0

This is the IP address that will be used to access the Nuxeo HTTP server. If you don't know what IP to use, leave it blank or use 0.0.0.0 to listen on all available IPs.

Log files directory /Users/solenguitter/Downloads/nuxeo-cap-5.

This is the directory where the log files are stored.

Data storage directory /Users/solenguitter/Downloads/nuxeo-cap-5.

Nuxeo data is stored on the filesystem. You can change the location of the directory where data will be stored.

Previous step **Next step**

Copyright © 2011 Nuxeo and its respective authors – Visit nuxeo.com – Join our community

Proxy settings

Some features of Nuxeo applications require access to the Internet. That's the case of the Update Center from which you can access to the Marketplace add-ons and plugins, updates for your application, your Studio customizations.

HTTP proxy settings

Your Nuxeo server must have access to the Internet for you to be able to use the Nuxeo Marketplace and Update Center features. If your organization uses a proxy to access Internet, you can set it here.

If you don't use a proxy, leave the default configuration.

HTTP proxy type: None

- None (selected)
- Anonymous proxy
- Authenticated proxy

Copyright © 2011 Nuxeo and its respective authors - Visit [nuxeo.com](#) - Join our community

Database settings

Nuxeo applications embed a database by default, called H2/Derby. This database enables you to fully test and evaluate the application. However it is not recommended to use this embedded database for production and load testing. Select the database you want to use and provide the connection information to the database.

Possible databases are:

- PostgreSQL,
- Oracle,
- MS SQL Server,
- MySQL.

Database settings

Nuxeo works with several databases. Here you can select and configure the database you want to use.

Nuxeo comes with an embedded database, designed for testing only. For testing and platform evaluation, leave the default database type.

Database type: PostgreSQL

Database name: nuxeo

Database user: nuxeo

Database password: *****

Database server host name: localhost

Database server port: 5432

For more details on database configuration, read the associated documentation.

Previous step Next step

Copyright © 2011 Nuxeo and its respective authors - Visit [nuxeo.com](#) - Join our community

Mail settings

Nuxeo applications include email alert features. By default, no SMTP configuration is enabled and therefore no email alerts will be sent to users. You can refer to the [email alerts section](#) for more information about the SMTP configuration.

The screenshot shows the Nuxeo Platform 5.5 Installation and Administration Guide. On the left, a sidebar lists steps: Welcome, Step 1, Step 2, Step 3, Step 4 (highlighted in blue), Step 5, Select packages, Download, and Summary. The main content area has a title "SMTP Settings". It says "To enable email alerts, Nuxeo needs a SMTP server." Below this is a yellow box containing a note: "Email server configuration is optional. If you don't want or can't provide the information below, leave the default values. The application will work normally. You just won't be able to use the notification system. You may need to ask your system administrator about the parameters below." There are three input fields: "SMTP Server host name" with value "localhost", "SMTP port" with value "25", and "SMTP authentication" set to "No authentication". Below these fields is a note: "Ask your network administrator for your SMTP gateway." and "Default port is usually 25." At the bottom are two buttons: "Previous step" and "Next step".

Connect settings

From this step, you can subscribe to a free 30 days trial offer of Nuxeo Connect which gives you the possibility to evaluate and fully leverage the Marketplace catalog and Nuxeo Studio, the online Nuxeo customization environment. If you subscribe to the trial offer of Nuxeo Connect, you will be sent an email confirming your subscription and your credentials to Nuxeo Connect and giving you the links to access the [Nuxeo Connect Portal](#) and [Nuxeo Studio](#).

nuxeo

- Welcome
- Step 1
- Step 2
- Step 3
- Step 4
- Step 5

Select packages

Download

Summary

Enable Nuxeo Connect & Nuxeo Studio for your installation

Create an account, enable Nuxeo Connect for your installation and get access to:

the latest patches and updates (delivered automatically via the Update Center)
add-ons and plug-ins available through the Nuxeo Marketplace
Nuxeo Studio, our intuitive graphical interface for customizing Nuxeo applications

This will provide you with a free 30-day trial with no commitment.

If you already have a Connect account and wish to link this installation to it, [click here](#)

First name	John	Last name	Doe
Username *	jdoe	eMail address *	jdoe@company.com
Password *	*****	Password (verify) *	*****
Company *	Company		

Be sure to double-check your email address.
Your connection information and links to the Nuxeo Studio documentation will be sent to you by email.

[Previous](#) [Validate & register](#) [Or skip and don't register](#)

Copyright © 2011 Nuxeo and its respective authors – Visit [nuxeo.com](#) – [Join our community](#)

 If you already have a Nuxeo Connect account, you can register your Nuxeo instance from this step to directly be able to apply your Nuxeo Studio customizations and the installation of Nuxeo Marketplace packages in your instance.

Module installation

Select the modules you want to install on the Platform. You can also just keep the naked Content Application Platform.

 You can install or uninstall modules afterwards from the [Admin Center](#).

nuxeo

- Welcome
- Step 1
- Step 2
- Step 3
- Step 4
- Step 5
- Nuxeo Connect
- Select packages

Download

Summary

Select modules

Select here the features you would like to install with Nuxeo Platform.

- Content Application Platform (CAP)
 Advanced Document Management (DM)
 Social Collaboration (SC)
 Digital Asset Management (DAM)
 Case Management Framework (CMF)

Once installed you will be able to pull more features from the Nuxeo Marketplace.

[NUXEOSTC](#) [NUXEODM](#) [NUXEODAM](#) [NUXEOCAP](#)

[Previous step](#) [Next step](#)

Copyright © 2011 Nuxeo and its respective authors – Visit [nuxeo.com](#) – [Join our community](#)

And if needed, download the module packages. Packages may be already included in the Platform.

Modules download

Download the packages needed for your installation.

Selected modules for your installation.

- nuxeo-dm-5.5-RC1.zip (already in local)
- nuxeo-content-browser-5.5-RC1.zip (already in local)

Previous step **Next step**

Summary

Copyright © 2011 Nuxeo and its respective authors – Visit [nuxeo.com](#) – [Join our community](#)

Summary

A final Summary step provides you with a screen on which you can see all the configuration parameters that you set in the previous steps so you can review them and possibly go back to a step to change them.

Summary

Congratulations, you have just configured your application. Here is the summary of the configuration parameters you have just set.

Non-default values are in bold.

SMTP authentication	false
SMTP Server host name	localhost
SMTP password	password
SMTP port	25
SMTP login	anonymous
IP address	0.0.0.0
Data storage directory	/Users/solenguitter/Downloads/nuxeo-cap-5.5-RC1-tomcat/nxserver/data
Database server host name	localhost
Database name	nuxeo
Database password	password
Database server port	5432
Database user	nuxeo
Database type	postgresql
Log files directory	/Users/solenguitter/Downloads/nuxeo-cap-5.5-RC1-tomcat/log

Previous step **Start Nuxeo**

Copyright © 2011 Nuxeo and its respective authors – Visit [nuxeo.com](#) – [Join our community](#)

To validate your configuration, click on the **Start Nuxeo** button. The server will automatically restart and your configuration will be applied. Once the server is restarted, you are displayed the login page. Log in to your application the **Administrator** user name and the **Administrator** passwor

d.

The screenshot shows the Nuxeo setup wizard interface. On the left, a vertical list of steps is shown with green checkmarks: Welcome, Step 1, Step 2, Step 3, Step 4, Step 5, Nuxeo Connect, Select packages, Download, and Summary. The 'Welcome' step is currently active. To its right, the text 'RESTARTING, PLEASE WAIT...' is displayed. Below this, there is a message: 'The page will automatically be refreshed when Nuxeo is ready.' and 'If you didn't change the authentication, the credentials will be Administrator/Administrator (login/password).'. At the bottom of the screen, a copyright notice reads: 'Copyright © 2011 Nuxeo and its respective authors – Visit nuxeo.com – Join our community'.

Update the application's configuration using the Admin Center

The **Admin Center** is the graphical interface that enables the application's administrators to edit the configuration of the application directly from the user interface, and prevents them from editing .xml and .conf files. They can edit the configuration of the application, monitor it, display messages to the users, and easily customize the application thanks to the Update Center.

To edit the configuration of the application using the Admin Center:

1. Log in with an administrator account.
Default administrator credentials are:
 - login: Administrator
 - password: Administrator
2. Click on the **Nuxeo Admin Center** link in the page header.
3. Click on the **Setup** tab, edit the configuration you want to change and Save.

The screenshot shows the Nuxeo Admin Center interface. The top navigation bar includes links for Home, Document Management, Admin Center (which is selected), Studio, and Administrator. A search bar is also present. The left sidebar contains links for System information, Activity, Nuxeo Connect, Update Center, Monitoring, OAuth / OpenSocial, Users & groups, Vocabularies, Dashboards, and Themes. The main content area is titled 'Setup' and contains several sections: 'Main information' (Application URL: http://localhost:8080/nuxeo, Product name: Nuxeo DM, Product version: 5.5-RC1, Bind address: 0.0.0.0, Configuration file: /Users/solenguitter/Downloads/nuxeo-cap-5.5-RC1-tomcat/bin/nuxeo.conf, Data directory: /Users/solenguitter/Downloads/nuxeo-cap-5.5-RC1-tomcat/nx, Log directory: /Users/solenguitter/Downloads/nuxeo-cap-5.5-RC1-tomcat/log, Database: Default (H2/Derby embedded databases)), 'Proxy Http' (Proxy type: No proxy), and 'Email information' (Email notifications subject prefix: [Nuxeo], User for mail authentication: empty). The tabs at the top of the main content area are Host, Nuxeo distribution, Setup (which is selected), Repository statistics, and Repository binaries.

4. If indicated as needed on top of the page, restart the server.



You can also take a look at the following pages for recommendations and examples:

- Recommended configurations,
- Configuration examples.

You can report to the Configuration parameters index ([nuxeo.conf](#)) for more information about the available parameters.

Manual edition of Nuxeo configuration file nuxeo.conf

By default, the `nuxeo.conf` file is located in `$NUXEO_HOME/bin`. If you installed your application using the Windows installer, the configuration is located in `%APPDATA%\Nuxeo DM\conf` (check the [corresponding Knowledge Base page](#) for more information). If you plan to use the application in production, you should [move the configuration file outside the Nuxeo home directory](#), to make upgrades easier and more secured: your data and configuration won't risk to be overridden or lost.

For Windows users

Do not use Office writers, nor Notepad.

Wordpad is fine, Notepad++ and SciTE are good text editors, there are a lot of [other text editors](#).

You can report to the Configuration parameters index (`nuxeo.conf`) for the list of available parameters.

Related content

-  [Setup](#)
-  [Recommended configurations](#)
-  [Configuration templates](#)
-  [Configuration examples](#)
-  [Configuration parameters index \(`nuxeo.conf`\)](#)
-  [Nuxeo clustering configuration](#)

Recommended configurations

Nuxeo applications come as ready-to-use applications, that you can quickly install and evaluate. However, if you plan to go in production, have several Nuxeo applications on the same machine or do some performance tests, here are some changes of configuration that we recommend to do, especially for advanced testing or before going into production:

The steps given below are given using the Admin Center. They can of course also be done by editing the `nuxeo.conf` file manually.



More configuration use cases on the [Configuration examples page](#).

On this page

- Move configuration, data and log directories outside Nuxeo
- Define environment variables
 - `NUXEO_HOME`
 - `NUXEO_CONF`
 - Windows specific case
- Enable email alerts
- Related content

Move configuration, data and log directories outside Nuxeo

The configuration of your application is saved in the `nuxeo.conf` configuration file, whatever the means you use to configure your application (manual edit, Startup Wizard or Admin Center). It is better, although not mandatory, to store your customized configuration outside Nuxeo. This way, you will be able to easily upgrade Nuxeo, keeping your configuration safely apart of Nuxeo directory.

To move the configuration file outside the Nuxeo directory:

1. Move the `nuxeo.conf` file from its default location.
2. After you moved `nuxeo.conf`, you need to [define its location as an environment variable](#).

By default, data and log directories are stored inside the Nuxeo tree. To ease backup and upgrades, it is highly recommended to move them outside the Nuxeo tree.

To move the data and log directories:

1. In the Admin Center, type the path to the location where you want the directories to be stored (see the table below).
2. Click on **Save**.
3. Restart your server.
The data and log directories are created at the location you typed.

Data and log directories configuration

Field / Property	Description
Data directory nuxeo.data.dir	Data directory (absolute or relative to NUXEO_HOME). It involves all data not being stored in the database. Linux recommended path: /var/lib/nuxeo/...
Log directory nuxeo.log.dir	Log directory (absolute or relative to NUXEO_HOME). Linux recommended path: /var/log/nuxeo/...

Define environment variables

When the server starts, it guesses where the Nuxeo home directory and the Nuxeo configuration file (`nuxeo.conf`) are located. If it doesn't find it or if you want to force it to use a specific home directory and/or a specific configuration file, you can define their location as environment variables.

NUXEO_HOME

Here is how Nuxeo home is guessed when the server starts:

- Before Nuxeo EP 5.4.1, if `NUXEO_HOME` is not set, then the following environment variables are used to find the Nuxeo Home, by order: `JBOSS_HOME`, `CATALINA_HOME`, `JETTY_HOME`.
- Since Nuxeo EP 5.4.1, if `NUXEO_HOME` is not set, then the parent directory of the called script (`nuxeoctl`) is used.

Setting the Nuxeo home directory as an environment variable is recommended in the following cases:

- if you installed several Nuxeo applications on the same machine (for evaluation or production purpose),
- if you want to use other scripts than the `$NUXEO_HOME/bin/nuxeoctl` script (such as a service in `/etc/init.d`).

You must then set `NUXEO_HOME=/path/to/nuxeo/` in the system environment variables:

Windows users must write "set `NUXEO_HOME=...`" or use the control panel interface to define user environment parameters (like it's done for `%PATH%`).

Linux and Mac OS X users will write "export `NUXEO_HOME=....`" in `~/.bashrc` or `~/.profile`.

NUXEO_CONF

You need to set the location of the `nuxeo.conf` file as an environment variable if you moved your configuration outside of the Nuxeo directory.

Moving the data and configuration outside the Nuxeo directory is recommended in a production environment because it makes upgrades easier and more secured: your data and configuration won't risk to be overridden or lost.

You must then set `NUXEO_CONF=/path/to/nuxeo.conf` in the system environment variables.

Windows specific case

Under Windows, the location of the `nuxeo.conf` is defined by that order of priority (ie first one of those found is used):

- Registry key `HKEY_LOCAL_MACHINE\SOFTWARE\%PRODNAME%\ConfigFile` with `%PRODNAME%` equals to "Nuxeo" (or in older versions, "Nuxeo CAP", "Nuxeo DM", "Nuxeo DAM", ...),
- Environment variable `NUXEO_CONF`,
- "`nuxeo.conf`" file in the working directory,
- "`nuxeo.conf`" file on the Desktop,
- "`nuxeo.conf`" file in the same location as `nuxeoctl.bat`.

Enable email alerts

Default Nuxeo DM email configuration is filled in with neutral values that you need to edit to make Nuxeo DM work with your mail server. Unless you do that, alerts emails won't be sent to users. Unless you do that, alerts emails won't be sent to users.

To make alerts available:

1. In the Admin Center, click on the **Setup** tab of system information section.
2. Edit and fill in the values of the Email information section (see below for expected parameters).



To enable alerts, filling in the SMTP parameters should be sufficient for most mail server configurations.

3. Click the button **Save**.
As indicated on top of the page, you need to restart your server so the new configuration is taken into account.

Email information configuration

Field / Property	Description

Email notifications subject prefix nuxeo.notification.eMailSubjectPrefix	Text displayed in the "Object" before the object of the alerts email to help users identify that the emails are coming from the application. Default value is "[Nuxeo]". You can change it to whatever value you like.
Mail store protocol mail.store.protocol	Name of the protocol used to store emails on the server. Default value is "pop3". You may need to change it to "IMAP".
Mail transport protocol mail.transport.protocol	Name of the protocol used to send emails. Default value is "smtp". This should work in most cases.
Host name for POP3 mail.pop3.host	Name of the mail server host used to receive and store emails. Default value is "pop3.nosuchhost.nosuchdomain.com". You need to change it.
Debug mode mail.debug	Default value is set to "false". Change it to "true" if you want to have the details of what the server is doing in the logs.
Host name for SMTP mail.smtp.host	Mail server host name for outgoing mails. Default value is "localhost". You need to change it so emails can be sent from the server.
Port number for SMTP mail.smtp.port	Mail server port for outgoing emails. Default value is 25.
Use authentication for SMTP mail.smtp.auth	Indicate if authentication is needed for the mail server to send emails. Default value is "true". You should change it to "false" if no authentication for sending email is required.
Use STARTTLS for SMTP mail.smtp.usetls	Indicate if STARTTLS is needed for the mail server. Default value is "false". You should change it to "true" if your SMTP requires STARTTLS
SMTP username mail.smtp.username	Type the username that will be used if you set the authentication for SMTP parameter to "true".
SMTP password mail.smtp.password	Type the password that will be used if you set the authentication for SMTP parameter to "true".
Sender address mail mail.from	Email address that will be displayed as the sender's address.



If you have complex mail server configurations, you may want to check the [Javamail API FAQ](#) for more information.

Related content

- [Recommended configurations](#)
- [Configuration templates](#)
- [Configuration examples](#)
- [Configuration parameters index \(nuxeo.conf\)](#)
- [Setup](#)
- [Nuxeo clustering configuration](#)

Configuration examples

Here are some configuration usecases:

- [Changing the Live Edit default version incrementation](#)
- [Changing the default port \(8080\)](#)
 - Related pages in this documentation
 - In other documentations



The use of the Admin Center is highlighted in the steps below. However, you can do the same configurations by editing the `nuxeo.conf` file manually.

Changing the Live Edit default version incrementation

When users edit documents with Live Edit, the default behavior is that no version incrementation occurs. This default behavior can be changed and you can set what version number should be incremented when users save a document with Live Edit.

Configure default Live Edit version incrementation:

1. In the Admin Center, click on the **Setup** tab of system information section.
2. In the Advanced Settings, edit the value of the parameter "org.nuxeo.ecm.platform.liveedit.autoversioning" :
 - **minor** will instruct the server to automatically increment the minor version of the document,
 - **major** will instruct the server to automatically increment the major version of the document,
 - **none** will instruct the server to not increment the version of the document (this is the default value).
3. Click the button **Save**.

Changing the default port (8080)

Nuxeo applications run on the 8080 port by default. As it may be used by another application, you may need to change it.

Change the default port:

1. In the Admin Center, click on the **Setup** tab of system information section.
2. In the Advanced Settings, edit the value of the parameter "nuxeo.server.http.port".
3. Click the button **Save**.
4. Restart the server as indicated on top of the page.

Related pages in this documentation

-  [Recommended configurations](#)
-  [Configuration templates](#)
-  [Configuration examples](#)
-  [Configuration parameters index \(nuxeo.conf\)](#)
-  [Setup](#)
-  [Nuxeo clustering configuration](#)

In other documentations

-  [Working with Live Edit \(Nuxeo Document Management \(DM\) - 5.5\)](#)
-  [Installing Live Edit \(Nuxeo Document Management \(DM\) - 5.5\)](#)
-  [LiveEdit makes MS Office slow to start \(Nuxeo Technical Knowledge Base \(FAQ\)\)](#)
-  [Setup Firefox protocol handler with LiveEdit 2 for MS Office and OpenOffice.org \(Nuxeo Technical Knowledge Base \(FAQ\)\)](#)
-  [LiveEdit icons are still available in Nuxeo after LiveEdit has been uninstalled \(Nuxeo Technical Knowledge Base \(FAQ\)\)](#)
-  [I can't view my websites and blogs \(displays a message "The HTTP header field "Accept" with value..."\) \(Nuxeo Technical Knowledge Base \(FAQ\)\)](#)
-  [Live Edit compatibility table \(Nuxeo Document Management \(DM\) - 5.5\)](#)
-  [Manage your own file type with LiveEdit \(Nuxeo Document Management \(DM\) - 5.5\)](#)

Configuration templates

Nuxeo applications integrate a configuration templates system to ease configuration and maintenance of configuration files.

Nuxeo comes with default templates which mainly provide database configurations, but the templates can be used for any configuration purpose.

Properly using that template system ensures your customization of Nuxeo exclusively resides in your nuxeo.conf, custom templates and plugin modules.

For instance, users can create templates for development, pre-production, and production environments; each template will include a different set of xml contributions (users, ldap integration, database used, ...).

Templates are located in the "templates" directory (`$NUXEO_HOME/templates`). To enable a configuration, such as database configuration, you just need to indicate which template to use in the Admin Center's Setup tab or in the nuxeo.conf configuration file.

Here are the templates provided by default:

- **common**: common template used by other templates;
- **common-binding**: (JBoss only), template used by other templates;
- **common-deploydir**: (JBoss only), template used by other templates;
- **default**: default Nuxeo configuration template for test purpose;
- **https**: (Tomcat only), not recommended template for making the server listen to port 443 (HTTPS);

- monitor: (JBoss only), activate the JBoss LogginMonitor service to log miscellaneous MBean informations;
- [postgresql](#): PostgreSQL configuration template;
- [postgresql-quartz-cluster](#)
- [mssql](#): MS SQL Server configuration template;
- [mssql-quartz-cluster](#)
- [mysql](#): MySQL configuration template;
- [oracle](#): Oracle configuration template;
- [oracle-quartz-cluster](#)
- [custom](#): sample custom templates. Of course, this template is empty by default. One should copy it outside \$NUXEO_HOME and adapt to his needs.



For production environment, it is recommended to define your own custom template outside \$NUXEO_HOME, as for `nuxeo.conf`. It must then be referenced in `nuxeo.conf` with its absolute path.

Technical overview

A server is considered as already configured when it has a "config" directory.

When the "config" directory doesn't exist, templates will be used to generate all configuration files (config and datasources).

The template files contain defined parameters such as \${sample.parameter}.

Values for parameters replacement are calculated this way:

- If `nuxeo.conf` does not define `nuxeo.templates`, then `nuxeo.templates` equals "default" (the deprecated parameter `nuxeo.template` is still read for backward compatibility).
- The `$(nuxeo.templates)` value is used for determining the chosen template(s).
- For each value "nuxeo.template" of `$(nuxeo.templates)` (comma separated values, relative to "templates/" directory or absolute path), the corresponding file `templates/$(nuxeo.template)/nuxeo.defaults` is read for defining new default values and maybe including other templates which are recursively parsed.
- The file `templates/nuxeo.defaults` is read for default values not already defined.
- The file `nuxeo.conf` is read for custom values (overwriting default values).

Configuration files are then generated by this way:

- For each comma separated value of `nuxeo.templates` and `nuxeo.template.includes` (let say `sample.template`), files in `templates/$(sample.template)` are copied using the previously calculated values for replacing parameters.
- Every included template will potentially overwrite its precedents.

Related content:

- [Adding custom templates](#)
- [Connecting Nuxeo to the database](#)
- [Configuration templates](#)
- [Configuration parameters index \(nuxeo.conf\)](#)

Case Management recommended configurations

Mailbox synchronization

Synchronization is triggered when the event 'syncMailbox' occurs. Default configuration is available in 'org.nuxeo.cm.schedule' component. The event is programmed for 1 am every day. If you don't need synchronization, you have to remove or modify `cm-sync-scheduler-config.xml` in `templates/cmf/config`:

```
<?xml version="1.0"?>
<component name="org.nuxeo.cm.schedule.custom.contrib">

<require>org.nuxeo.cm.schedule</require>

<extension
  target="org.nuxeo.ecm.platform.scheduler.core.service.SchedulerRegistryService"
  point="schedule">

<documentation>
  Change the frequency of syncMailbox event so it is triggered
  at 01:00am on the last day of every month.
  Details about the cronExpression syntax can be found here:
  [http://www.quartz-scheduler.org/docs/tutorials/crontrigger.html]
</documentation>

<schedule id="syncMailbox">
  <username>Administrator</username>
  <password>Administrator</password>
  <eventId>syncMailbox</eventId>
  <eventCategory>cmSync</eventCategory>
  <cronExpression>0 0 1 L * ?</cronExpression>
</schedule>

</extension>
</component>
```

Synchronization Service

Default configuration will synchronize user and group directories.

If you need to deactivate one of them, or change the mailbox titles are generated, you need to override default configuration. Simply add an xml file like cm-sync-mailbox-custom-config.xml in your template folder with the following content:

```

<?xml version="1.0"?>
<component name="org.nuxeo.cm.service.synchronization.contrib.custom">

    <require>org.nuxeo.cm.service.synchronization.contrib</require>
    <extension
target="org.nuxeo.cm.core.service.synchronization.MailboxSynchronizationService"
point="directoryToMailbox">

        <!-- deactivate user synchronization -->
        <userToMailbox enabled="false" />

        <!-- use a custom title generator for group synchronization -->
        <groupToMailbox>

            <titleGenerator>com.company.project.cmf.synchronization.CustomGroupMailboxTitleGenerator</titleGenerator>
            </groupToMailbox>

        </extension>
    </component>

```

If you need more details on synchronization implementation, please visit the [Mailbox synchronization page](#).

Configuration parameters index (nuxeo.conf)

Here is a list of available parameters for nuxeo.conf. This list may not be exhaustive but it will be often updated.

Those parameters can be either environment parameters used by Nuxeo runtime or template parameters used for values replacement in configuration files.

Parameter	Default value (" " separates possible values)	Description
JAVA_HOME	None. If undefined nuxeoctl script will try to discover it.	Path to Java home directory.
JAVA_OPTS	-Xms512m -Xmx1024m -XX:MaxPermSize=512m -Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000 -Dfile.encoding=UTF-8	Optional values passed to the JVM. Nuxeo requires at least 1024 Mo in JVM heap size and 256Mo as maximum permanent size (512 recommended). Decreasing garbage collector frequency avoid having too much CPU usage (Sun Java specific options, recommended by JBoss).
launcher.start.max.wait	300	Since 5.4.1. Maximum time to wait for effective Nuxeo server start before giving up (applies on commands "start" and "restart").
launcher.stop.max.wait	60	Since 5.5. Maximum time to wait for effective Nuxeo server stop cleanly before using forced stop.

launcher.override.java.tmpdir	true	Since 5.4.1. Possible values: true or false If true, will set <code>java.io.tmpdir = nuxeo.tmp.dir.</code>
nuxeo.log.dir	log	Log directory (absolute or relative to <code>NUXEO_HOME</code>). Linux recommended path: <code>/var/log/nuxeo/...</code>
nuxeo.pid.dir	bin	Directory where to store Nuxeo PID file.
nuxeo.data.dir	data	Data directory (absolute or relative to <code>NUXEO_HOME</code>). It involves all data not being stored in database. Linux recommended path: <code>/var/lib/nuxeo/...</code>
nuxeo.tmp.dir	server/default/tmp (JBoss) temp (Tomcat) tmp (Jetty)	Location of the temporary files.
nuxeo.force.generation	true once	If "true", will force generation of configuration files; otherwise they are only generated when not existing. If "once", will force one time and switch to false after successful generation. If "false", configuration changes are ignored.
nuxeo.templates	default	Comma separated list of templates to include. Templates paths are absolute or relative to <code>\$NUXEO_HOME/templates/</code> . Available templates: postgresql, mysql, mssql, oracle, custom, ...
nuxeo.bind.address	0.0.0.0	Server binding address. "0.0.0.0" means "all available network interfaces". WARNING: when changing "nuxeo.bind.address", you must accordingly change "nuxeo.loopback.url".
nuxeo.server.http.port	8080	Server HTTP listen port.
nuxeo.server.ajp.port	8009	Server AJP listen port. This is not available on Jetty.
nuxeo.server.jvmRoute	nuxeo	Server AJP route for load-balancing (since 5.4.2)
nuxeo.server.tomcat-admin.port	8005	Tomcat server's "admin" port. This is only useful if you have another tomcat server running and want to avoid port conflicts.
nuxeo.server.https.port	8443	Server HTTPS listen port. This is only useful if you have modified the application server to use HTTPS.

nuxeo.server.emptySessionPath	false	(Tomcat only) Since 5.5. If set to true, all paths for session cookies will be set to /. May be useful to enable authentication on proxyfied WebEngine applications (see HTTP and HTTPS reverse-proxy configuration).
org.nuxeo.ecm.instance.name	Nuxeo 5.4	Server name.
org.nuxeo.ecm.instance.description	Nuxeo ECM server	Server description.
org.nuxeo.ecm.product.name	Nuxeo Platform	Product name, displayed in the page title on your browser.
org.nuxeo.ecm.product.version	5.5	
org.nuxeo.ecm.webapp.dashboard.mode	auto	Defines the dashboard mode. There are 3 modes: <ul style="list-style-type: none"> • auto: (default value) let Nuxeo choose dashboard based on user browser capabilities. • old: force usage of the 'old' JSF based dashboard for all users. • opensocial: force usage of the new OpenSocial based dashboard for all users.
templateName.target	server/default/deploy/nuxeo.ear	Directory where <i>templateName</i> files will be deployed.
mailservice.user	nobody	(JBoss only) User for mail authentication.
mailservice.password	password	(JBoss only) Password for mail authentication.
mail.store.protocol mail.transport.protocol	pop3 smtp	Server protocol parameters for mailing.
mail.user	nobody	User who will receive mail (unused in Nuxeo).
mail.pop3.host	pop3.nosuchhost.nosuchdomain.com	Mail server.
mail.debug	false	Enable debugging output from the javamail classes.
nuxeo.notification.eMailSubjectPrefix	[Nuxeo]	Subject prefix in Nuxeo notification mails.
mail.smtp.host	localhost	SMTP gateway server.
mail.smtp.port	25	Mail server port.
mail.smtp.useTLS	false	Use TLS for the SMTP connection.
mail.smtp.auth	true	
mail.smtp.username	anonymous	
mail.smtp.password	password	
mail.from	noreply@nuxeo.com	The address mails will be sent from.
nuxeo.db.name	nuxeo NUXEO	Database name.

nuxeo.db.user	sa nuxeo	Database username.
nuxeo.db.password	(empty value) password	Database password.
nuxeo.db.host	localhost	Database host URL.
nuxeo.db.port	3306 1521 5432	Database host port.
nuxeo.db.jdbc.url	(database-dependent)	Database JDBC connection URL for Nuxeo datasources, for instance <code>jdbc:postgresql:/\${nuxeo.db.host}:\${nuxeo.db.port}/\${nuxeo.db.name}</code> .
nuxeo.db.validationQuery		Database validation query, a <code>SELECT</code> statement used to check connections before using them, usually <code>SELECT 1</code> . Using this has a noticeable speed impact but makes connections resilient to network or sever problems.
nuxeo.db.min-pool-size	5	Database minimum pool size for Nuxeo datasources.
nuxeo.db.max-pool-size	20 (JBoss) 100 (Tomcat)	Database maximum pool size for Nuxeo datasources.
nuxeo.vcs.min-pool-size	0	Database minimum pool size for Nuxeo repository (VCS).
nuxeo.vcs.max-pool-size	20	Database maximum pool size for Nuxeo repository (VCS).
nuxeo.url	http://localhost:8080/nuxeo	Application URL (without final slash)
nuxeo.loopback.url	http://localhost:8080/nuxeo	Since 5.4.1. Nuxeo URL, for connections from Nuxeo to itself (theme banks default). The port should be the same as <code>nuxeo.server.http.port</code> . Since 5.5, if not explicitly configured, the loop back URL is generated from <code>nuxeo.bind.address</code> , <code>nuxeo.server.http.port</code> and <code>org.nuxeo.ecm.contextPath</code> values.
org.nuxeo.ecm.contextPath	/nuxeo	Application context path. You also have to accordingly rename all occurrences of <code>nuxeo.xml</code> (for Tomcat)
org.nuxeo.ecm.platform.transformer.ooo.host.name	127.0.0.1	DEPRECATED.
org.nuxeo.ecm.platform.transformer.ooo.host.port	8100	DEPRECATED.
org.nuxeo.ecm.platform.transformer.ooo.version	2.2.1	DEPRECATED.
org.nuxeo.ecm.platform.transformer.ooo.enableDaemon	true	DEPRECATED.
jod.connection.protocol	SOCKET	OpenOffice Connection protocol, either PIPE or SOCKET.
jod.max.tasks.per.process	200	Maximum task per Office instance before restarting it

jod.task.execution.timeout	120000	Will stop the task if it's not completed after timeout.
jod.task.queue.timeout	30000	Will stop looking for the next task in the queue after timeout.
jod.office.home		Home directory of OpenOffice or LibreOffice.
jod.office.ports	2003	Since 5.5. When running in SOCKET mode, comma-separated list of ports used for the office connection.
jod.office.pipes		Since 5.5. When running in PIPE mode, comma-separated list of pipe names used for the office connection.
jod.jpipe.lib.path		Path to Jpipe library. Only used when connecting to OO through PIPE.
jod.template.profile.dir		Path to custom OO template directory.
opensocial.gadgets.host	localhost	
opensocial.gadgets.port	8080	
opensocial.proxy.proxySet	false	
opensocial.proxy.proxyHost		
opensocial.proxy.proxyPort		
opensocial.proxy.user		
opensocial.proxy.password		
repository.clustering.enabled	false	Activate clustering mode.
repository.clustering.delay	1000	When clustering is activated, defines the delay during which invalidations don't need to be processed (expressed in milliseconds).
repository.binary.store		Defines the folder where binaries are stored. Useful when using clustering or to change the location of binaries to another location.
nuxeo.templates.parsing.extensions	xml,properties	Files extensions being parsed for parameters replacement when copying templates.
org.nuxeo.ecm.jboss.configuration	default	JBoss configuration to use ("default", "minimal", "all", ...) Pay attention to the fact that this won't apply to templates defining their own " <i>template.target</i> " value (for instance, "default" template sets "default.target=server/default/deploy" without being aware of "org.nuxeo.ecm.jboss.configuration" value).
zip.entry.encoding	ascii	Choose how to encode filename when exporting documents to zip in the worklist.

org.nuxeo.ecm.platform.liveedit.autoversioning	none,minor,major	see Configuration examples
nuxeo.wizard.done	true or false depending on the package	If set to false, will start a setup wizard before starting Nuxeo.
nuxeo.http.proxy.host		HTTP proxy host.
nuxeo.http.proxy.port		HTTP proxy port.
nuxeo.http.proxy.login		HTTP proxy login.
nuxeo.http.proxy.password		HTTP proxy password.
facelets.REFRESH_PERIOD	-1	Indicates to the compiler the number of seconds to wait between subsequent checks for changes in modified JSF facelets in a running application. Useful for facelet debugging. To disable this compiler check use a value of -1 which is a recommended value for production deployments as compiler checks have an impact on application performance.
nuxeo.db.transactiontimeout	300	Database transaction timeout in seconds (available for server tomcat only - Since 5.5)
server.status.key		Since 5.5. Secure key for connecting to server status monitoring servlet. It is randomly generated if not set.
session.timeout	60	Since 5.5. Session timeout (see web.xml session-timeout).
nuxeo.updatecenter.disabled	false (unset)	Since 5.5, Disable the Update Center feature.
theme.useOldLocalConfiguration	false (unset)	Since 5.5, Use the old local configuration for theme, selecting a theme page instead of a flavor.
org.nuxeo.big.file.size.limit	5Mi (unset)	Since 5.4.1, redirect onto the big file download url if size exceed limit
org.nuxeo.ecm.platform.ui.web.auth.NuxeoAuthenticationFilter.isLoginNotSynchronized	false (unset)	Since 5.5, disable login synchronization
nuxeo.wizard.packages.url		Since 5.5, defines the base url used by the Setup Wizard to get the packages.xml file describing the available software packages options
nuxeo.wizard.skippedpages	null	Since 5.5, comma separated list of pages that should be skipped inside the wizard
nuxeo.pageprovider.default-max-page-size	100	Since 5.5-HF06, defines the default maxPageSize to use in pageProvider when no value is defined in the pageProvider contribution. Value '0' means no limit.

Related content

-  Recommended configurations
-  Configuration templates
-  Configuration examples
-  Nuxeo clustering configuration

Making CM and DM work with the same repository

Deployment

Case Management is a framework based on the Nuxeo Platform. That means that all the document management features available on the Platform are available or can be easily integrated to CMF.

Despite that, an interesting use case might be to have two separate instances, one for document management and one case management running on the same repository.

For this configuration, the clustering needs to be enabled on the two servers.

The documentation to enable the clustering can be found here: [Nuxeo clustering configuration](#).

We also need to have the same document types available on both instances and a common structure of the root documents.

For this you need to deploy the following case management JARs on DM:

- nuxeo-case-management-api
- nuxeo-case-management-core
- nuxeo-case-management-lang

By default, we have the following structure for the Default Domain:

Case Management
-CaseRoots
-MailboxRoots
-SectionRoots

If you need to modify it , in order to make available for example the Workspaces root (as in standard DM), you should override the contentTemplate defined in cm-content-template-contrib.xml and deploy your contribution on both instances.

Database

Nuxeo applications store most of their data in a SQL database. Several databases are supported, but they must be configured to work correctly.

This takes two steps:

1. Configure the database:
 - PostgreSQL (8.4 and 9.0),
 - Oracle (10g R2 (10.2.0.5) and 11g),
 - MS SQL Server (2005 or 2008).
2. Connect Nuxeo to the database.

Configuring PostgreSQL

Nuxeo supports

PostgreSQL 8.4, 9.0 and 9.1

. We always recommend that you use the latest stable version, which is PostgreSQL 9.1 at the time of this writing.

(PostgreSQL 8.3 and earlier are not supported anymore and shouldn't be used in production anyway as these versions have lots of performance issues and missing features.)

The database needs to be configured to work properly with Nuxeo. Some settings **must** be changed in order for Nuxeo to work. Other settings *should* be changed in order for Nuxeo to have good performance.

This FAQ will give you some hints to configure your database, but please refer to your DBA or the PostgreSQL documentation for more information (http://wiki.postgresql.org/wiki/Tuning_Your_PostgreSQL_Server).

Most settings have to be changed in the `postgresql.conf` file. Some SQL commands may have to be executed directly at the PostgreSQL console (`psql`).

On this page

- Mandatory changes
 - Two-phase commit
 - Implicit casts
 - Language plpgsql
 - Create the role and database for Nuxeo
- Specific configuration
 - Accent-insensitive fulltext search
- Performance tuning
 - Shared buffers and system cache
 - Memory for workers
 - Buffering writes
 - Mass import specific tuning
 - Index vs table scan
 - Updating the planner statistics
 - Vacuuming
 - Monitoring
 - Reporting problems
 - Related content in this documentation
 - In other documentation

Mandatory changes

Two-phase commit

Nuxeo uses two-phase commit and needs to have the `max_prepared_transactions` settings updated.

```
max_prepared_transactions = 64
```

! This change is mandatory for PostgreSQL >= 8.4 since prepared transactions are disabled by default. If you don't change this option you will have `javax.transaction.HeuristicMixedException` exceptions.

Implicit casts

Jena (used for relations and comments) and jBPM (used for workflows) assume some implicit value casting in the SQL they generate. However since PostgreSQL 8.3 the database is much stricter than PostgreSQL 8.2 with respect to value casting.

To make Nuxeo work with PostgreSQL >= 8.3, you must therefore execute the following commands in your PostgreSQL console when connected to the `template1` database, so that any database created afterward will come with the required CASTs (if your database is already created, execute the commands in your database as well):

```
CREATE FUNCTION pg_catalog.text(integer) RETURNS text STRICT IMMUTABLE LANGUAGE SQL AS
'SELECT textin(int4out($1));';
CREATE CAST (integer AS text) WITH FUNCTION pg_catalog.text(integer) AS IMPLICIT;
COMMENT ON FUNCTION pg_catalog.text(integer) IS 'convert integer to text';

CREATE FUNCTION pg_catalog.text(bigint) RETURNS text STRICT IMMUTABLE LANGUAGE SQL AS
'SELECT textin(int8out($1));';
CREATE CAST (bigint AS text) WITH FUNCTION pg_catalog.text(bigint) AS IMPLICIT;
COMMENT ON FUNCTION pg_catalog.text(bigint) IS 'convert bigint to text';
```

! This change is mandatory for PostgreSQL >= 8.3 since casts have been simplified. If you don't change this option you will have `operator or does not exist` exceptions.

Possible errors if you don't update the casts as described above are:

```
org.postgresql.util.PSQLException: ERROR: operator does not exist: integer = character
varying

org.postgresql.util.PSQLException: ERROR: operator does not exist: bigint = character
varying

com.hp.hpl.jena.db.RDFRDBException: Exception while checking db format -
com.hp.hpl.jena.db.RDFRDBException: Internal SQL error in driver -
org.postgresql.util.PSQLException: ERROR: current transaction is aborted, commands
ignored until end of transaction block

com.hp.hpl.jena.db.RDFRDBException: Internal SQL error in driver -
org.postgresql.util.PSQLException: ERROR: current transaction is aborted, commands
ignored until end of transaction block
```

For further details, please see [this url](#). You might also be interested in this migration helper.

Language plpgsql

If not already done, if you have PostgreSQL < 9.0 you must enable the `plpgsql` language:

```
CREATE LANGUAGE 'plpgsql';
```

Execute this on the `template1` database, so that any database created afterward will get the required language. If your database is already created, execute the command in your database as well.

If you get the following error then it just means that the language is already created (which is the case since PostgreSQL 9.0) and there is nothing further to do:

```
ERROR: language "plpgsql" already exists
```

Create the role and database for Nuxeo

For instance (please change the password and the `nuxeo.conf` file of your instance accordingly):

```
$ createuser -U postgres -W -P nuxeo
$ createdb -U postgres -W -O nuxeo -E UTF8 nuxeo
```

Or from the `psql` command prompt:

```
CREATE ROLE nuxeo WITH PASSWORD 'nuxeo' LOGIN;
CREATE DATABASE nuxeo ENCODING 'UTF8' OWNER nuxeo;
```



Note that using the `UTF8` encoding for your database is important.

Specific configuration

Accent-insensitive fulltext search

If you want accent-insensitive fulltext search, you'll need to install the *unaccent* contribution, create a new text search configuration, and specify its use in Nuxeo.

Unaccent is described here <http://www.postgresql.org/docs/9.0/static/unaccent.html>.

Install it by running `unaccent.sql` script. For Ubuntu users, this file is located at `/usr/share/postgresql/9.0/contrib/unaccent.sql`

Connect to your database and run the following instructions:

```
CREATE TEXT SEARCH CONFIGURATION fr ( COPY = french );
ALTER TEXT SEARCH CONFIGURATION fr ALTER MAPPING FOR asciihword, asciiword,
hword_asciipart, hword, hword_part, word WITH unaccent, french_stem;
```

Then replace in your `default-repository-config.xml` file the `french` analyzer by the one you just defined (`fr` in this example).

Performance tuning

Shared buffers and system cache

One of the most important thing for PostgreSQL is to have lots of shared buffers along with free memory that can be used by the system cache.

If you plan to use 1 GB of shared buffers, update the following property in your `postgresql.conf` file:

```
shared_buffers = 1GB
```

The shared memory is dedicated to PostgreSQL and must be available on the system side using `sysctl`. You need to enable a little bit more at the OS level, for instance try 1 GB + 128 MB:

```
sysctl kernel.shmmax=1207959552
```

Then restart the PostgreSQL.

If there is no enough shared memory you will have an explicit error message and you should try with a bigger `kernel.shmmax` value.

Once PostgreSQL is started the retained `shmmax` value, should be registered in the `/etc/sysctl.conf` file by adding the following line.

```
kernel.shmmax = <SHMMAX_VALUE>
```

PostgreSQL needs to know how much memory the system will use for disk caching. This is used as a hint when executing queries, this memory is not allocated by PostgreSQL.

To set `effective_cache_size` value, you need to run your application once and check how much memory is used by system cache. This can be done using the `free` command and by summing buffers and cached values. The value is `shared_buffers` plus the amount of OS cache.

```
effective_cache_size = 2GB
```

Memory for workers

Increasing the `work_mem` parameter allows PostgreSQL to do larger in-memory sorts which is much faster than disk sorts. Have in mind that wor

`work_mem` size will be taken by each connection (a pool of 20 connections will take $20 * \text{work_mem}$).

```
work_mem = 32MB
```

Increasing the `maintenance_work_mem` will speed up the vacuum procedure.

```
maintenance_work_mem = 128MB
```

Buffering writes

The default `wal_buffers` can be increased to improve write access time:

```
wal_buffers = 8MB
```

Mass import specific tuning

When doing mass import you can disable the fulltext trigger and fulltext index. They must be dropped after a successful login on a running Nuxeo DM because DDL SQL commands are executed on the first access.

```
ALTER TABLE fulltext DISABLE TRIGGER nx_trig_ft_update;
DROP INDEX IF EXISTS fulltext_fulltext_idx;
```

After the import you can update the fulltext column like this:

```
ALTER TABLE fulltext ENABLE TRIGGER nx_trig_ft_update;
-- Let the trigger update the fulltext column
UPDATE fulltext SET fulltext = ''::TSVECTOR WHERE length(fulltext) is NULL;
CREATE INDEX fulltext_fulltext_idx ON fulltext USING gin (fulltext);
```

Changing temporary the PostgreSQL configuration during the import provide performance benefits:

```
checkpoint_segments = 16
full_page_writes = off
fsync = off
synchronous_commit = off
```

Please refer to the [PostgreSQL reference manual](#).

Index vs table scan

The `random_page_cost` parameter influences this query planner's choice. The default value 4 is too high and can result in a wrong bet on large table. You can lower the cost to 2.

```
random_page_cost = 2
```

Updating the planner statistics

PostgreSQL computes statistics on table content in order to plan for the best performance when executing queries with joins and complex filters. The default configuration in PostgreSQL <= 8.3 is `default_statistics_target` set to the value 10 which can lead to not accurate enough estimates. In 8.4 this value is now set to 100 by default. To set it to 100 in 8.3 instances, just use:

```
SET default_statistics_target = 100;
```

On a running instance you can check that settings with:

```
SHOW default_statistics_target;
```

And then re-execute `ANALYZE` to update the statistics.

Vacuuming

The autovacuum is enabled by default since PostgreSQL 8.3.

Exceptionally, a full vacuum can be done at downtime to recover disk space, it should be followed with a `reindexdb` command.

Monitoring

You can monitor the slowest request with the following configuration:

```
log_line_prefix = '%t [%p]: [%l-1] user=%u,db=%d'
log_min_duration_statement = 400
```

See the PostgreSQL section of the Monitoring page.

Reporting problems

If you have a database configuration problem and you want to fill a JIRA ticket, there are some information to report:

- The PostgreSQL server state: is it dedicated or shared, which OS, how many CPU, RAM, is it a virtual machine...
- How much memory is available on the database server (`free -m` output).
- Amount of Nuxeo documents and PostgreSQL configuration. Using the "psql" command line tool **connect to your Nuxeo database** (and not the default database named `postgres`) and execute the following commands:

```
\o /tmp/pgconf.txt
\timing
SELECT now(), Version();
SELECT current_database() AS db_name,
pg_size_pretty(pg_database_size(current_database())) AS db_size,
pg_size_pretty(SUM(pg_relation_size(indexrelid))::int8) AS index_size FROM pg_index;
SELECT COUNT(*) AS documents_count FROM hierarchy WHERE NOT isproperty;
SELECT primarytype, COUNT(*) AS count FROM hierarchy WHERE NOT isproperty GROUP BY
primarytype ORDER BY count DESC;
```

```
SELECT COUNT(*) AS hierarchy_count FROM hierarchy;
SELECT COUNT(*) AS aces_count FROM acls;
SELECT COUNT(DISTINCT(id)) AS accls_count FROM accls;
SELECT COUNT(*) AS read_accls_count FROM read_accls;
SELECT (SELECT COUNT(*) FROM users) AS users, (SELECT COUNT(*) FROM user2group) AS user2groups,
      (SELECT COUNT(*) FROM groups) AS group, (SELECT COUNT(*) FROM group2group) AS group2group;
SELECT stat.relname AS "Table",
      pg_size.pretty(pg_total_relation_size(stat.relid)) AS "Total size",
      pg_size.pretty(pg_relation_size(stat.relid)) AS "Table size",
      CASE WHEN cl.reltoastrelid = 0 THEN 'None' ELSE
            pg_size.pretty(pg_relation_size(cl.reltoastrelid)+
            COALESCE((SELECT SUM(pg_relation_size(indexrelid)) FROM pg_index WHERE
indrelid=cl.reltoastrelid)::int8, 0)) END AS "TOAST table size",
      pg_size.pretty(COALESCE((SELECT SUM(pg_relation_size(indexrelid)) FROM pg_index WHERE
indrelid=stat.relid)::int8, 0)) AS "Index size",
      CASE WHEN pg_relation_size(stat.relid) = 0 THEN 0.0 ELSE
            round(100 * COALESCE((SELECT SUM(pg_relation_size(indexrelid)) FROM pg_index WHERE
indrelid=stat.relid)::int8, 0) / pg_relation_size(stat.relid)) / 100 END AS "Index
ratio"
FROM pg_stat_all_tables stat
JOIN pg_statio_all_tables statio ON stat.relid = statio.relid
JOIN pg_class cl ON cl.oid=stat.relid AND stat.schemaname='public'
ORDER BY pg_total_relation_size(stat.relid) DESC
LIMIT 20;
SELECT nspname,relname,
      round(100 * pg_relation_size(indexrelid) / pg_relation_size(indrelid)) / 100 AS
index_ratio, pg_size.pretty(pg_relation_size(indexrelid)) AS index_size,
pg_size.pretty(pg_relation_size(indrelid)) AS table_size
FROM pg_index I
LEFT JOIN pg_class C ON (C.oid = I.indexrelid)
LEFT JOIN pg_namespace N ON (N.oid = C.relnamespace)
WHERE
  nspname NOT IN ('pg_catalog', 'information_schema', 'pg_toast') AND C.relkind='i'
AND pg_relation_size(indrelid) > 0
ORDER BY pg_relation_size(indexrelid) DESC LIMIT 15;
SELECT relname, idx_tup_fetch + seq_tup_read AS total_reads
FROM pg_stat_all_tables WHERE idx_tup_fetch + seq_tup_read != 0
ORDER BY total_reads desc LIMIT 15;
\di+
SELECT sum(generate_series) AS "speedTest" FROM generate_series(1,1000000);
EXPLAIN ANALYZE SELECT sum(generate_series) AS "speedTest" FROM
generate_series(1,1000000);
SELECT now() - query_start AS duration, current_query FROM pg_stat_activity
  WHERE current_query != '<IDLE>' ORDER BY duration DESC;
SELECT database, gid FROM pg_prepared_xacts;
SELECT name, unit, current_setting(name), source FROM pg_settings WHERE
source != 'default';
```

```
SHOW ALL;  
\q
```

and attach the output file located in `/tmp/pgconf.txt` into the JIRA ticket.

- If you are monitoring the slowest queries (See monitoring section) you can zip and attach the `postgresql` log file to the JIRA ticket.`\o /tmp/pgconf.txt`

Related content in this documentation

- [Connecting Nuxeo to the database](#)
- [Configuring Oracle](#)
- [Configuring MS SQL Server](#)
- [Configuring PostgreSQL](#)

In other documentation

- [Configure Nuxeo 5.3 with VCS and PostgreSQL](#) (Nuxeo Technical Knowledge Base (FAQ))

Configuring Oracle

Nuxeo supports

Oracle 10g R2 (10.2.0.5) and Oracle 11g R2 (11.2.0.1)

On this page

- [Oracle Text \(fulltext\)](#)
- [Grant on DBMS_CRYPTO](#)
- [Grant on V\\$SESSION and GV\\$SESSION](#)
- [Grants for CREATE TABLE and ALTER TABLE](#)
- [Other grants](#)
- [Character set](#)
- [Import/Export](#)
- [JDBC driver](#)
 - [Related content in this documentation](#)
 - [In other documentation](#)

Oracle Text (fulltext)

Oracle Text needs to be enabled in your database for fulltext indexing, please consult your Oracle documentation.

If you fail to install Oracle Text, you will get on startup the following error:

```
java.sql.SQLException: ORA-29833: indextype does not exist
```

In addition, if you want to configure specific lexers or word lists then check http://download.oracle.com/docs/cd/B19306_01/text.102/b14218/cdata dic.htm for configuration parameters and syntax. Lexers and word lists are used by Nuxeo when configured in its `default-repository-confi g.xml` file.

Grant on DBMS_CRYPTO

Since Nuxeo 5.3.2, you need to grant `DBMS_CRYPTO` execution (replace `nuxeo` with the database user):

```
GRANT EXECUTE ON SYS.DBMS_CRYPTO TO nuxeo;
```

Note that for Oracle running on Amazon RDS, `DBMS_CRYPTO` is now directly accessible and you should simply do:

```
GRANT EXECUTE ON DBMS_CRYPTO TO nuxeo;
```

This is due to optimizations now enabled on Oracle that need hashing functions (MD5), available in this package.

Possible errors if you don't do this grant are:

```
java.sql.SQLException: ORA-06550: line 1, column 7: PLS-00905: object  
NUXEO.NX_REBUILD_READ_ACLS is invalid
```

Grant on V\$SESSION and GV\$SESSION

If you use Nuxeo Clustering (<clustering enabled="true"/>), then you must make sure that your database user has access to the system views V\$SESSION and GV\$SESSION (replace nuxeo with the database user):

```
GRANT SELECT ON SYS.V_$SESSION TO nuxeo;  
GRANT SELECT ON SYS.GV_$SESSION TO nuxeo;
```

You can check that this works as intended by doing, as the database user:

```
SELECT SID FROM V$SESSION WHERE SID = SYS_CONTEXT('USERENV', 'SID');  
SELECT SID FROM GV$SESSION WHERE SID = SYS_CONTEXT('USERENV', 'SID');
```

(V\$SESSION is a public synonym for SYS.V_\$SESSION.)

Note: the view GV\$SESSION is used in recent Nuxeo version instead of V\$SESSION to allow working with Oracle RAC.

Possible errors if you don't do this grant are:

```
java.sql.SQLException: ORA-00942: table or view does not exist
```

Grants for CREATE TABLE and ALTER TABLE

As Nuxeo creates tables in the database at first startup, you need to grant CREATE TABLE and ALTER TABLE to your database user.

```
GRANT CREATE TABLE TO nuxeo;  
GRANT ALTER TABLE TO nuxeo;
```

Other grants

The following more standard grants must also be executed :

```
GRANT CONNECT TO nuxeo;
GRANT RESOURCE TO nuxeo;
```

The following is sometimes needed, if you have several schemas:

```
GRANT SELECT ANY TABLE TO nuxeo;
```

Character set

Your database must be configured with `NLS_CHARACTERSET` set to `AL32UTF8`. If your database character set is not `AL32UTF8`, you may observe incorrect behavior including:

- error while trying to insert null values into `acl_user` table (`ORA-01400: cannot insert NULL into ("HUDSON"."ACLR_USER"."USER_ID")`)
- incorrect storage of accented or special characters,
- no tree structure visible on the left of Nuxeo DM,
- queries returning no document.

To check the character set on your server, execute:

```
SELECT value FROM NLS_DATABASE_PARAMETERS WHERE parameter = 'NLS_CHARACTERSET';
```

If you need to change the character set of your database, please check http://download.oracle.com/docs/cd/B19306_01/server.102/b14225/ch11c_harsetmig.htm.

If for some reason you must use an unsupported character set that is not in the list: `AL32UTF8`, `UTF8`, `US7ASCII`, `WE8DEC`, `WE8ISO8859P1`, `WE8MSWIN1252`, then you will need an additional `orai18n.jar` JAR in your Java class path. Download `orai18n.jar` at http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/htdocs/jdbc_10201.html.

Then add it in the class path for your Nuxeo server. For instance, in JBoss, you just put the jar in `$JBOSS/server/default/lib`. (The file `ora18n.jar` replaces the `nls_charset*.jar` files in the Oracle 9i and older releases.)



Technical explanation

Internally, for security checks, Nuxeo executes SQL statements that need to be passed `ARRAY` objects (for the list of principals and permissions), but if the correct character set is not installed then the Oracle JDBC driver behaves incorrectly and Oracle actually receives empty strings. This in turn results in empty results for the queries as none of the documents will match due to incorrect security checks. The `orai18n.jar` fixes this.

Import/Export

Starting 11gR2, Oracle does not allocate space for a table until the first row gets inserted into the table. What happens is if you take an export of the schema/database, the dump would not include any of the tables that hasn't got any space allocations yet. A configuration change needs to be done to allocate space even more tables with no records.

```
alter system set deferred_segment_creation=false;
```

JDBC driver

Nuxeo needs the Oracle JDBC driver to establish a connection to the database.

The driver can be downloaded from the [Oracle JDBC driver downloads site](#).

We recommend the latest version for 11.2.0.* : `ojdbc6.jar` for JDK 1.6. It is compliant with Oracle 10g.

The driver must be in the `$NUXEHOME/lib` directory.

If you are using the `oracle` template (`nuxeo.templates=oracle` in `nuxeo.conf`), just put the driver in the `$NUXEHOME/templates/ora`

—cle/lib directory.

Related content in this documentation

-  Connecting Nuxeo to the database
-  Configuring Oracle
-  Configuring MS SQL Server
-  Configuring PostgreSQL

In other documentation

-  I can't see tree structure with Nuxeo on Oracle!Queries returns no document on Oracle (Nuxeo Technical Knowledge Base (FAQ))

Configuring MS SQL Server

Nuxeo supports

Microsoft SQL Server 2005 or 2008



Note that currently SQL Server 2008 may crash when working with fulltext queries, see [NXP-6143](#) for details. (This is a bug due to Microsoft, not Nuxeo.)

On this page

- Database collation
- Row Versioning-Based Transaction Isolation
- Fulltext
 - Fulltext catalog
- Additional maintenance operation
 - Related content in this documentation
 - In other documentation

Database collation

To work properly Nuxeo need to have some column with a case sensitive collation.

You need have case sensitive collation (a name with "CS" inside), if this is not the case for your existing database you can change it like this:

```
ALTER DATABASE nuxeo COLLATE French_CS_AS
```

Row Versioning-Based Transaction Isolation

To prevent locking and deadlocking problems you need to enable the row versioning-based isolation levels, this can be done with the following SQL commands:

```
ALTER DATABASE nuxeo SET ALLOW_SNAPSHOT_ISOLATION ON;
ALTER DATABASE nuxeo SET READ_COMMITTED_SNAPSHOT ON;
```

Note that there must be no other open connection in the database until ALTER DATABASE is complete.

If you don't do this you will get the following error at startup:

```
Snapshot isolation transaction failed accessing database 'nuxeo' because snapshot isolation is not allowed in this database. Use ALTER DATABASE to allow snapshot isolation.
```

Fulltext

If you configure a fulltext index in Nuxeo (which is the default), you will need to make sure that your SQL Server instance has Full-Text Search configured (it's an optional component during installation). See <http://msdn.microsoft.com/en-us/library/ms142571.aspx> for details.

Failing to do this will provoke errors like:

SQL Server Msg 7601

Cannot use a CONTAINS or FREETEXT predicate on table or indexed view 'fulltext' because it is not full-text indexed.

SQL Server Msg 7616

Full-Text Search is not enabled for the current database. Use sp_fulltext_database to enable full-text search for the database. The functionality to disable and enable full-text search for a database is deprecated. Please change your application.

The French version of these messages, for reference:

SQL Server Msg 7601

Impossible d'utiliser le prédictat CONTAINS ou FREETEXT sur table ou vue indexée 'fulltext', car il n'y a pas d'index de texte intégral.

SQL Server Msg 7616

La recherche en texte intégral n'est pas activée dans la base de données en cours. Utilisez sp_fulltext_database pour l'activer sur cette base de données. La fonctionnalité de désactivation et d'activation d'une recherche en texte intégral pour une base de données est désapprouvée. Modifiez votre application.

Fulltext catalog

Nuxeo uses a fulltext catalog named `nuxeo` by default, this can be changed in the Nuxeo configuration files (see [configuration details](#)).

Additional maintenance operation

Since 5.4.2 HF05 the SQL Server backend comes with ACL (Access Control List) optimization. This optimization works with cache tables to store rights for each users and keep tracking of documents and rights changes. These data are reset when the server is started.

For long running instance or if you want to perform a hot backup without these unnecessary data, you can invoke the following stored procedure:

```
USE nuxeo;
EXEC dbo.nx_vacuum_read_acls;
```

— Or you can exclude the following tables from your backup:

- aclr
- aclr_modified
- aclr_permissions
- aclr_user_map
- aclr_user

Related content in this documentation

- [Connecting Nuxeo to the database](#)
- [Configuring Oracle](#)
- [Configuring MS SQL Server](#)
- [Configuring PostgreSQL](#)

In other documentation

No content found for label(s) sqlserver.

Connecting Nuxeo to the database

To connect Nuxeo to your database, you need to tell Nuxeo which database template to use and provide the database connection information.

On this page

- [Connecting Nuxeo to the database from the Admin Center](#)
- [Connecting Nuxeo to the database from the Startup wizard](#)
- [Connecting Nuxeo to the database from the nuxeo.conf file](#)
- [Database templates](#)
 - [default](#)
 - [postgresql \(recommended\)](#)
 - [oracle](#)
 - [mssql](#)
 - [mysql](#)
- [Related content in this documentation](#)
- [In other documentation](#)

Connecting Nuxeo to the database from the Admin Center

1. In the Admin Center, click on the **Setup** tab of system information section.
 2. In the **Main information** section, select the target database in the drop down menu.
- A new **Database Information** section is displayed on the page.

The screenshot shows the 'Setup' tab of the Nuxeo Admin Center. At the top, there is a navigation bar with tabs: Host, Nuxeo distribution, Setup (which is selected), Repository statistics, and Users activity. Below the navigation bar, the 'Setup' section is divided into two main sections: 'Main information' and 'DataBase information'.

Main information section fields:

Application URL	http://localhost:8080/nuxeo
Product name	Nuxeo DM
Product version	5.4.0-SNAPSHOT
Bind address	0.0.0.0
Configuration file	/Applications/nuxeo-dm/bin/nuxeo.conf
Data directory	/Applications/nuxeo-dm/nxserver/data
Log directory	/Applications/nuxeo-dm/log
Database	PostgreSQL

DataBase information section fields:

Database name	nuxeo
Database username	nuxeo
Database password	password
Database host URL	localhost
Database host port	5432
Database minimum pool size for Nuxeo datasources	
Database maximum pool size for Nuxeo datasources	20
Database minimum pool size for Nuxeo repository (VCS)	0
Database maximum pool size for Nuxeo repository (VCS)	20

3. Fill in the database connection information.

4. Click on the **Save** button.
5. Restart your server.

Connecting Nuxeo to the database from the Startup wizard

The first time you start your Nuxeo server, a wizard is displayed to help you setup your application. Step 3 is about the database: select the database you want to use in the drop down list and provide the connection information to the database.

nuxeo

Database settings

Nuxeo works with several databases. Here you can select and configure the database you want to use.

Nuxeo comes with an embedded database, designed for testing only. For testing and platform evaluation, leave the default database type.

Step 4	Database type	PostgreSQL
Step 5	Database name	nuxeo
Select packages	Database user	nuxeo
Download	Database password	*****
Summary	Database server host name	localhost
	Database server port	5432

Previous step **Next step**

Copyright © 2011 Nuxeo and its respective authors – Visit [nuxeo.com](#) – Join our [community](#)

Connecting Nuxeo to the database from the `nuxeo.conf` file

By default, the "default" template is enabled on your Nuxeo server (see the [#Database templates](#) section for more information on this template). You need to edit it to change the template to be used.

1. Open your `nuxeo.conf` file with a text editor.



For Windows users

Do not use Office writers, nor Notepad.
Wordpad is fine, Notepad++ and SciTE are good text editors, there are a lot of other text editors.

2. If needed, uncomment or edit the `nuxeo.templates` parameter and replace `default` with the wanted database template's name.
3. Uncomment or edit the parameters below and provide their values:

- `nuxeo.db.name`
- `nuxeo.db.user`
- `nuxeo.db.password`
- `nuxeo.db.host`
- `nuxeo.db.port`



For production or load testing environments

These are the minimum required parameters to enable the Nuxeo server to communicate with the database. For a production or load testing environment, you may need to provide the other commented parameters.

4. Save your modifications.
5. Restart the server.

Database templates

The default available database templates are:

- #default
- #postgresql (recommended)
- #oracle
- #mssql
- #mysql

default

This is the default Nuxeo configuration. It is designed for development or test purpose.

Repository backend: H2

Services backend: Derby

postgresql (recommended)

This is the recommended configuration for production, based on PostgreSQL.

Repository backend: PostgreSQL XA

Services backend: PostgreSQL XA

The PostgreSQL driver is included in the Nuxeo applications by default. However, if needed you can download it from <http://jdbc.postgresql.org/download.html#current>.

The JAR (for instance postgresql-9.0-801.jdbc4.jar) is located in \$JBOSS/server/default/lib/ or \$TOMCAT/lib/.



You can use a later driver with an earlier database version, for instance the 9.0 driver still works with PostgreSQL 8.3 or 8.4.

See the page [Configuring PostgreSQL](#) for more information on the database configuration.

oracle

Repository backend: Oracle XA

Services backend: Oracle

The driver is not included in Nuxeo applications. To install it:

1. Download the appropriate JDBC driver from: <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.
2. The JAR must be placed in \$JBOSS/server/default/lib/ or \$TOMCAT/lib/.



Nuxeo applications have been tested with the JDBC drivers available at http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/htdocs/jdbc_10201.html.

See the page [Configuring Oracle](#) for more information on the database configuration.

mssql

Repository backend: Ms SQL Server XA

Services backend: Ms SQL Server XA

The Open Source JTDS driver must be used (the official Microsoft JDBC driver has problems).

You can download it from: <http://repo2.maven.org/maven2/net/sourceforge/jtds/1.2.2/jtds-1.2.2.jar>.

This JAR must then be placed in \$JBOSS/server/default/lib/ or \$TOMCAT/lib/.

See the page [Configuring MS SQL Server](#) for more information on the database configuration.

mysql

We generally don't recommend to use MySQL (see [Why avoid MySQL?](#)).

Repository backend: MySQL XA

Services backend: MySQL

The JDBC driver (downloadable from <http://www.mysql.com/downloads/connector/j/>) is included in Nuxeo applications and is located in \$JBOSS/server/default/lib/ or \$TOMCAT/lib/.

Related content in this documentation



[Connecting Nuxeo to the database](#)

- Adding custom templates
- Configuration templates
- Configuring Oracle
- Configuring MS SQL Server
- Configuring PostgreSQL
- Configuration parameters index (nuxeo.conf)

In other documentation

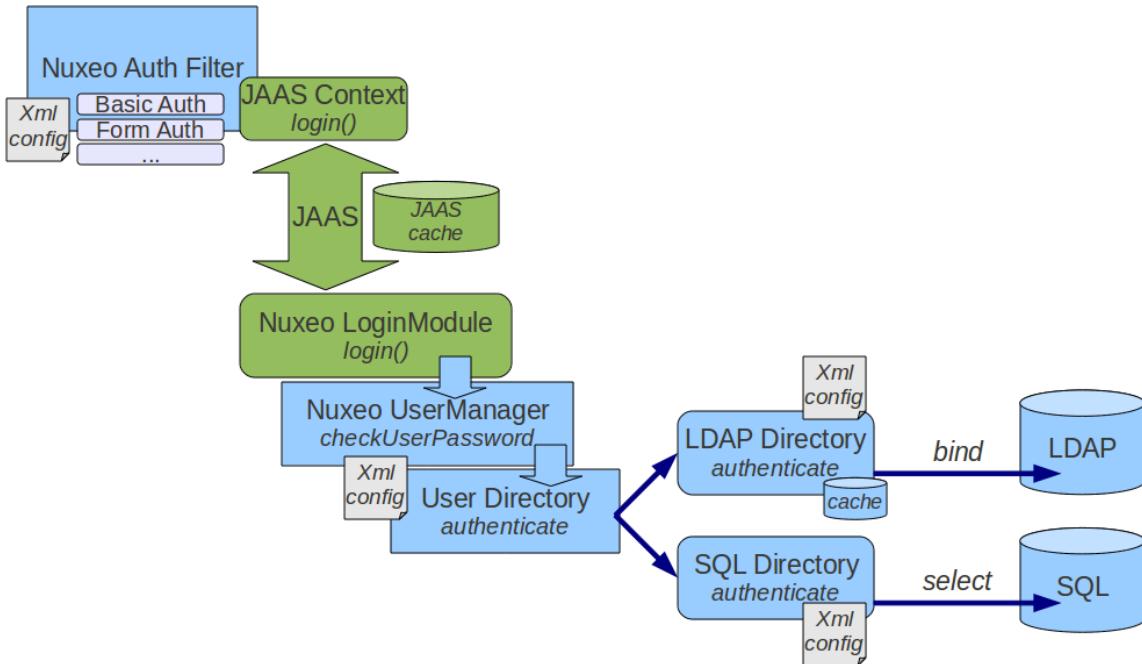
No content found for label(s) database.

Authentication, users and groups

Authentication and Nuxeo EP

Nuxeo EP authentication infrastructure is based on the JAAS standard and has been designed as pluggable as possible so that you can choose how you retrieve user information (identification) and how you validate (authentication).

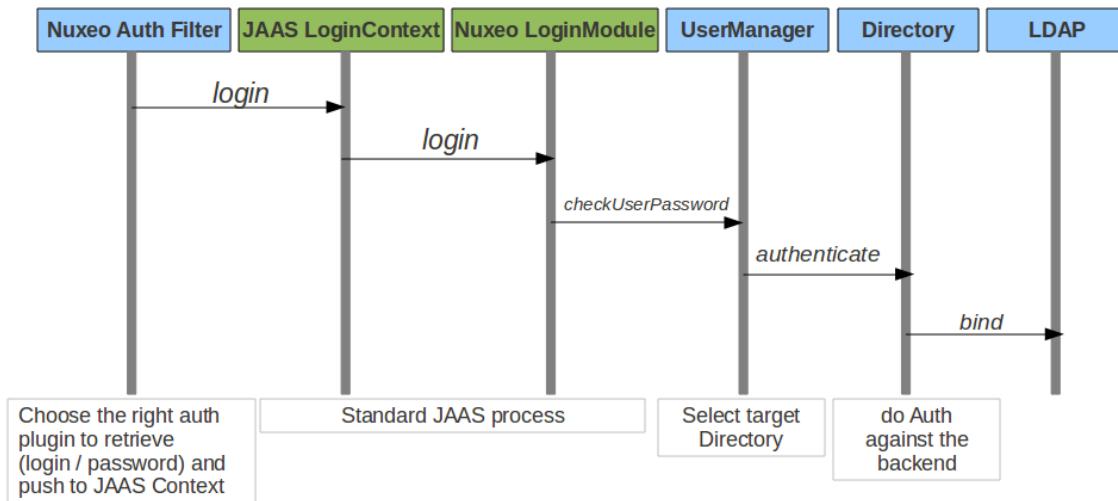
You can see below a schema showing how the global auth process works:



The blue blocks represent the pluggability points:

- retrieving user related information (getting login/password, getting a SSO ticket ...),
- validating user credentials against a backend (SQL DB, LDAP directory, external application ...).

You can see below the flow chart for an authentication.



Built-in Login Plugins

Login plugins are responsible for retrieving the user informations.

It's usually a negotiation between the Nuxeo server and the client browser, but a SSO server may also be part of the process.

By default Nuxeo includes 3 Login Plugins:

- HTTP Basic authentication,
- Form based authentication,
- Anonymous authentication.

Additional Login plugins are available as addons.

When needed, the security filter will determine the right login plugin to use according to:

- what the client browser can provide,
- the server configuration (that can be server wide or specific for some urls).

Built-in LoginModule Plugins

Nuxeo EP uses extension points to define `LoginModule` plugins in order to avoid having to define several `LoginModules`.

By default there are two implementations of the `LoginModule` plugins:

- one that checks Login/Password against the declared directories (SQL or LDAP),
- one that does not check the password and only checks that user exists and fetch user properties (this one is useful when Nuxeo is behind a portal or a SSO server and Nuxeo can not check any password).

OAuth support

Since version 5.4.1, Nuxeo EP provides a built-in support for OAuth.

Please see the dedicated [OAuth Page](#) for more info.

Available authentication modules

nuxeo-platform-login-cas2

The [Central Authentication Service \(CAS\)](#) is a [single sign-on](#) protocol for the [web](#). Its purpose is to permit a user to access multiple applications while providing their credentials (such as user ID and password) only once. It also allows web applications to authenticate users without gaining access to a user's security credentials, such as a password. The name CAS also refers to a software package that implements this protocol. (extracted from wikipedia)

The `nuxeo-platform-login-cas2` defines an authentication plugin to validate the identity using the CAS server. For further information, see [Using CAS2 authentication](#).

nuxeo-platform-login-mod_sso

This plugin is used when Nuxeo is behind a reverse proxy that manage the authentication and simply transmit user information as a set of HTTP headers to Nuxeo.

This is typically the case when:

- Client Certificate authentication is used (Apache does the certificate validation and only transmit to Nuxeo a DN),
- a custom proxy-SSO is used.

nuxeo-platform-login-ntlm

This plugin allows NTLM V1 challenge/response over HTTP.

This plugin does not support NTLM V2 over HTTP and for recent MS Windows auth integration, you should probably use a CAS server with Kerberos.

nuxeo-platform-login-portal-sso

This plugin is used when Nuxeo EP is accessed via an external app (like a portal) that wants to access Nuxeo data in the name of a given user.

Because in most cases the external app does not know the password of the user, this plugin allow to define a shared secret between the app and Nuxeo EP so that the app can access Nuxeo as if it was a given user.

nuxeo-platform-login-shibboleth

The Shibboleth® System is a standards based, open source software package for web single sign-on across or within organizational boundaries. It allows sites to make informed authorization decisions for individual access of protected online resources in a privacy-preserving manner.

The nuxeo-platform-login-shibboleth bundle defines:

- an authentication plugin to map the user metadata from HTTP headers,
- a NuxeoExceptionHandler to force the login of an anonymous user trying to access a restricted resource,
- ShibGroups, virtual groups based on Shibboleth attributes manageable from the UI,
- a hierarchical group suggestion widget for the access rights management tab.

For further information, see [Using Shibboleth](#).

Using a LDAP directory

Principles

In Nuxeo, users and groups are managed by *directories*. If you want your Nuxeo instance to use a LDAP directory you will need to:

- configure a user directory pointing to your LDAP server(s),
- configure a group directory pointing to your LDAP server(s) (if you need LDAP groups).

Of course you can have a specific custom config where:

- you use a custom user / group schema,
- you use several LDAP directories, or a mix of SQL and LDAP directories.

But for the most common use case, all you want to do is map the default `userDirectory` to your LDAP Server. Since groups are used in Nuxeo to associate permissions with content, fetching groups from LDAP is usually not very efficient: LDAP groups are usually not designed for that.

Simple configuration example

Create a file called `default-ldap-users-directory-config.xml` in your config directory

- `server/default/deploy/nuxeo.ear/config/` in JBoss,
- `nxserver/config/` in Tomcat.

Then copy this content (make sure it's valid XML, sometimes what you think is a space character is actually a non-breaking space (U+00A0) which is invalid in XML):

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.directory.ldap.storage.users">

  <require>org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory</require>

  <!-- the groups SQL directories are required to make this bundle work -->
```

```
<require>org.nuxeo.ecm.directory.sql.storage</require>

<extension target="org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory"
point="servers">

    <!-- Configuration of a server connection
    A single server declaration can point to a cluster of replicated
    servers (using OpenLDAP's slapd + slurpd for instance). To leverage
    such a cluster and improve availability, please provide one
    <ldapUrl/> tag for each replica of the cluster.
    -->
    <server name="default">
        <ldapUrl>ldap://localhost:389</ldapUrl>
        <!-- Optional servers from the same cluster for failover
            and load balancing:

            <ldapUrl>ldap://server2:389</ldapUrl>
            <ldapUrl>ldaps://server3:389</ldapUrl>

            "ldaps" means TLS/SSL connection.
        -->

        <!-- Credentials used by Nuxeo5 to browse the directory, create
            and modify entries.

            Only the authentication of users (bind) use the credentials entered
            through the login form if any.
        -->
        <bindDn>cn=nuxeo5,ou=applications,dc=example,dc=com</bindDn>
        <bindPassword>changeme</bindPassword>
    </server>
</extension>

<extension target="org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory"
point="directories">
    <directory name="userDirectory">
        <server>default</server>
        <schema>user</schema>
        <idField>username</idField>
        <passwordField>password</passwordField>

        <searchBaseDn>ou=people,dc=example,dc=com</searchBaseDn>
        <searchClass>person</searchClass>
        <!-- To additionally restrict entries you can add an
            arbitrary search filter such as the following:

            <searchFilter>(&(sn=toto*)(myCustomAttribute=somevalue))</searchFilter>

            Beware that "&" writes "&" in XML.
        -->

        <!-- use subtree if the people branch is nested -->
        <searchScope>onelevel</searchScope>

        <!-- using 'subany', search will match *toto*. use 'subfinal' to
            match *toto and 'subinitial' to match toto*. subinitial is the
            default behaviour-->
        <substringMatchType>subany</substringMatchType>
    </directory>
</extension>
```

```
<readOnly>false</readOnly>

<!-- comment <cache* /> tags to disable the cache -->
<!-- cache timeout in seconds -->
<cacheTimeout>3600</cacheTimeout>

<!-- maximum number of cached entries before global invalidation -->
<cacheMaxSize>1000</cacheMaxSize>

<!--
    If the id field is not returned by the search, we set it with the searched
    entry, probably the login.
    Before setting it, you can change its case. Accepted values are 'lower' and
    'upper',
        anything else will not change the case.
-->
<missingIdFieldCase>lower</missingIdFieldCase>

<!-- Maximum number of entries returned by the search -->
<querySizeLimit>200</querySizeLimit>

<!-- Time to wait for a search to finish. 0 to wait indefinitely -->
<queryTimeLimit>0</queryTimeLimit>

<creationBaseDn>ou=people,dc=example,dc=com</creationBaseDn>
<creationClass>top</creationClass>
<creationClass>person</creationClass>
<creationClass>organizationalPerson</creationClass>
<creationClass>inetOrgPerson</creationClass>

<rdnAttribute>uid</rdnAttribute>
<fieldMapping name="username">uid</fieldMapping>
<fieldMapping name="password">userPassword</fieldMapping>
<fieldMapping name="firstName">givenName</fieldMapping>
<fieldMapping name="lastName">sn</fieldMapping>
<fieldMapping name="company">o</fieldMapping>
<fieldMapping name="email">mail</fieldMapping>

<references>
    <inverseReference field="groups" directory="groupDirectory"
        dualReferenceField="members" />
</references>
</directory>
</extension>

<extension target="org.nuxeo.ecm.platform.usermanager.UserService"
point="userManager">
    <userManager>
        <defaultAdministratorId>johndoe</defaultAdministratorId>
        <defaultGroup>members</defaultGroup>
    </userManager>
```

```
</extension>
</component>
```

Then you should edit this file:

- to set the correct server
 - <ldapUrl>
 - <bindDn> and <bindPassword>
- to set the correct LDAP config:
 - <searchBaseDN>
 - <searchClass>
 - <fieldMapping>

If you want Nuxeo to be able to create users in the LDAP directory, you will need to:

- make sure the user you use to access LDAP has write access,
- define the <creationBaseDn> and associated parameters.

The third contribution is here to define the default mapping:

- since the *Administrator* user won't exists anymore, you should assign at least one user to be administrator using <defaultAdministratorId>,
- you can also choose to make all users members of the default "members" group using <defaultGroup>.

Once done, restart the Nuxeo server, and you should now be able to authenticate against LDAP.



If you want to roll back the changes, simply delete the `default-ldap-users-directory-config.xml` file and restart the server.

For a more detailed view about possible configuration see:

- [LDAPDirectory and associated extension points](#)
- [UserManager extension point](#)

The `ldaptools/` folder in source code of the `nuxeo-platform-directory-ldap` module further provides sample ldif files and OpenLDAP configuration file to help you setup a sample OpenLDAP server you can use as a base setup to build your corporate directory.

Using Active Directory

If you use Active Directory and want to use it with Nuxeo you will need to:

- be sure that LDAP mode is enabled on the Active Directory server,
- get the schema info (because Active Directory schema changes depending on a lot of external factors).

Once you have this information, you can connect Nuxeo to Active Directory as it was a real LDAP server.

Active Directory users are advised to use the aggregated global catalog port number (3268 by default) instead of the default LDAP port (389) in order to avoid getting referrals request to sub directories blocked by corporate firewalls.

Usually with AD you will have to map the field "username" to "sAMAccountName".

Known issues

LDAP contribution not activated

Since Nuxeo 5.4.2, <https://jira.nuxeo.com/browse/NXP-6574> prevents the LDAP contribution from being activated.

A quick solution is to put in comments the "`<directory name="userDirectory">...</directory>`" part in `templates/common/config/default-sql-directories-bundle.xml` (or overwrite that file with a custom template).

A cleaner workaround is to define directories whose name are different from the default ones (userDirectory for users, groupDirectory for groups). Then you need to use the user manager to specify the name of the directories which will be used for authentication, searching, ... Therefore you should apply the changes described below to your existing LDAP contributions:

```
<!-- directory for users -->
<directory name="userLdapDirectory">
  (...)

  <inverseReference field="groups" directory="groupLdapDirectory"
    dualReferenceField="members" />
</directory>

<!-- directory for groups -->
<directory name="groupLdapDirectory">
  (...)

  <ldapReference field="members" directory="userLdapDirectory"
    forceDnConsistencyCheck="false" staticAttributeId="uniqueMember"
    dynamicAttributeId="memberURL" />

    <ldapReference field="subGroups" directory="groupLdapDirectory"
      forceDnConsistencyCheck="false" staticAttributeId="uniqueMember"
      dynamicAttributeId="memberURL" />
  (...)

</directory>

<!-- definition in the user manager -->
<extension target="org.nuxeo.ecm.platform.usermanager.UserService"
  point="userManager">
  <userManager>
    (...)

    <users>
      <directory>userLdapDirectory</directory>
    </users>
    (...)

    <groups>
      <directory>groupLdapDirectory</directory>
    </groups>
    (...)

  </userManager>
</extension>
```

See attached files for templates of LDAP configuration.

This method applies to multi-directories too.

Debug information

If you encounter some difficulties configuring LDAP, the first step is to get more details about what happens.

In the Log4J configuration, increase the log level for org.nuxeo.ecm.directory and org.nuxeo.runtime.model.impl:

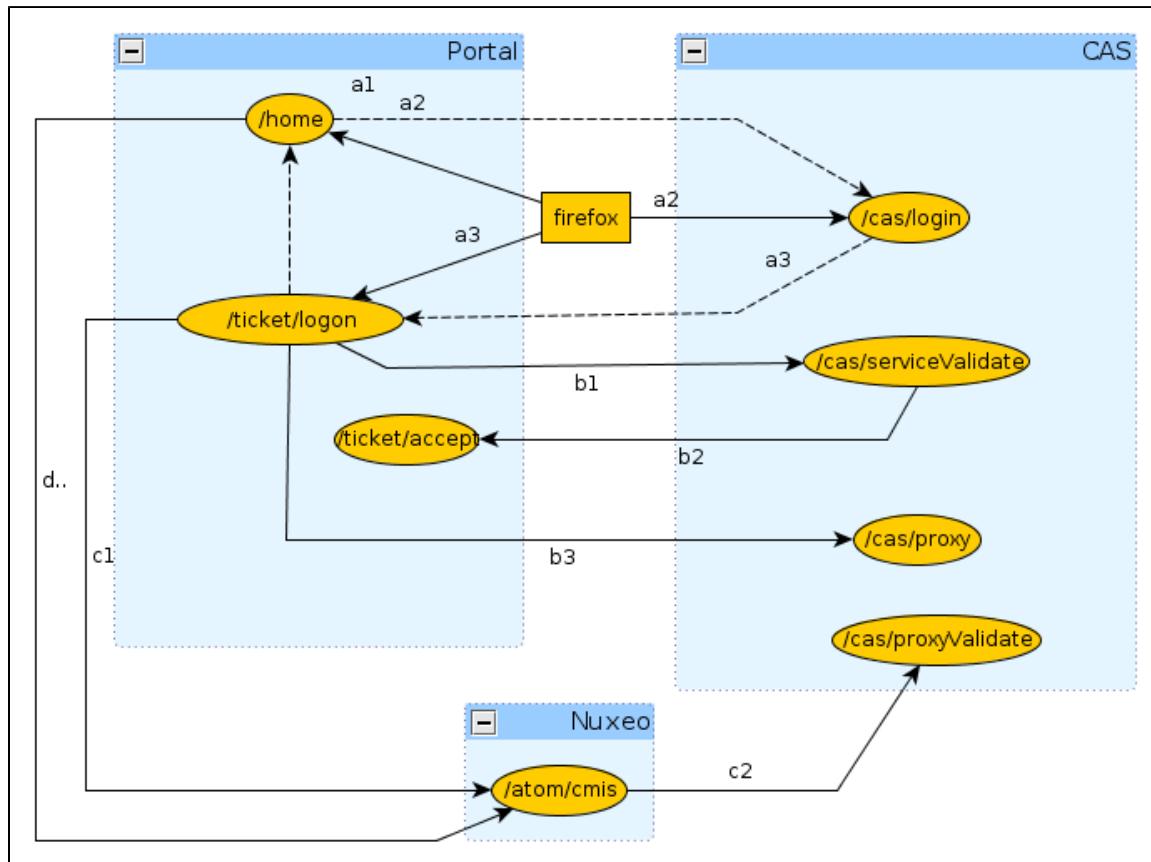
```
<category name="org.nuxeo.ecm.directory">
  <priority value="DEBUG" />
</category>
<category name="org.nuxeo.runtime.model.impl">
  <priority value="INFO" />
</category>
```

This will give you more informations such as:

- Is your XML contribution properly loaded ? Search for the component name of your contribution in the log file (for instance "org.nuxeo.ecm.directory.ldap.storage.users").
- Did the LDAP directory initialized ? If so, your "servers" extension point is working.
- What is the LDAP request sent when you try to log in Nuxeo ? You must be run the same request outside Nuxeo, using your preferred LDAP tool.

Using CAS2 authentication

A typical CAS use case would be the portal. In this n-tiers architecture, the identity is to be shared between the components. The following diagram depicts the interactions between a client, a portal, a CAS server and nuxeo for establishing the authentication context.



(a) Portal service ticket

The first phase is the portal authentication (a1).

```
GET /home
```

The client is redirected to the CAS server for entering its credentials (a2).

```
GET /cas/login?service=http://127.0.0.1:9090/ticket/validate
```

Once the credentials are entered, if they are valid, the CAS server generates a service ticket `ST`. The client is redirected by the CAS server back onto the portal using the service URL (a3).

i In the same time, CAS server generates a ticket granting and registers it client side using the cookie `CASTGC`. If the cookie is already present in the request headers, then the client is automatically redirected to the portal.

```
http://127.0.0.1:9090/ticket/validate?ticket=ST-81-rCbbm5oj9geCKjvhNCvJ-cas
```

(b) Proxy granting ticket

In the second phase, the portal validate the service ticket and request for a proxy granting ticket PGT (b1).

```
GET /cas/serviceValidate?ticket=ST-81-rCbbm5oj9geCKjvhNCvJ-cas&
service=http://127.0.0.1:9090/ticket/validate&pgtUrl=http://127.0.0.1:9090
/ticket/accept
```

If the ticket is valid, the CAS server invoke the `pgtUrl` callback with two parameters `pgtIou` and `pgtId` (b2).

```
GET /ticket/accept?pgtIou=PGTIOU-34-jJZH23r2wbKUqbc3dLft-cas&
pgtId=TGT-45-sSnfcQ7A0TXGsQR2gJONm74rObZ0qRQzhENJWTdZJG5rcGN2T5-cas
```

In case of success, the server respond to the portal with the following content

```
<cas:serviceResponse xmlns:cas='http://www.yale.edu/tp/cas'>
.<cas:authenticationSuccess>
..<cas:user>slacoin</cas:user>
..<cas:proxyGrantingTicket>PGTIOU-34-jJZH23r2wbKUqbc3dLft-cas</cas:proxyGrantingTicket>
.</cas:authenticationSuccess>
</cas:serviceResponse>
```

The `pgtIou` is used portal side for retrieving the accepted PGT.

(c) Nuxeo proxy ticket

In the third phase, the portal ask the CAS server for a new service ticket that will give him access to the nuxeo server using the client identity (c1).

```
GET
/cas/proxy?pgt=TGT-45-sSnfcQ7A0TXGsQR2gJONm74rObZ0qRQzhENJWTdZJG5rcGN2T5-c
as&
targetService=http://127.0.0.1:8080/nuxeo/atom/cmis
```

The CAS server generate a new ST and respond to the portal with the following content

```
<cas:serviceResponse xmlns:cas='http://www.yale.edu/tp/cas'>
.<cas:proxySuccess>
..<cas:proxyTicket>ST-82-20eCHgCqvMCvnP6AmZmz-cas</cas:proxyTicket>
.</cas:proxySuccess>
</cas:serviceResponse>
```

Then the proxy ticket is used by the portal for login into nuxeo (c2).

```
GET /nuxeo/atom/cmis?ticket=ST-82-20eCHgCqvMCvnP6AmZmz-cas
&proxy=http://127.0.0.1:9090/ticket/accept
&service=http:127.0.0.1:8080/nuxeo/atom/cmis
```

The nuxeo server validate the ticket by invoking the portal server (c3).

```
GET
/cas/proxyValidate?ticket=ST-82-20eCHgCqvMCvnP6AmZmz-cas&service=http:127.
0.0.1:8080/nuxeo/atom/cmis
```

If the ticket is valid, the CAS server send the following response

```
<cas:serviceResponse xmlns:cas='http://www.yale.edu/tp/cas'>
.<cas:authenticationSuccess>
..<cas:user>slacoin</cas>
..<cas:proxyGrantingTicket>PGTIOU-34-jJZH23r2wbKUqbc3dLft-cas</cas:proxyGrantingTicket
>
..<cas:proxies>
...<cas:proxy>http://127.0.0.1:9090/ticket/accept</cas:proxy>
..</cas:proxies>
.</cas:authenticationSuccess>
</cas>
```

The nuxeo server create an http session and send the atom pub response message

```
<?xml version='1.0' encoding='UTF-8'?>
<app:service xmlns:app="http://www.w3.org/2007/app"
    xmlns:atom="http://www.w3.org/2005/Atom"
    xmlns:cmis="http://docs.oasis-open.org/ns/cmis/core/200908/"
    xmlns:cmisra="http://docs.oasis-open.org/ns/cmis/restatom/200908/">
    <app:workspace>...</app:workspace>
</app:service>
```

The portal save the client context for being able to invoke Nuxeo using the same HTTP session.

(d) Invoking nuxeo

The nuxeo HTTP session id is retrieved from the portal session context and invoked.

```
GET /nuxeo/atom/cmis?repositoryId=default
```

Using OAuth

Table of contents:

- Background information about OAuth
 - Problem to solve
 - 3-Legged OAuth
 - 2-Legged OAuth (Signed Fetch)

- OAuth Signature
- OAuth in Nuxeo EP
 - Managing Service Consumers in Nuxeo
 - Generic consumers
 - The special case of Nuxeo/Shindig
 - Nuxeo OAuth URLs
 - Managing Service providers in Nuxeo
 - Managing Tokens

Background information about OAuth

This section proposes a quick introduction to OAuth concepts.

For a detailed presentation of OAuth, you should read the OAuth 1.0 Protocol specification.

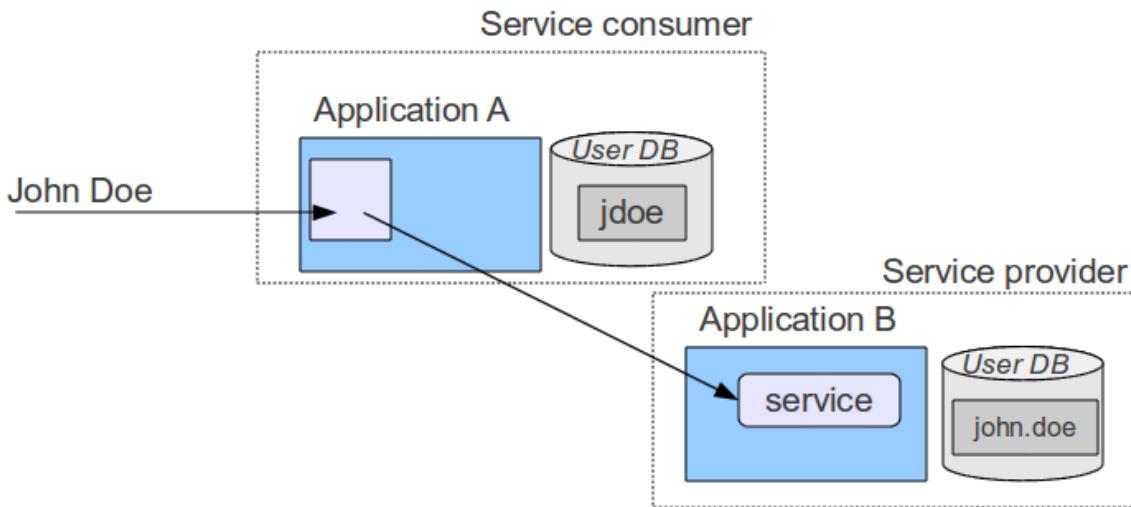
Problem to solve

OAuth addresses authentication issues in cases that include:

- a Service Provider that exposes a web service,
- a User that wants to access the service,
- a Service Consumer that will access the Service Provider on behalf of the user.

In this context, the end user may have different accounts on each application.

The User may also want to grant more or less rights to the consumer. Typically, when Application A access to services of Application B, it may be granted only read accesses. But when using Application C, users wants to grant Read/Write.



3-Legged OAuth

OAuth provides a solution to handle this problem.

The synopsis of the authentication goes like that:

- User wants to access a module of Application A that needs to access a Service hosted by B (for example an OpenSocial Gadget)
- Application A will request a "Request Token" from Application B
- Application B will respond with:
 - a Request Token
 - an authentication URL
- User will be redirected to the authentication URL on the Application B
- User will then
 - authenticate against Application B
 - accept the access request from Application A displayed by Application B (and maybe specify some associated security)
- User will then be redirected to Application A (on the callback url, with the token and a random verifier)
- Application A will receive the request
 - call Application B to exchange the Request Token for an Access Token
 - Application B will only allow the exchange if
 - the verifier is OK
 - the Request Token has been granted by user
 - Application A will then
 - store the Access Token

- use it to access Service in B on behalf of User

The interesting part is that:

- User does not have to give Application A his login/password on Application B
- User can choose how Application A can access his data on Application B
 - define security if needed
 - he should also be able to revoke the grant at any time
- Application A can store the token and reuse it as long as it is valid (no more authentication needed)
- Application A and B do not need to share the same user DB

2-Legged OAuth (Signed Fetch)

Of course 3-Legged OAuth is not always the best option. In some situations, you only want to do a server-to-server authentication.

To answer this use case, OAuth supports a simple mode that will only use the OAuth signature system (see below) to have a trusted communication between two servers.

OAuth Signature

An important part of OAuth is that the http requests have to be signed:

- to be able to verify the identity of the caller
- to be sure the request was not modified during the transfer

For that OAuth can use 2 different signing methods:

- HMAC symmetric signature (Shared Secret)
- RSA1 asymmetric signature (Public /Private keys)

All requests using OAuth (2-Legged or 3-Legged) are signed using one of these 2 signature method.

The signature method is part of the request and may be chosen by the client, but the server may also enforce one of them.

In terms of security both methods have pros and cons, but RSA1 is generally more expensive.

OAuth in Nuxeo EP

Managing Service Consumers in Nuxeo

Generic consumers

Consumers are declared in Nuxeo OAuthConsumerRegistry.

You can access this service via the Admin Center / OAuth / Consumers.



To define a OAuth Consumer you will need to enter:

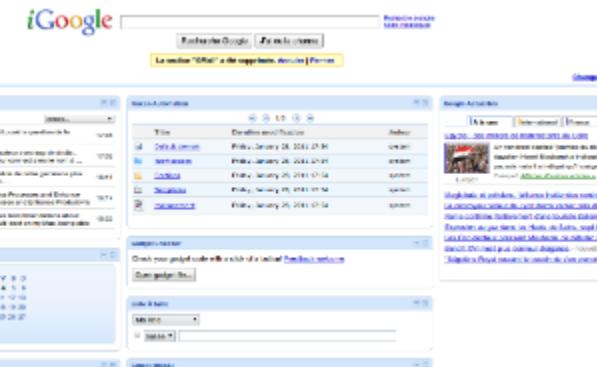
- a Consumer Key: identifier of the application
- a Consumer Secret (HMAC key) and/or a RSA1 public key
- a callback url
 - only needed for 3-Legged OAuth
 - most consumers will automatically provide this callback url as part of the OAuth dialog
- a security mode
 - you can restrain consumer usage to 3-Legged OAuth
 - you can accept 2-Legged OAuth (i.e. only server-to-server authentication), in this case you will need to choose how the nuxeo user will be chosen:
 - OpenSocial viewer
 - OpenSocial owner
 - Dedicated login

Depending on the consumer, it may not allow all possibilities:

- may be only HMAC signature will be supported
- may be only 3-Legged OAuth will be supported
- ...

For example, you may want to use iGoogle as a consumer of some of your Nuxeo Gadgets and services. In this case, you will need to a consumer with:

- www.google.com as consumer key
- use RSA1 key as documented [here](#)
- provide in Nuxeo the callback url <http://oauth.gmodules.com/gadgets/oauthcallback>



The special case of Nuxeo/Shindig

When Nuxeo is both the provider and the consumer, 2-Legged OAuth (signed Fetch) is used with a HMAC Signature. The HMAC secret is dynamically generated at runtime and shared between Shindig and Nuxeo in memory.

Nuxeo OAuth Urls

Access URL: /nuxeo/oauth/access-token (url to request an Access Token)

Request URL: /nuxeo/oauth/request-token (url to request a Request Token)

Authorization URL: /nuxeo/oauth/authorize (url used by a User to grant access)

Managing Service providers in Nuxeo

Service providers are declared in OAuthServiceProviderRegistry.

You can access it via Admin Cener / OAuth / Service providers.

You will have to enter the following information:

- a gadget URL (AppId): the identifier of the gadget inside Nuxeo that will use the service
- a service Name: the name of the service as used in the makeRequest call
- a consumer Key: the consumer key (login) that must be used to call the service
- a consumer Secret (HMAC key): a shared secret used to sign OAuth requests
- OAuth callback urls

OAuth callback urls will only be used if the gadget does not provide the needed information as part of the Gadget Spec (it should). If the consumer secret is empty, Nuxeo will use the server global RSA1 key.

When looking up the right service provider, Nuxeo will do the search based on the service Name and the AppId and fallback to the closest match.

Managing Tokens

OAuth requires Nuxeo to manage Tokens:

- Request Tokens are managed in memory (transient)
- Access Token are managed in the SQL DB

You can use the Admin Center screens to:

- see the Tokens that Nuxeo granted or was granted
- revoke (delete) the Tokens

Using Shibboleth

Introduction

Shibboleth has two major halves: an identity provider (IdP), which authenticates users and releases selected information about them, and a service provider (SP) that accepts and processes the user data before making access control decisions or passing the information to protected applications. These entities trust each other to properly safeguard user data and sensitive resources.

Here is the process of authentication :

1. the user accesses a protected resource
2. the SP determines an IdP and issues the authentication request
3. the user authenticates to the IdP
4. the IdP issues a response to the SP
5. the SP decodes the message and extracts the attributes
6. the browser is redirected to the protected resource

For details, see <https://spaces.internet2.edu/display/SHIB2/FlowsAndConfig>.

Authentication plugin

The SHIB_AUTH plugin is implemented by the class `org.nuxeo.ecm.platform.shibboleth.auth.ShibbolethAuthenticationPlugin`. It authenticates the user based on HTTP headers received from the SP. It also creates (or updates) an entry in the `userDirectory` for this user.

As the Shibboleth attributes values are passed by HTTP headers, the service `org.nuxeo.ecm.platform.shibboleth.service.ShibbolethAuthenticationService` has been added to configure the mapping between the user metadata and the headers names.

Installation

Module installation

The easy way

The Shibboleth authentication module could be installed through the Marketplace with the Shibboleth Authentication package. See [Shibboleth Authentication](#).

The manual way

You can also download directly the built jar, and deploy it into your Tomcat or JBoss instance, in the `bundles` directory.

Download link: [nuxeo-platform-login-shibboleth.jar](#)

If you did the installation by just deploying the `nuxeo-platform-login-shibboleth.jar` into your Nuxeo instance, you need to add a new configuration file to define the login and logout URLs, the mapping between the user metadata and the headers names.

Add a new file named `shibboleth-config.xml` in the `config/` directory of your server.

- `$NUXEO/nxserver/config` for a Tomcat distribution
- `$NUXEO/server/default/deploy/nuxeo.ear/config` for a JBoss distribution

`shibboleth-config.xml`

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.login.shibboleth.config">
<extension
  target="org.nuxeo.ecm.platform.shibboleth.service.ShibbolethAuthenticationService"
  point="config">
<config>
  <uidHeaders>
    <default>uid</default>
  </uidHeaders>

  <loginURL>https://host/Shibboleth.sso/WAYF</loginURL>
  <logoutURL>https://host/Shibboleth.sso/Logout</logoutURL>

  <fieldMapping header="uid">username</fieldMapping>
  <fieldMapping header="mail">email</fieldMapping>
</config>
</extension>
</component>
```

Overriding the default Authentication chain

To enable the Shibboleth authentication, you need to add the Shibboleth plugin to the authentication chain.

To override the default authentication chain in Nuxeo DM, add a new file named `authentication-chain-config.xml` in the `config/` directory of your server.

`authentication-chain-config.xml`

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.your.authentication.chain.config">
<extension
  target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
  point="chain">
<authenticationChain>
  <plugins>
    <plugin>BASIC_AUTH</plugin>
    <plugin>SHIB_AUTH</plugin>
    <plugin>FORM_AUTH</plugin>
    <plugin>WEBENGINE_FORM_AUTH</plugin>
    <plugin>ANONYMOUS_AUTH</plugin>
    <plugin>WEBSERVICES_AUTH</plugin>
  </plugins>
</authenticationChain>
</extension>
</component>
```

If you already define your own authentication chain in any of your config or contrib files, you just need to add the `SHIB_AUTH` plugin into your own chain.

Anonymous authentication compatibility

As it is not possible with Nuxeo to write access rules based on resource URLs, Shibboleth `LazySession` has to be enabled. By adding the `ANONYMOUS_AUTH` in the authentication chain, the Shibboleth login will only be asked when accessing a restricted resource.

For that, the `org.nuxeo.ecm.platform.shibboleth.auth.exceptionhandling.ShibbolethSecurityExceptionHandler` will redirect to the login URL when a `NuxeoSecurityException` is thrown while the current user is anonymous.

To activate it, add a new file named `login-anonymous-config.xml` in the `config/` directory of your server.

`login-anonymous-config.xml`

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.your.anonymous.user.config">
  <extension target="org.nuxeo.ecm.platform.usermodeler.UserService"
    point="userManager">
    <userManager>
      <users>
        <anonymousUser id="Guest">
          <property name="firstName">Guest</property>
          <property name="lastName">User</property>
        </anonymousUser>
      </users>
    </userManager>
  </extension>
</component>
```

Full sample configuration file

Here is a sample configuration file containing everything you need to set up the shibboleth authentication module:

```
<component name="sample.shibboleth.config">

    <require>org.nuxeo.ecm.platform.ui.web.auth.WebEngineConfig</require>
    <require>org.nuxeo.ecm.platform.usermanager.UserManagerImpl</require>

    <extension

target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
    point="chain">
        <authenticationChain>
            <plugins>
                <plugin>BASIC_AUTH</plugin>
                <plugin>SHIB_AUTH</plugin>
                <plugin>ANONYMOUS_AUTH</plugin>
            </plugins>
        </authenticationChain>
    </extension>

    <extension
target="org.nuxeo.ecm.platform.shibboleth.service.ShibbolethAuthenticationService"
    point="config">
        <config>
            <uidHeaders>
                <default>uid</default>
            </uidHeaders>

            <loginURL>https://host/Shibboleth.sso/WAYF</loginURL>
            <logoutURL>https://host/Shibboleth.sso/Logout</logoutURL>

            <fieldMapping header="uid">username</fieldMapping>
            <fieldMapping header="mail">email</fieldMapping>
        </config>
    </extension>

    <!-- Add an Anonymous user -->
    <extension target="org.nuxeo.ecm.platform.usermanager.UserService"
    point="userManager">
        <userManager>
            <users>
                <anonymousUser id="Guest">
                    <property name="firstName">Guest</property>
                    <property name="lastName">User</property>
                </anonymousUser>
            </users>
        </userManager>
    </extension>

</component>
```

Source code

The source code of the Shibboleth authentication module can be found as part of the `nuxeo-platform-login` addon, [here](#).

ShibGroups addon

ShibGroups are virtual groups based on an EL expression with Shibboleth attributes. A new user management tab is added to create and edit

them. The definitions are stored in the `shibbGroup` directory.

The class `org.nuxeo.ecm.platform.shibboleth.computedgroups.ShibbolethGroupComputer` computes at login time the ShibbolethGroups that the current user is member of.

The source code of this addon can be found [here](#).

HTTP and HTTPS reverse-proxy configuration

The Nuxeo webapp can be virtual hosted behind a HTTP/HTTPS reverse proxy, like Apache.

Motivations for virtual hosting

Virtual hosting provides several advantages:

- Support for HTTPS
HTTPS support in Apache is easy and flexible to setup. Apache can also be used to handle certificate authentication.
- URL filtering
You can use Apache filtering tools to limit the URLs that can be accessed via the reverse proxy.
- Handle HTTP cache for static resources
Nuxeo EP generates standard HTTP cache headers for all static resources (images, JavaScript...). These resources are by default cached on the client side (in the browser cache). For performance reasons, it can be useful to host these resources in the reverse proxy cache.

Virtual hosting configuration for Apache 2.x

Reverse proxy with mod_proxy

For this configuration, you will need to load and enable the `mod_proxy` and `mod_proxy_http` modules.

```
ProxyPass /nuxeo/ http://Nuxeo5ServerInternalIp:8080/nuxeo/
ProxyPassReverse /nuxeo/ http://Nuxeo5ServerInternalIp:8080/nuxeo
ProxyPreserveHost On
```

You can also use rewrite rules to achieve the same result:

```
ProxyVia On
ProxyRequests Off
RewriteEngine On
RewriteRule /nuxeo(.*) http://Nuxeo5ServerInternalIp:8080/nuxeo$1 [P,L]
```

This configuration will allow you to access the Nuxeo EP webapp via <http://ApacheServer/nuxeo/>.

The Nuxeo webapp will generate the URLs after reading the `http` header `x-forwarded-host`.

Unfortunately, this header does not specify the protocol used, so if your Apache is responding to HTTPS, you will need to send Nuxeo EP a specific header to indicate the base URL that the webapp must use when generating links.

```
RequestHeader append nuxeo-virtual-host "https://myDomainName/"
```

This will require you to load and activate the `mod_headers` module.

And if you have "client denied by server configuration" error, you must check the access control directives of `mod_proxy`:

```
<Proxy *>
  Order Deny,Allow
  Deny from all
  Allow from 192.168
</Proxy>
```

Reverse proxy a Webengine site to a myexample.com/mysite url

You need the same configuration from the first section. It is advised to first get it to work before proxying exclusively a Webengine site.

A site request queries both from its own url (/nuxeo/site/mysite) but also gets static resources from the root (/nuxeo/nxthemes ...). A rewrite configuration for mysite would look like:

```
RewriteRule ^/nuxeo$    /nuxeo/  [P,L]
RewriteRule ^/mysite$   /mysite/  [P,L]

RewriteCond %{REQUEST_URI}  ^/mysite/skin/mysite/.*
RewriteRule ^/mysite/skin/mysite/(.*) http://127.0.0.1:8080/nuxeo/site/skin/mysite/$1
[P,L]

RewriteCond %{REQUEST_URI}  ^/mysite/skin/.*
RewriteRule ^/mysite/skin/(.*) http://127.0.0.1:8080/nuxeo/site/skin/mysite/$1 [P,L]

RewriteCond %{REQUEST_URI}  ^/mysite/nxthemes-(lib|css)/.*
RewriteRule ^/mysite/(.*) http://127.0.0.1:8080/nuxeo/$1 [P,L]

RewriteCond %{REQUEST_URI}  ^/mysite/nxthemes-webwidgets/.*
RewriteRule ^/mysite/(.*) http://127.0.0.1:8080/nuxeo/site/$1 [P,L]

RewriteRule ^/mysite/(.*) http://127.0.0.1:8080/nuxeo/site/mysite/$1 [P,L]
```

Webengine also needs to know the base of the site, in this case, an empty string instead of /nuxeo/site. This information is passed using the mod_headers:

```
RequestHeader append nuxeo-webengine-base-path ""
```

You can also fetch the static resources from a different path. To do so add a properties to the nuxeo.properties file

```
org.nuxeo.ecm.webengine.skinPathPrefix=/skin/
```

Reverse proxy with mod_jk

! The AJP connector may lock threads if you're not using the APR implementation. Please read the native tomcat documentation for activating the APR implementation on your system. On linux you just have to install the package libtcnative-1.

mod_jk allows you to communicate between Apache and Tomcat via the ajp1.3 protocol.

```
JkWorkersFile /etc/apache2/jk/workers.properties
JkLogFile      /var/log/mod_jk.log
JkLogLevel     info
JkMount /nuxeo ajp13
JkMount /nuxeo/* ajp13
```

The workers.properties file will contain the list of Nuxeo EP Tomcat servers. The AJP13 tomcat listener should be enabled by default on port 8009.

```
worker.list=ajp13
worker.ajp13.port=8009
worker.ajp13.host=Nuxeo5ServerInternalIp
worker.ajp13.type=ajp13
worker.ajp13.socket_keepalive=1
worker.ajp13.connection_pool_size=50
```

Once again, if you use HTTPS, you will need to set the Nuxeo-specific header to tell the webapp how to generate URLs:

```
RequestHeader append nuxeo-virtual-host "https://myDomainName/"
```

This will require you to load and activate the mod_header module.

Configuring http cache

The Simple cache filter is deprecated, we recommend using the `filterConfig` extension point of `RequestControllerService`.

RequestControllerService's filter extension point

This XP lets you contribute customized filter for a given pattern URL.

Example of a filterConfig Registration

```
<extension
target="org.nuxeo.ecm.platform.web.common.requestcontroller.service.RequestControllerService"
point="filterConfig">

<filterConfig name="filterName" transactional="true" synchronize="true"
cached="true" private="true" cachetime="3600">
<pattern>/nuxeo/urlPattern/.*</pattern>
</filterConfig>
</extension>
```

This contribution will ensure that every pattern matching url will go through `NuxeoRequestControllerFilter`. The header of the corresponding request will be modified according to the XP configuration. Here is a list of the possible options:

- `filterConfig`
 - `name`: name of the Filter.
 - `transactional`: use transaction.
 - `synchronize`: is synchronized

- cached: if true, add cache-control to header
- cacheTime: cache duration.
- private: if true, cache is private, public if false.
- pattern: url pattern to match

Using Simple Cache Filter

The Nuxeo webapp includes a Servlet Filter that will automatically add header cache to some resources returned by the server.

By using the deployment-fragment.xml you can also put some specific resources behind this filter:

```
<extension target="web#FILTER">
  <filter-mapping>
    <filter-name>simpleCacheFilter</filter-name>
    <url-pattern>/MyURLsToCache/*</url-pattern>
  </filter-mapping>
</extension>
```

When Nuxeo EP is virtual hosted with apache you can use mod_cache to use the reverse-proxy as cache server.

You can also use Squid or any other caching system that is compliant with the standard HTTP caching policy.

Advanced configuration

The pages below are about advanced configurations that are not mandatory and are used to integrate new features or change low level behaviors.

- Adding custom templates
- Changing context path
- Configure User & Group storage and Authentication
- Firewall consideration
- How to create a Nuxeo static WAR?
- JDBC
- Nuxeo clustering configuration
- VCS Configuration

Adding custom templates

The "custom" template folder allows you to add customization such as using multiple databases, configuring services, ...

1. Add your own template files in `templates/custom` directory.
You can use either existing or new parameters in these new template files.
2. From the Admin Center or by manually editing the `nuxeo.conf` file, set your parameters' values and set `nuxeo.templates=custom`.
You can refer to custom templates directory with a relative path or to your own custom templates directory with an absolute path.
3. Edit `custom/nuxeo.defaults` and set `nuxeo.template.includes` parameter to define the list of existing templates to include (comma separated values); your custom template will be used at last.
`nuxeo.defaults` files from included templates are also read.

In case you need multiple customizations, create multiple directories and reference them in a dedicated `nuxeo.conf` for each server.

Changing context path

Nuxeo offers the capability to change the context path, ie `/nuxeo` in the URL of your application.

This configuration is done with two steps:

1. edit the configuration file `nuxeo.conf` to change the property `org.nuxeo.ecm.contextPath`

```
org.nuxeo.ecm.contextPath=/myapp
```

1. rename the file `$NUXEHOME/templates/default/conf/Catalina/localhost/nuxeo.xml` into `$NUXEHOME/templates/default/conf/Catalina/localhost/myapp.xml`

If you have already started Nuxeo with the old context path, you'll have to remove `$NUXEHOME/conf/Catalina/localhost/nuxeo.xml`

Finally, if you're using one of the quartz-cluster templates, you'll also have to rename the file `nuxeo.xml` under `templates/*-quartz-cluster/conf/Catalina/localhost/`

**Restriction**

this documentation is about Tomcat distribution

Configure User & Group storage and Authentication

Nuxeo EP provides large possibilities about User and Group storage, you can:

- bind Nuxeo EP users and group to ones defined into your LDAP
- bind Nuxeo EP users and group to ones defined into SQL tables (as default configuration)
- bind Nuxeo EP users to ones defined into your LDAP and groups to ones locally stored into a SQL table
- bind Nuxeo EP users to ones defined into an aggregation of 2 LDAP servers and groups to ones locally stored into a SQL table
- bind Nuxeo EP users to ones defined into an aggregation of your LDAP server and a local SQL table and groups to ones locally stored into a SQL table
- ...

Many other more complex possibilities are possible, but this is the most usual ones. If you have other exposition than LDAP and SQL, this is also possible to bind to it.

You have all the documentation about the User Management configuration into [this page](#).

If the authentication against your infrastructure is particular - you have an SSO system, or others - Nuxeo EP have extension points that will let you adapt your Nuxeo application to your infrastructure. The documentation about that is [here](#).

Firewall consideration

The page Firewall consideration does not exist.

How to create a Nuxeo static WAR?

Some important remarks before reading this documentation:

- This documentation is valid since Nuxeo EP 5.5.
- All dynamic features will **not** work (we are into a **static war**):
 - Nuxeo Studio hot reload,
 - Nuxeo IDE hot reload,
 - Nuxeo Marketplace integration (Hotfixes and packages installation).

Nuxeo static WAR distribution creation

Here is the way to create your own static war distribution:

1. Download a Nuxeo Tomcat distribution.
2. Unzip it in somewhere (let's call it \$NUXEO_HOME) .
3. Copy your specific bundles into \$NUXEO_HOME/nxserver/bundles/.
4. Create your [own templates](#) and configuration files as needed.
5. Start your Nuxeo Tomcat instance as usual.
6. Check that your Nuxeo is well configured (database configuration, ldap configuration, etc...).
7. Stop your Nuxeo instance.
8. Launch the following command:

```
$NUXEO_HOME/bin/nuxeoctl pack /tmp/nuxeo-war.zip
```

Your static WAR distribution will be generated into /tmp/nuxeo-war.zip.

The generated ZIP contains what needs to be copied to your Tomcat installation:

```
nuexo-war.zip/
|-- endorsed (jaxb and jaxws api libs that should replace packages provided by default
  JDK that are outdated)
|-- lib (nuxeo common libs)
|-- README-NUXEO.txt
`-- webapps
    `-- nuxeo (exploded WAR that you can zip if needed)
```

Installing Nuxeo in the target Tomcat

The `README-NUXEO.txt` gives you the needed instructions, but basically what you need to do is:

- Copy the structure (endorsed, lib, webapp) inside your target Tomcat.
- Copy the `<Resource>` tag as described in the `README-NUXEO.txt` in your `server.xml` in order to declare the JDBC datasources.

The target database will be the one you defined in your source Nuxeo instance, and the default for the Nuxeo directories are:

- `nuxeo.config.dir`: WEB-INF directory,
- `nuxeo.runtime.home`: `$TOMCAT_HOME/nuxeo`,
- `nuxeo.data.dir`: `$TOMCAT_HOME/nuxeo/data`,
- `nuxeo.tmp.dir`: `$TOMCAT_HOME/nuxeo/tmp`,
- `nuxeo.web.dir`: `$TOMCAT_HOME/nuxeo/web` (for WebEngine modules).

See [Configuration parameters index \(nuxeo.conf\)](#) for more details about these config parameters of Nuxeo.

If you need to change these values, you can add parameters in the file `webapps/nuxeo/WEB-INF/web.xml`.

See [JavaDoc of NuxeoStarter](#) for more details.

JDBC

Nuxeo clustering configuration

Nuxeo can be clustered between several nodes (a.k.a. instances or machines) with the appropriate configuration. In addition, a HTTP load balancer with session affinity must be used in front of the nodes.

Requirements

To enable clustering, you must have at least two nodes with:

- a shared database,
- a shared filesystem (unless you use an external binary store like S3),
- a load-balancer with stick sessions.

The shared filesystem is usually a NFS mount. You **must not** share the whole Nuxeo installation tree, but only the `nxserver/data/binaries` directory.

The load balancer **must** use sticky sessions if the clustering delay is not 0. Having a non-0 clustering delay is recommended for performance reasons. See below for more.

In this section

- Requirements
- VCS cluster configuration
 - Setup
 - Checking the SQL tables initialization
 - Checking VCS cache invalidations
- Quartz scheduler cluster configuration
- HTTP load balancer configuration
 - Troubleshooting session affinity problems
 - Related pages

VCS cluster configuration

Setup

To set up clustering, please update the `repository.clustering.enabled`, `repository.clustering.delay` and `repository.binary.store` in the [nuxeo.conf parameters](#). On all Nuxeo instances, the `repository.binary.store` should point to a shared filesystem unless you use an external binary store like S3.

- `repository.clustering.enabled` must be `true` to enable clustering.
- `repository.clustering.delay` is expressed in milliseconds, and specifies a delay during which invalidations don't need to be processed. Using a non-0 value is an important optimization as otherwise every single transaction, even a read-only one, would have to hit the database to check invalidations between several nodes. However this means that one node may not see immediately the changes made on another node, which is a problem if you don't use sticky session on the load balancer.
- `repository.binary.store` must point to a shared storage. Under Windows, the path value can be UNC formatted, for instance `\servername\sharename`.

There is a dedicated page detailing all the [VCS configuration options](#).

Checking the SQL tables initialization

Start the SQL server, all Nuxeo nodes (the first alone and the other afterwards to avoid concurrent initialization of the SQL tables) and the load balancer and login on the HTTP user interface on each cluster node, then check that on the database that the `cluster_nodes` table is initialized with one line per node:

```
nuxeo-db=# select * from cluster_nodes;
nodeid | created
-----+-----
 25767 | 2009-07-29 14:36:08.769657
 32546 | 2009-07-29 14:39:18.437264
(2 lines)
```

Checking VCS cache invalidations

Create a document and browse it from two different nodes. Edit the title from one node and navigate back to the document from second node to check that the change is visible. You can also monitor what's happening in the `cluster_invals` table to see cache invalidation information.

Quartz scheduler cluster configuration

The Quartz scheduler should be configured to run in a cluster. This is needed in order for scheduled events, like periodic cleanups or periodic imports, to be executed only one node and not on all nodes at the same time, which could cause problems.

Standard configuration is available from Nuxeo templates for Tomcat for PostgreSQL, Oracle and SQL Server.

First, you should populate the database with the tables needed by quartz (names `QRTZ_*`). The DDL scripts come from the standard Quartz distribution and are available in the Nuxeo templates in `$NUXEHOME/templates/<database>-quartz-cluster/bin/create-quartz-tables.sql`.

Second, you should enable the quartz-specific cluster templates by adding the template `<database>-quartz-cluster`.



In cluster mode the schedule contributions (deployed from plugins or configuration files) **must** be the same on all nodes.

HTTP load balancer configuration

Set up an HTTP or AJP load balancer such as Apache with `mod_proxy` or `mod_proxy_ajp` or Pound, and configure it to keep session affinity by tracking the value of the `JSESSIONID` cookie and the "`;jsessionid`" URL parameter.

If you use a stateless load balancer such as apache modules such as `mod_jk` and `mod_proxy_balancer`, you need to make the HTTP server generate `JSESSIONID` cookies with values that end with `.nxworkern`, where `nxworkern` is a string suffix specific to each node (you can use any string).

To do so, in `nuxeo.conf` specify a different `nuxeo.server.jvmRoute` for each node, for instance `nuxeo.server.jvmRoute=nxworker1`. This will instruct the Nuxeo preprocessing phase to correctly fill the `jvmRoute` attribute of the `Engine` element in the generated `server.xml`.

Then configure your stateless balancer to follow these routes, for instance here is the relevant configuration fragment when using `mod_proxy_balancer`:

```
ProxyPass /nuxeo balancer://sticky-balancer stickysession=JSESSIONID|jsessionid
nofailover=On

<Proxy balancer://sticky-balancer>
  BalancerMember http://192.168.2.101:8080/nuxeo route=nxworker1
  BalancerMember http://192.168.2.102:8080/nuxeo route=nxworker2
</Proxy>
```

Troubleshooting session affinity problems

To test that the load balancer forwards the HTTP requests of a given session to the same node you can add a new file on each node (after Tomcat started), `$NUXEHOME/nxserver/nuxeo.war/clusterinfo.html`, with on the first node:

```
<html><body>Node 1</body></html>
```

and on the second node:

```
<html><body>Node 2</body></html>
```

Using a browser with an active Nuxeo session (an already logged-in user), then go to `http://yourloadbalancer/nuxeo/clusterinfo.html` and check that you always return to the same node when hitting the refresh button of the browser.

Related pages

-  [Nuxeo clustering configuration \(Nuxeo Installation and Administration - 5.5\)](#)
-  [VCS Architecture \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)

VCS Configuration

VCS (Visible Content Store) is the default storage engine for Nuxeo documents.

The following are the options available to configure VCS repository in Nuxeo Platform. They usually go in a file named `default-repository-config.xml`.

Example file

This file contains many more options than are necessary by default.

```
<?xml version="1.0"?>
<component name="default-repository-config">
  <extension target="org.nuxeo.ecm.core.repository.RepositoryService"
  point="repository">
    <repository name="default"
      factory="org.nuxeo.ecm.core.storage.sql.coremodel.SQLRepositoryFactory">
      <repository name="default">
        <clustering enabled="true" delay="1000" />
        <schema>
          <field type="largetext">note</field>
        </schema>
        <indexing>
          <includedTypes>
            <type>File</type>
            <type>Note</type>
          </includedTypes>
          <!-- sample for excluded types -->
          <!--
          <excludedTypes>
            <type>Root</type>
            <type>Workspace</type>
          </excludedTypes>
          -->
          <fulltext analyzer="english"> <!-- PostgreSQL -->
            <index name="default">
              <!-- all props implied -->
            </index>
            <index name="title">
              <field>dc:title</field>
            </index>
            <index name="description">
              <field>dc:description</field>
            </index>
          </fulltext>
          <queryMaker class="org.nuxeo.ecm.core.storage.sql.NXQLQueryMaker" />
          <queryMaker class="org.nuxeo.ecm.core.chemistry.impl.CMISQLQueryMaker" />
        </indexing>
        <binaryStore path="binaries"/>
        <binaryManager class="org.nuxeo.ecm.core.storage.sql.DefaultBinaryManager"/>
        <usersSeparator key="," />
        <aclOptimizations enabled="true" readAclMaxSize="4096" />
        <pathOptimizations enabled="true" />
        <noDDL>false</noDDL>
      </repository>
    </repository>
  </extension>
</component>
```

Clustering options

```
<clustering enabled="true" delay="1000" />
```

- clustering enabled: use **true** to activate Nuxeo clustering (default is **false**, i.e., no clustering).
- clustering delay: a configurable delay in milliseconds between two checks at the start of each transaction, to know if there are any remote invalidations.

Schema options

```
<schema>
  <field type="largetext">note</field>
  <field type="largetext">my:field</field>
  ...
</schema>
```

- field type="largetext": a field that should be stored as a CLOB column inside the database instead of a standard VARCHAR column.

This is important for your large text fields, especially for MySQL, Oracle and SQL Server which have very small defaults for standard text fields.

Using Oracle, if you attempt to save a string too big for the standard NVARCHAR2(2000) field, you will get the error:

```
java.sql.SQLException: ORA-01461: can bind a LONG value only for insert into a LONG
column
```

Indexing options

Configure which types will be indexed

Since Nuxeo DM 5.5 it is possible to configure the document types you want to index or you want to exclude from fulltext indexing. This is possible using the tags `includedTypes` and `excludedTypes` inside the `indexing` tag:

```
<includedTypes>
  <type>File</type>
  <type>Note</type>
</includedTypes>
```

or

```
<excludedTypes>
  <type>Root</type>
  <type>Workspace</type>
</excludedTypes>
```

If you set both included and excluded types, only the included types configuration will be taken into account.

Fulltext

```
<fulltext disabled="true" analyzer="english" catalog="...">
  ...
</fulltext>
```

- fulltext disabled: use **true** to disable fulltext support, the repository configuration must be updated to have (default is **false**, i.e., fulltext enabled).
- fulltext analyzer: a fulltext analyzer, the content of this attribute depends on the backend used:
 - H2: a Lucene analyzer, for instance `org.apache.lucene.analysis.fr.FrenchAnalyzer`. The default is an english analyzer.
 - PostgreSQL: a Text Search configuration, for instance `french`. The default is `english`. See <http://www.postgresql.org/docs/8.3/static/textsearch-configuration.html>
 - Oracle: an Oracle PARAMETERS for fulltext, as defined by http://download.oracle.com/docs/cd/B19306_01/text.102/b14218/cda_tadic.htm (see NXP-4035 for details).
 - Microsoft SQL Server: a fulltext LANGUAGE, for instance `english`, as defined in [http://msdn.microsoft.com/en-us/library/ms187787\(v=SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms187787(v=SQL.90).aspx). The default is `english`.
 - other backends don't have configurable fulltext analyzers.
- fulltext catalog: a fulltext catalog, the content of this attribute depends on the backend used:
 - Microsoft SQL Server: a fulltext CATALOG, the default is `nuxeo`.
 - other backends don't need a catalog.

Fulltext indexes are queried in NXQL through the `ecm:fulltext` pseudo-field. A non-default index "foo" can be queried using `ecm:fulltext_foo`.

If no `<index>` elements are present, then a **default** index with all string and blob fields is used.

```
<fulltext ...>
  <index name="title" analyzer="..." catalog="...">
    <field>dc:title</field>
    <field>dc:description</field>
  </index>
  <index name="blobs">
    <fieldType>blob</fieldType>
  </index>
  <index name="other">
    <fieldType>string</fieldType>
    <excludeField>dc:title</excludeField>
  </index>
</fulltext>
```

- index name: the name of the index (the default is **default**).
- index analyzer: a fulltext analyzer just for this index. See fulltext options above for details.
- index catalog: a fulltext catalog just for this index. See fulltext options above for details.
- fieldType: **string** or **blob**, the default being both. This selects all these fields for indexing.
- field: the name of a field that should be selected for indexing.
- excludeField: the name of a field that should not be in the index.

If no `<fieldType>`, `<field>` or `<excludeField>` is present, then all string and blob fields are used.

Query Maker

```
<queryMaker class="..."/>
```

- queryMaker class: registers a `QueryMaker`, to extend the query system. The class must implement `org.nuxeo.ecm.core.storage.sql.QueryMaker` (the default is `org.nuxeo.ecm.core.storage.sql.NXQLQueryMaker`, i.e., the standard NXQL QueryMaker).

There can be several `<queryMaker>` elements defined.

This is not needed (and deprecated) starting with Nuxeo EP 5.3.2.

Binary Store

```
<binaryManager class="org.nuxeo.ecm.core.storage.sql.XORBinaryManager" key="abc" />
```

- **binaryManager class:** the default Binary Manager can be changed using this (the default is to use the standard binary manager that stores files normally). A new XORBinaryManager has been added, it knows how to do XOR with a pattern on read/write (see the key below). The on-disk binary store is unchanged (the hash of the files is still the filename), but of course it's now unreadable by humans by default. One consequence is that for the same file the application-level digest in the Binary object is now different if encryption is enabled.
- **binaryManager key:** the encryption key for the binary manager (if it's doing any encryption). Changing this value will of course render existing binaries unreadable.

```
<binaryStore path="/foo/bar" />
```

- **binaryStore path:** the filesystem path where the binary store should live. A relative path is interpreted relative to the Nuxeo Framework home. The default is the **binaries** directory.

Optimizations

```
<pathOptimizations enabled="false" />
```

- **pathOptimizations enabled:** for PostgreSQL, Oracle and MS SQL Server (and H2), it is possible to disable the path-based optimizations by using **false** (the default is **true**, i.e., path optimizations enabled).

```
<aclOptimizations enabled="false" />
```

- **aclOptimizations enabled:** for PostgreSQL, Oracle and MS SQL Server (and H2), you can disable the read ACL optimizations by using **false** (the default is **true**, i.e., ACL optimizations enabled).
- since DM 5.4.1, you can set the property `readAclMaxSize` to define the size of the larger ACL for a document : this may be useful if you have mainly affected permissions to a lot of users, instead of using groups (do not set this attribute if you disable ACL optimizations).

```
<usersSeparator key="," />
```

- in case the user/group names in your directories contains the separator character used in the Read ACL cache(comma by default), you can change this value using the attribute `usersSeparator`
- if you change this value on an existing database, you will need to rebuild the ACL cache with the SQL command:
`SELECT nx_rebuild_read_acls();`

Database creation option

```
<noDDL>true</noDDL>
```

Set the value `noDDL` to **true** to execute no DDL (Data Definition Language). The default is **false**.

When this is **true**, VCS will assume that no new structure has to be created in the database. This means that none of these statements will be executed:

- `CREATE TABLE`, `CREATE INDEX`, `ALTER TABLE ADD CONSTRAINT` for a new schema or complex property,
- `ALTER TABLE ADD column` for a new property in a schema,
- `CREATE FUNCTION`, `CREATE PROCEDURE`, `CREATE TRIGGER` for VCS internal stored procedures and migration steps.

The only statements that VCS will execute are:

- INSERT, UPDATE, DELETE for data changes,
- calling of stored procedures.

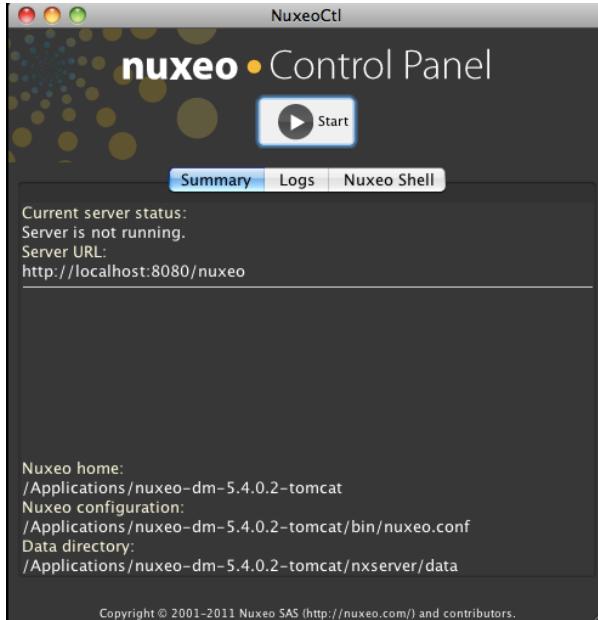
This means that all tables, indexes, triggers and stored procedures needed by VCS have to be created beforehand, either by a previous execution when the flag was **false**, or by a manual execution of a SQL script from a previously-created Nuxeo instance.

This option is typically needed if you configure the VCS connection with a username who is not the owner of the database, usually for security considerations.

Server start and stop

On this page, you will see how to start and stop your Nuxeo application.

Nuxeo applications come with a Control Panel that allows you to start and stop the server easily, and to access more administration features.



The Control Panel gives you access to:

- A summary of the server status: is it running, is stopped, etc...
- The logs of the server: the console and server logs are information of the tasks the server is doing and messages on how it is processing these tasks.
- The **Nuxeo Shell**: the administrators' Swiss Army knife.

Here are the different ways to start and stop your Nuxeo application, depending on your OS:

- Starting your Nuxeo application
 - Starting your application on Windows
 - Starting your application on Linux
 - Starting your application on Mac OS X
- Stopping your Nuxeo application

Starting your Nuxeo application



By default, you cannot run two Nuxeo applications at the same time. If you want to run two Nuxeo applications at the same time (for instance a Nuxeo DM and a Nuxeo DAM), you need to [change the default port](#) used by one of the Nuxeo servers.

Depending on your OS, there are different ways to start the application. The steps below show how to use the Control Panel to start the server. However, you can use the command below in a terminal if you prefer. From the `$NUXEO_HOME/bin`, execute `nuxeoctl start`. You can refer to the [nuxeoctl](#) and [Control Panel usage](#) for more information on the `nuxeoctl` command and the Control Panel.

Starting your application on Windows

1. Open the Nuxeo Control Panel:
 - In the folder `C:\Nuxeo application`, double-click on `Start Nuxeo.bat`.
 - In the folder `C:\Nuxeo application\bin`, double-click on `nuxeoctl.bat`.

The Nuxeo Control Panel opens.

2. Click on the **Start** button.

Starting the Nuxeo server takes between a few seconds and several minutes, depending on your hardware and the distribution you have chosen to install.

When the server is started, the **Start** button becomes a **Stop** button.

3. Open a browser and type the URL <http://localhost:8080/nuxeo/>.

If the server is started for the first time after the installation, the **startup wizard** is displayed so you can select what module you want to install on the platform and help you configure it.

Otherwise, the login page is displayed so you can use the application.



On Windows 7, you need to run the `nuxeoctl.bat` and `Start Nuxeo.bat` commands as an administrator if you haven't installed your Nuxeo application at the root of C: (for instance in C:\Program Files). To run them as an administrator, right-click on the command and click on "Run as administrator".

On Windows, it is possible to start Nuxeo as a service. Please report the [Installing the Nuxeo Platform as a Windows service](#) page for guidelines and examples.

Starting your application on Linux

Nuxeo applications are started using scripts.

1. Launch a terminal and go to your installation directory.
2. Start the server using the `nuxeoctl` script (located in the `bin` directory):

```
./bin/nuxeoctl gui
```



The command used to launch the Control Panel may not be executable by default. If it is the case, in the terminal go to the `bin` directory of Nuxeo and type the line below to be able to use it:

```
chmod +x *.sh *ctl
```

The Control Panel opens.

3. Click on the **Start** button.

Starting the Nuxeo server takes between 30 sec and several minutes, depending on your hardware and the distribution you have chosen to install.

When the server is started, the **Start** button becomes a **Stop** button.

4. Open a browser and type the URL <http://localhost:8080/nuxeo/>.

If the server is started for the first time after the installation, the **startup wizard** is displayed so you can select what module you want to install on the platform and help you configure it.

Otherwise, the login page is displayed so you can use the application.

Starting your application on Mac OS X

Mac OS users can use either the [same steps as Linux users](#) or some Mac OS convenient commands (see below).

1. From the Finder, click on "Start Nuxeo.command". You can also drag and drop the start script in the terminal and press Enter.



The command may not be executable by default. If it is the case, in the terminal go to the `bin` directory of Nuxeo and type the line below:

```
chmod +x *.command
```

The Control Panel opens.

2. Click on the **Start** button.

Starting the Nuxeo server takes between 30 sec and several minutes, depending on your hardware and the distribution you have chosen to install.

When the server is started, the **Start** button becomes a **Stop** button.

3. Open a browser and type the URL <http://localhost:8080/nuxeo/>.

If the server is started for the first time after the installation, the **startup wizard** is displayed so you can select what module you want to install on the platform and help you configure it.

Otherwise, the login page is displayed so you can use the application.

Stopping your Nuxeo application

The steps to stop your Nuxeo application are the same for all operating systems.

To stop your server:

1. On the Control Panel, click on the **Stop** button.
Stopping the server takes several seconds. When the server is stopped, the **Stop** button becomes a **Start** button.
2. Close the Control Panel.

 If you started the server using the `nuxeoctl start` command in the terminal, use the `nuxeoctl stop` command to stop it.

nuxeoctl and Control Panel usage

nuxeoctl usage

The `nuxeoctl` script enables various options and commands (explained in details below).

Here is the Shell/Batch script usage:

```
nuxeoctl [options] [gui|nogui]
(help|start|stop|restart|configure|wizard|console|status|startbg|restartbg|pack)
[additional parameters]
```

 Issue "nuxeoctl help" to print this information.
The options between square brackets are optional.
The options and commands separated by "|" are choices ("|" means "exclusive or").

Options

Option	Description
<code>-d</code> <code>--debug</code>	(Since 5.5) Activate Launcher DEBUG logs.
<code>-q</code> <code>--quiet</code>	(Since 5.5) Activate quiet mode.

GUI options

Option	Description
<code>gui</code>	Launcher with a graphical user interface. On Linux/Mac OS X, default is in headless/console mode. On Windows, the <code>gui</code> option is activated by default.
<code>nogui</code>	Windows only. This option deactivates the <code>gui</code> option which is set by default under Windows.

See [the Environment variables page](#) for setting Nuxeo Home and Configuration paths.

Commands

Command	Description
<code>help</code>	Print this message.
<code>start</code>	Start Nuxeo server in background, waiting for effective start. Useful for batch executions requiring the server being immediately available after the script returned.
<code>stop</code>	Stop any Nuxeo server started with the same <code>nuxeo.conf</code> file.
<code>restart</code>	Restart Nuxeo server.

configure	Configure Nuxeo server with parameters from <code>nuxeo.conf</code> .
wizard	Enable the wizard (force the wizard to be played again in case the wizard configuration has already been done).
console	Start Nuxeo server in a console mode. Ctrl+C will stop it.
status	Print server status (running or not).
startbg	Start Nuxeo server in background, without waiting for effective start. Useful for starting Nuxeo as a service.
restartbg	Restart Nuxeo server with a call to <code>startbg</code> after <code>stop</code> .
pack	Build a static archive. Same as the "pack" Shell script.
mp-list	List marketplace packages.
mp-add	Add marketplace package(s) to local cache. You must provide the package file(s) as parameter.
mp-install	Run marketplace package installation. It is automatically called at startup if <code>installAfterRestart.log</code> exists. Else you must provide the package file(s) or ID(s) as parameter.
mp-uninstall	Uninstall marketplace package(s). You must provide the package ID(s) as parameter (see "mp-list" command).
mp-remove	Remove marketplace package(s). You must provide the package ID(s) as parameter (see "mp-list" command).
mp-reset	Reset all packages to DOWNLOADED state. May be useful after a manual server upgrade.

Additional parameters

All parameters following a command are passed to the Java process when executing the command. That can be used for specific installs on which you rely on server specific parameters.

Java usage

Launcher can be run as a Java command, without using the Shell (`nuxeoctl`) or Batch (`nuxeoctl.bat`) script. The equivalent Java command to Shell command is printed at startup. See the line starting with "Launcher command:". It can be reused for writing your own scripts.

Here is the Java usage:

```
java [-Dlauncher.java.opts="JVM options"] [-Dnuxeo.home="/path/to/nuxeo"] [-Dnuxeo.conf="/path/to/nuxeo.conf"] \
[-Djvmcheck=nofail] -jar "path/to/nuxeo-launcher.jar" \
[options] [gui]
(help|start|stop|restart|configure|wizard|console|status|startbg|restartbg|pack)
[additional parameters]
```

Java options

Option	Description
<code>launcher.java.opts</code>	Parameters for the server JVM (default are <code>-Xms512m -Xmx1024m -XX:MaxPermSize=512m</code>).
<code>nuxeo.home</code>	Nuxeo server root path (default is parent of called script).
<code>nuxeo.conf</code>	Path to <code>nuxeo.conf</code> file (default is <code>\$NUXEHOME/bin/nuxeo.conf</code>)

jvmcheck	If equals to "nofail", will continue execution even if JVM does not fit requirements, else will exit.
gui	Launcher with a graphical user interface. On Linux/Mac OS X, default is in headless/console mode. On Windows, the <code>gui</code> option is activated by default.

Commands

See [the `nuxeoctl` commands](#).

Additional parameters

See [the `nuxeoctl` additional parameters](#).

Exit code values

Exit code values are following the [Linux Standard Base Core Specification 4.1](#).

If the status command was requested, NuxeoCtl will return the following exit status codes:

0	program is running or service is OK
3	program is not running
4	program or service status is unknown

In case of an error while processing any action except for status, NuxeoCtl shall print an error message and exit with a non-zero status code:

1	generic or unspecified error
2	invalid or excess argument(s)
3	unimplemented feature
4	user had insufficient privilege
5	program is not installed
6	program is not configured
7	program is not running

Monitoring

Monitoring Nuxeo

Nuxeo server running and components loading statuses.

Since 5.5, Nuxeo provides an URL for monitoring the server status. This method is actually also used by the Launcher to follow the server startup status, after checking the Java process status.

<http://localhost:8080/nuxeo/runningstatus> will be available at last. While it isn't reachable, the server is stopped or still starting.

<http://localhost:8080/nuxeo/runningstatus?info=started> returns `true` if the server finished starting and the Nuxeo runtime is fine with its components.

<http://localhost:8080/nuxeo/runningstatus?info=summary&key=xxx> returns `true` or `false` (see "info=started") and a detailed summary about components. Access to this URL is restricted by an access key configurable in `nuxeo.conf` (see "server.status.key" in [Configuration parameters index \(nuxeo.conf\)](#)).

Sample output if something was wrong at startup (for instance, missing RelationService)

```
false
=====
= Nuxeo EP Started
=====
= Component Loading Status: Pending: 7 / Unstarted: 0 / Total: 462
  * service:org.nuxeo.ecm.webengine.sites.wiki.relation requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
  * service:org.nuxeo.ecm.annotations.graph requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
  * service:org.nuxeo.ecm.platform.relations.jena requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
  * service:org.nuxeo.ecm.annotations.repository.graph requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
  * service:org.nuxeo.ecm.platform.publisher.relations.contrib requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
  * service:org.nuxeo.ecm.platform.relations.services.DefaultJenaGraph requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
  * service:org.nuxeo.ecm.platform.comment.service.CommentServiceConfig requires
[service:org.nuxeo.ecm.platform.relations.services.RelationService]
=====
```

You can get that information with `./bin/nuxeoctl status` (see [nuxeoctl and Control Panel usage](#)).

JVM Garbage collector

The garbage collector attempts to reclaim memory used by objects that are no longer in use by the application.

Monitoring the garbage collector can be very useful when tuning the JVM or setting the initial heap size.

Edit `$NUXEO_HOME/bin/nuxeo.conf` and uncomment the following options titled "Log Garbage Collector informations into a file":

```
JAVA_OPTS=$JAVA_OPTS -Xloggc:${nuxeo.log.dir}/gc.log -verbose:gc -XX:+PrintGCDetails
-XX:+PrintGCTimeStamps
```

JBoss

The JBoss LoggingMonitor service can monitor specific attributes of a MBean periodically and log their value to the filename specified.

More info on the LoggingMonitor:<http://wiki.jboss.org/wiki/Wiki.jsp?page=JBossLoggingMonitor>

Edit `$NUXEO_HOME/bin/nuxeo.conf` and add "monitor" to the `nuxeo.templates` parameter (uncomment it if needed).

The logging-monitor jar file is deployed by default in `$JBOSS_HOME/server/default/lib`.

JBoss JVM information

The "monitor" template will deploy a file named `jvm-monitor-service.xml` which will produce a `jvm.log` file.

JBoss thread pool

The "monitor" template will deploy a file named `webthreads-monitor-service.xml` which will produce a `webthreads.log` file.

Nuxeo unified datasource connection pool

The "monitor" template will deploy a file named `default-ds-monitor-service.xml` which will produce a `nuxeo-ds.log` file.

PostgreSQL

The PostgreSQL logs can be setup like in the pgFouine tutorial:

<http://pgfouine.projects.postgresql.org/tutorial.html>

For instance to log only request slower than 100ms change your postgresql.conf file:

```
log_min_duration_statement = 100
```

After a test you can catch the vacuum output like this:

```
vacuumdb -fzv $DBNAME &> vacuum.log
```

OS

The sysstat utilities are a collection of performance monitoring tools for Linux that is easy to setup.

You can monitor the system activity like this:

```
sar -d -o $JBOSS_HOME/server/default/log/sysstat-sar.log 5 720 >/dev/null 2>&1 &
```

This will monitor the activity every 5s during 1h.

For more information on systat, visit <http://pagesperso-orange.fr/sebastien.godard/>.

Log analysis

logchart

This is a small script that process the following log files:

- Garbage collector logging (`gc.log`)
- Java Virtual Machine logging (`jvm.log`)
- JBoss threads logging (`webthreads.log`)
- NuxeoDS Data source usage (`nuxeo-ds.log`)
- Sysstat sar logging, cpu and disk activity (`sysstat-sar.log`)
- PostgreSQL logs (`pgsql.log`)
- PostgreSQL vacuum output (`vacuum.log`)

to generate an html report with charts:

<http://public.dev.nuxeo.com/~ben/logchart/monitor.html>

More information on logchart can be found on the `README.txt` of the project:

<https://hg.nuxeo.org/tools/logchart/trunk>

Other reporting tools

- GCViewer a tool to visualize data produced by the garbage collector logs: <http://www.tagtraum.com/gcviewer.html>
- kSar a sar grapher that can produce detail PDF report of your system activity: <http://ksar.atomique.net/linux.html>
- pgfouine the PostgreSQL log analyzer which is used by logchart: <http://pgfouine.projects.postgresql.org/>

Nuxeo Management

About Nuxeo Management infrastructure

Nuxeo Runtime Management

The bundle `nuxeo-runtime-management` provides a management infrastructure based on Java Simon.

Two main types of objects are managed:

- Counters,
- StopWatches.

`nuxeo-runtime-management` also provides a way to publish JMX Resources (in addition of the Counters and StopWatches that are automatically published): the `factories` extension point can be used to contribute new JMX resources.

Counters and StopWatches are accessible:

- via JMX,
- via the Admin Center.

Nuxeo Core Management

Nuxeo Core Management uses Nuxeo Runtime Management:

- to expose counters on Events,
- to publish a JMX Bean to manage events and Events handler,
- to publish a JMX Bean to manage "Administrative Status" and Probes.

Administrative Status are a way to define cluster-wide or instance-wide named variable that can be used to manage the status of a running platform:

- turn on/off a node of the cluster,
- display a message to all users of the platform,
- ...

Administrative Status can be configured and declared via an extension point.

Core Management also adds the concept of probes that can be used to run a test on the target deployed platform. Probes can be used to check that all part of the architecture are actually running for real:

- check LDAP access,
- check instance availability,
- check VCS access,
- ...

Probes can be defined via a dedicated extension point.

Probes and Administrative status are accessible:

- via JMX,
- via REST,
- via the Admin Center.

On this page

- About Nuxeo Management infrastructure
 - Nuxeo Runtime Management
 - Nuxeo Core Management
- Existing metrics and counters
 - Event System
 - Repository
 - Transactions management
 - Web Layer
- Probes and Administrative Status
 - Administrative Status
 - Probes
- JMX Access
- Using the Admin Center
 - Activity
 - Monitoring
- Rest Access
 - Counters

Existing metrics and counters

By default, Nuxeo Platform comes with a set of counters, StopWatches and probes.

Event System

All operations executed against the repository will trigger events. For this reason, monitoring events is a good way to have an idea about the activity of the platform.

There are by default three defined counters:

- `event.create`: counts the number of created Document/Version,
- `event.update`: counts the number of updated Documents,
- `event.remove`: counts the number of removed Documents/Version.

A `EventMonitoring` service is also exposed via JMX (and as a Java Nuxeo Service). This `EventMonitoring` service allows to:

- enable/disable all synchronous event listeners,
- enable/disable all asynchronous event listeners,
- enable/disable tracking of listeners execution,
- enable/disable bulk mode,
- retrieve statistics about event handler execution.

This `EventMonitoring` service can be used:

- for tuning the platform during mass import,
- for profiling execution of listeners,
- for monitoring.

Repository

Counters and Metrics

The repository exposes by default three counters about cache usage:

- `cache.access`,
- `cache.hits`,
- `cache.size`.

These counters are updated every 200 access to avoid having bad performance impacts.

If `org.nuxeo.vcs.cache.statistics` is set to true (via the [System information](#) tab of the Admin Center or via `nuxeo.conf`), there are also two StopWatch that are defined:

- `cache.gets`: time to get an entry from cache,
- `sor.gets`: time to bulk get entries from DB.

Services

`SQLStorage` service provides a JMX API to:

- gather information about the repository (active sessions, ...),
- do administration operations (Flush caches, Start Binaries GC ...).

Transactions management

To have more information about the repository and transaction management, you can also deploy the additional bundle `nuxeo-core-management-jta:jca`.

This bundle provides:

- StopWatch for transactions,
- TransactionMonitoring Service,
- JCA Connectionpool Service.

Web Layer

Two counters are published by default in JMX:

- `web.requests`: number of HTTP requests served,
- `we.sessions`: number of HTTP Sessions.

Probes and Administrative Status

Nuxeo comes with a set of default probes and administrative status. Both can be seen and managed from JMX and from Admin Center.

Administrative Status

By default only three status are defined :

- nuxeoInstance: indicates if a Nuxeo instance (cluster node) is active or not,
- adminMessage: message to be displayed to all users,
- smtpService: defines if SMTP gateway can be used.

Probes

By default four probes are defined:

- adminStatus: checks local instance enable flag (checks nuxeoInstance administrative status),
- activeRepositorySession: returns the number of active sessions per repository,
- ldapDirectory: check LDAP connectivity,
- remoteSQLStorageSession: number of remove VCS client connected (only used in VCS client/server mode that is not enabled by default).

JMX Access

You can use JVisualVM or similar tool to access Nuxeo JMX interface.

Name	Value
BlockingTimeoutMilliseconds	100
ConnectionCount	5
IdleConnectionCount	5
IdleTimeoutMinutes	10
PartitionCount	1
PartitionMaxSize	20
PartitionMinSize	0

JMX Remote access is by default disabled. You can activate it by adding the required option (for example in `nuxeo.conf`)

```
-Dcom.sun.management.jmxremote
```

However, you will then have to manage security for this access, since there is no authentication by default.

Using the Admin Center

Inside the Admin Center there are two sections that are related to monitoring: Activity and Monitoring

Activity

The Activity section provides access to:

- a view that displays HTTP counters (requests and sessions),
- a view on audit logs,
- activity charts based on web and repository counters.

Monitoring

The Monitoring sections provides access to:

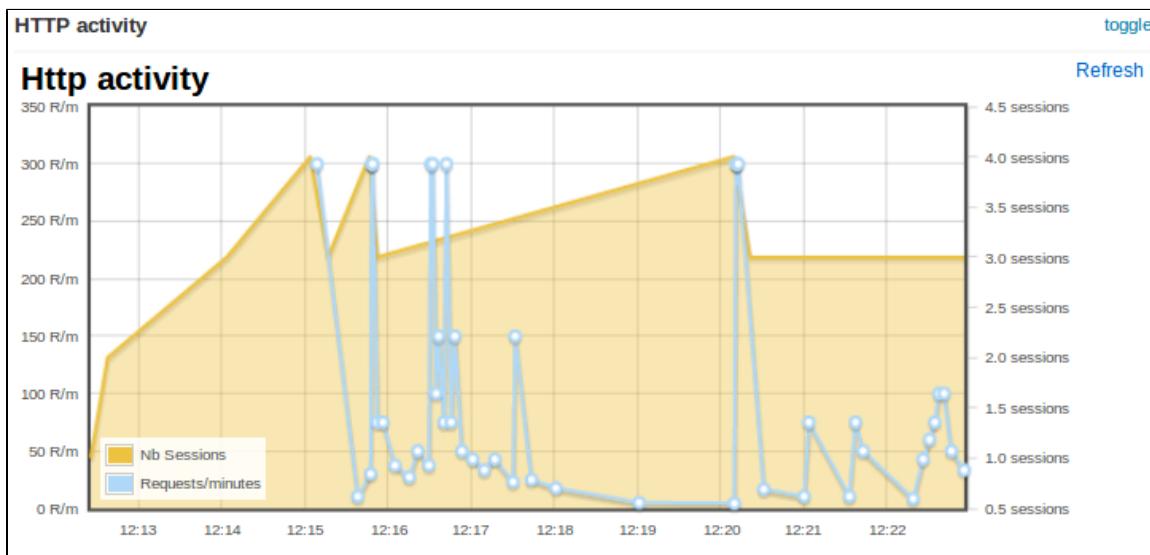
- a view on Administrative Status (view / edit),
- a view on probes (view/run),
- a view that allows to enable Event Listener statistic gathering.

Rest Access

Counters

Counter are exposed via Automation API [Counters.GET](#)

This API is used inside the Admin Center to be able to generate the small graphs with an OpenSocial gadget.



Sample CURL call:

```
curl -H 'Content-Type:application/json+nxrequest' -X POST -d
'{"params": {"counterNames": "org.nuxeo.web.requests"} }' -u
Administrator:Administrator
http://localhost:8080/nuxeo/site/automation/Counters.GET
```

Remote monitoring through a SSH tunnel

At some time, you may need to monitor your server trough your SSH access. We assume you have a direct connection to your remote server host. If you're using a gateway, this works too, you just have to configure the right ports forwarding.

Here is how to monitor your server through a SSH tunnel:

1. On server host, run jstated with the privileges of the nuxeo's user:

```
jstated -J-Djava.security.policy=jstat.permissions
```

2. On your SSH connection, configure a local “dynamic” application-level port forwarding:

```
ssh -D 6767 remote-server-host
```

3. Run jvisualvm on your local host and in the network options, enable the manual proxy settings and configure a socks proxy:

```
localhost:6767
```

4. Now in jvisualvm, add a remote host for the remote-server-host.

You should get the list of Java processes ran by Nuxeo’s user remotely.

5. Identify Nuxeo’s Tomcat and launch a connection.

In the given jstat command line, we reference the Java security configuration file jstat.permissions. Here is its content:

```
grant codebase "file:${java.home}/../lib/tools.jar" {
    permission java.security.AllPermission;
};
```

Transactions and connections

Troubleshooting issues with connections with transactional resources (databases) can be done with the optional bundle nuxeo-core-management-jtajca. Just download it from our Nexus it from our nexus and put it in NUXEO_HOME/nxserver/plugins. Then restart the Nuxeo server.

Logging transactional events

Configure log4j in NUXEO_HOME/lib/log4j.xml by adding the following keywords to your appender conversion pattern %t for the thread name and %x for the logging context map:

```
<param name="ConversionPattern" value="%d{ISO8601} %t %-5p [%c] %m %X%n" />
```

You should also add a new category if you want the traces being enabled:

```
<category name="org.nuxeo.ecm.core.management.jtajca">
  <priority value="TRACE" />
</category>
```

At this stage, once a transaction is started or a connection is opened, their identifiers are put in a context map for the logger. By adding the %x key word, you’ve requested to print them each a message is logged. The transactions and connections will also be logged. You should add additional log statements at level debug or trace around the code you want to monitor.

Monitoring transactional resources

You should enable JMX for being able to poll the mbean attributes. In NUXEO_HOME/bin/nuxeo.conf uncomment the JMX options.

! You should note that these settings open a **security hole** on the server and should not be left as this in production.

```
# Enable JMX
JAVA_OPTS=$JAVA_OPTS -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=1089 -Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.authenticate=false
```

You've got two new beans in JMX that will provide you the way of monitoring the transactions and connections being used in your system. Just point your JMX browser (such as JVisualVM which is part of the JDK) to the server and looks to the following beans :

- org.nuxeo:name=ConnectionMonitoring,type=service,repositoryName=default
- org.nuxeo:name=TransactionMonitoring,type=service

The transaction monitoring provides useful information about the transactions in progress in your application. Also it gives you global counters and information about the last committed and rolledback transaction in the system. The connection monitoring provide you a way of configuring the pool connection size used by the Nuxeo storage. It gives you also access to global counters about the connection usage.

Nuxeo Shell (admin's best friend)



Availability Info

The Nuxeo Shell can work against any Nuxeo distributions since the **5.4.0.1** version of Nuxeo EP and deprecates the old RMI-based Shell.

Nuxeo Shell is working on top of Nuxeo Content Automation Client and is using the REST API provided by the Automation framework. Because it is using HTTP to communicate with the server the Shell can be used against any Nuxeo distribution which includes the automation framework.

This also means, most of the Shell commands are implemented as remote Automation Operations.



Downloads

The Nuxeo Shell is available at the Admin Center, in the "Monitor" tab and also as a WebEngine site at <http://host:port/nuxeo/site/shell>.

Nuxeo Shell can be installed in Eclipse adding the following repository: Nuxeo ECR - <http://osgi.nuxeo.org/p2/ecr/ide/>, then install "ECR Shell Feature".

Download the last released version of Nuxeo Shell for a manual install. You can browse available versions at <https://maven.nuxeo.org> (or pick latest).

Overview

Nuxeo Shell comes with a lot of features to improve your experience using the command line. Also, it provides high pluggability so you can extend it and add new commands in a simple way.

Here is a list of the most important features.

Interactive Execution

The Nuxeo Shell uses `jline` library for interactive mode. When you launch the Shell you will automatically enter the interactive mode. This mode provides a prompt and auto completion so you can easily browse a Nuxeo repository and execute commands.

Batch Execution

You can also launch the Shell in batch mode. In this mode you can specify a list of commands to be executed and then to return the control back to the terminal.

There are 3 batch modes available:

1. Run commands from a file
2. Run commands from standard input.
3. Run commands specified on the command line - this is a convenient way to run a short list of commands.

See [the Shell Batch Mode page](#) for more details.

Namespaces

Namespaces are commands registry that you can register in the Shell. A command registry is useful to bring inside the Shell a new set of commands for a specific context without having name clash problems with already existing commands.

To switch to a different namespace use the `use` command.

So for example, you may want to have `ls` command that is listing the content of a Nuxeo Folder when connected to a remote server but still use the `ls` command when switching to the `local` context to be able to list the content of a directory on the local file system.

Also, namespaces are hierarchical so a namespace my extend another one to adds more commands. Available namespaces are setup by the

features installed in the Shell. By default, Nuxeo Shell provides the following default namespaces:

1. global - the global namespace. This provides global commands like `help`, `'use'`, `'cmds'` etc.
2. local - provides file system commands like: `ls`, `cd`, `cat`, etc. Extends the global namespace.
3. remote - provides remote commands against a Nuxeo server. Extends the global namespace. Available only after connecting to a remote server.
4. automation - expose remote automation operations as commands. Available only after connecting to a remote server.

Auto Completion

Auto completion support is provided by `jline`. The Shell is leveraging this so that you can auto complete:

1. command names.
2. parameter names.
3. parameter values (when the command provide completion).

There are several type of objects that supports completion (like, file, document, etc.) but you can add more by implementing new completors.

Scripting



Security Warning

Because of potential security issues the scripting feature is available only when logged in as Administrator

Nuxeo Shell is providing scripting capabilities through Groovy or Mvel. You can execute thus your Groovy scripts on the server JVM. There are two ways of executing scripts:

1. Either invoke a script file from your file system using the `script` command.
2. Either write a command that execute a script (instead of executing a remote automation operation).

Documentation Generation

Command documentation pages are automatically generated by the Shell by introspecting the command annotations.

Also, if you want to add more details like in-depth explanation of the command, examples etc. then you can write this in a text file and the Shell will automatically append it to the generated documentation.

When writing custom documentation you can use tags like `{header}`, `{bold}`, `{red}`, `{green}` etc. to control the formatting on the terminal.

Automation Chain Execution

You can execute any automation chain exposed by the server by using the `run` or `runonfile` commands. This is useful since you can create your administration commands in [Nuxeo Studio](#) by creating Automation Chains.

Direct Operation execution

The automation mode provides Shell commands for any operation available on the server. These commands are automatically generated from operation descriptors.

This mode should be used only to debug operations or to execute some operations not exposed through specialized commands by the Shell.



Experimental

Because of the Shell environment not every automation operation command will correctly work. You should use regular operations instead.

Exception Handling

In the interactive mode unchecked exceptions are printed in red on the screen. Checked exceptions are not printed - only the exception message is printed out. If you need to see the stack trace of the last exception type `trace`.

Extensibility

A Shell feature is providing Shell configuration such as namespaces, commands, completors, etc. To extend the capabilities of the Shell you can register a new feature which will install the new capabilities into the Shell at boot time.

The Shell can be extended wither by modifying the Nuxeo Shell JAR either by adding on the classpath JARs containing additional Nuxeo Shell features.

See [Extending the Shell](#) section for more details.

Usage

Launching the Shell

You can launch the Shell either by running the command:

```
java -jar nuxeo-shell.jar
```

either by running:

```
java -cp nuxeo-shell.jar org.nuxeo.shell.Main
```

The difference is that the first command will launch the Shell in a Swing based terminal (i.e. a desktop application) while the second one will launch the Shell inside your current terminal.

On some Operating Systems like Windows double clicking the nuxeo-shell.jar will launch the Shell in a Swing based terminal.

When launching the Shell the **local** namespace is automatically activated.

When connecting to a remote server the Shell will automatically switch to **remote** namespace. When disconnecting it will go back to the **local** namespace.

To switch between namespaces just use the **use** command. For example if you are in the **remote** namespace you can switch to the **local** one by executing the command:

```
use local
```

Getting the Shell Version

You can launch the Shell using

```
java -jar nuxeo-shell.jar --version
```

to get information about the Shell version and the minimal server version required by the Shell to correctly run against a remote server.

You can also have this information by using the **version** command.

Connecting to a Server

To connect to a remote Nuxeo Server you must specify the Automation Service URL.

This URL is in the form:

```
http://host:port/nuxeo/site/automation
```

For example, in the case of a local server your URL will be:

```
http://localhost:8080/nuxeo/site/automation
```

To connect to a server you should use the global **connect** command. This command require 3 parameters:

1. the URL of the remote automation service.
2. the username to login
3. the password to login

You can either pass these arguments to the connect command itself or through the command line when starting the Shell.

You can use `-u` for the username, `-p` for the password and the URL should be given as argument.

If password is not specified you will be prompted for a password when in interactive mode.

Example

Here is an example of a short session:

After launching the Shell you are in **local** mode. So you can use the provided file system commands like:

- `ls` - to list the content of the current directory.
- `cd, pushd, popd` - to change the current directory.
- `cat, mv, cp, etc.` for other file based operations.

To connect to a remote server type:

```
connect http://localhost:8080/nuxeo/site/automation -u Administrator
```

After the connection is done you are automatically switched in **remote** namespace.

So doing now a `ls` will list the content of the current document. (which for now the root document).

To switch back in **local** namespace type:

```
use local
```

To show the current namespace name type:

```
use
```

**Note**

When doing file based auto-completion this will be relative to the current directory (that you can change using `cd, pushd, popd` when in local namespace). The same for document based auto-completion.

Shell Command Index

Here is a list of namespaces available in the Nuxeo Shell, each namespace providing an index of its commands.



The command index is generated using the shell command

```
help -export [file] -ns [namespace]
```

Built-in Commands

**Namespace: *global***

Basic commands provided by the shell

Index:

- `commands`
- `exit`
- `help`
- `install`
- `interactive`
- `settings`
- `trace`

- use
- version

commands

NAME

commands – Print a list of available commands

SYNTAX

```
commands
```

ALIASES

cmds

exit

NAME

exit – Exit the interactive shell

SYNTAX

```
exit [code]
```

ALIASES

quit

ARGUMENTS

- code - [optional] -The exit code. Must be a positive number otherwise 0 is assumed. Defaults to 0.

help

NAME

help – The help command

SYNTAX

```
help [options] [command]
```

OPTIONS

- -export - If used export all the commands available in a wiki format to the given file. If a directory is given the export will be made in file help.wiki in that directory.
- -ns - [optional] - to be used to filter commands by namespaces when generating the documentation. By default all namespaces are dumped.

ARGUMENTS

- command - [optional] -the name of the command to get help for

install

NAME

install – Install a SH script to launch the shell in the terminal. Available only for UNIX systems.

SYNTAX

```
install [file]
```

ARGUMENTS

- file - [optional] -the file where to install the shell script. If not used the script will be printed on stdout.

interactive

NAME

interactive – Interactive shell

SYNTAX

```
interactive
```

settings

NAME

settings – Print or modify the shell settings.

SYNTAX

```
settings [options] [name] [value]
```

OPTIONS

- -reset - [flag] - Reset settings to their defaults. Need to restart shell.

ARGUMENTS

- name - [optional] -The variable to print or set.
- value - [optional] -The variable value to set.

trace

NAME

trace – Print the last error stack trace if any

SYNTAX

```
trace
```

use

NAME

use – Switch the current command namespace. If no namespace is specified the current namespace name is printed.

SYNTAX

```
use [name]
```

ARGUMENTS

- name - [optional] -The command namespace to use

version

NAME

version – Print Nuxeo Shell Version

SYNTAX

```
version
```

Related pages



[Built-in Commands \(Nuxeo Installation and Administration - 5.5\)](#)



[File System Commands \(Nuxeo Installation and Administration - 5.5\)](#)



[Nuxeo Server Commands \(Nuxeo Installation and Administration - 5.5\)](#)



[Shell Batch Mode \(Nuxeo Installation and Administration - 5.5\)](#)



[Shell Commands \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Shell Features \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Extending The Shell \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Shell Documentation \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Shell Namespaces \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Configuration Commands \(Nuxeo Installation and Administration - 5.5\)](#)

• File System Commands

Namespace: *local*

Commands available on the local file system

Index:

- cat
- cd
- cp
- ls
- mkdir
- mv
- popd
- pushd
- pwd
- rm
- touch

cat

NAME

cat – Print the content of a file

SYNTAX

```
cat [file]
```

ARGUMENTS

- file - [optional] -The file to print

cd

NAME

cd – Change the local working directory

SYNTAX

```
cd file
```

ARGUMENTS

- file - [required] - A local directory to change to

cp

NAME

cp – Copy a file or directory

SYNTAX

```
cp [options] source destination
```

OPTIONS

- -r - [flag] - Recursive copy directories

ARGUMENTS

- source - [required] - The file to copy
- destination - [required] - The target file

ls

NAME

ls – List file names in a local directory

SYNTAX

```
ls [file]
```

ARGUMENTS

- file - [optional] -A local directory to list its content. Defaults to the working directory

mkdir

NAME

mkdir – Create a directory in local file system

SYNTAX

```
mkdir file
```

ARGUMENTS

- file - [required] - The directory path to create

mv

NAME

mv – Rename a file or directory

SYNTAX

```
mv source destination
```

ARGUMENTS

- source - [required] - The file to rename
- destination - [required] - The target file

popd

NAME

popd – Pop working directory

SYNTAX

```
popd
```

pushd

NAME

pushd – Push a new local working directory

SYNTAX

```
pushd file
```

ARGUMENTS

- file - [required] - A local directory to push

pwd**NAME**

pwd – Print the local working directory

SYNTAX

```
pwd [options]
```

OPTIONS

- -s - [flag] - Use this flag to show the working directory stack

rm**NAME**

rm – Remove a file or directory

SYNTAX

```
rm [options] file
```

OPTIONS

- -r - [flag] - Recursive remove directories

ARGUMENTS

- file - [required] - The directory path to create

touch**NAME**

touch – Touch a file

SYNTAX

```
touch file
```

ARGUMENTS

- file - [required] - The file to touch

Related pages[Built-in Commands \(Nuxeo Installation and Administration - 5.5\)](#)**• Nuxeo Server Commands****Namespace: *remote***

High level commands exposed by a remote Nuxeo Server

[File System Commands \(Nuxeo Installation and Administration - 5.5\)](#)**Index:**

- audit
- cat
- cd
- connect
- cp
- disconnect
- fire
- getFile

[Nuxeo Server Commands \(Nuxeo Installation and Administration - 5.5\)](#)[Shell Batch Mode \(Nuxeo Installation and Administration - 5.5\)](#)[Shell Commands \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)[Shell Features \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)

- Extending The Shell (Nuxeo Enterprise Platform (EP) - 5.5)
- Shell Documentation (Nuxeo Enterprise Platform (EP) - 5.5)
- Shell Namespaces (Nuxeo Enterprise Platform (EP) - 5.5)
- Configuration Commands (Nuxeo Installation and Administration - 5.5)
 - popd
 - publish
 - pushd
 - putfile
 - pwd
 - query
 - rename
 - rm
 - rmfile
 - run
 - runonfile
 - script
 - setp
 - tree
 - unlock
 - update
- getfiles
- getp
- getrel
- lstate
- lock
- ls
- mkdir
- mkrel
- mv
- perms

audit

NAME

audit – Run a query against audit service

SYNTAX

```
audit [options] query
```

OPTIONS

- -s - Use this to change the separator used in query variables. The default is ','
- -ctx - Use this to set query variables. Syntax is: "k1=v1,k1=v2"
- -max - The max number of rows to return.
- -page - The current page to return. To be used in conjunction with -max.
- -f - Use this to save results in a file. Otherwise results are printed on the screen.

ARGUMENTS

- query - [required] - The query to run. Query is in JPQL format

USAGE

The -page parameter can be used in conjunction with -max parameter to paginate the query result.
To specify the first page use 1 as value, for the second page use 2 and so on.

When saving results in a file - they are in JSON format - and dates are specified using a long value timestamp.
Results printed on the screen are printed in tab separated columns:
eventId category eventDate principal docUUID docType docLifeCycle comment

EXAMPLES

Using date literals in your query:

```
audit "FROM LogEntry log WHERE log.eventDate > timestamp('2010-11-10 00:00:00')"
```

Using pagination:

```
audit "FROM LogEntry log WHERE log.category='studio' ORDER BY log.eventDate DESC" -max 20 -page 1
```

Using query variables:

```
audit "FROM LogEntry log WHERE log.category='studio' AND log.eventDate > :startDate" -ctx "startDate={d 2010-11-10}"
```

or

```
audit "FROM LogEntry log WHERE log.category='studio' AND log.eventDate > :startDate" -ctx "startDate={d 2010-11-10 00:00:00}"
```

Note that query variable keys must be prefixed with "audit.query." to avoid name clash with other keys in the context.

cat

NAME

cat – Print document details

SYNTAX

```
cat [options] [doc]
```

OPTIONS

- -all - [flag] - Include all schemas. The -schemas attribute will be ignored if used in conjunction with this one.
- -schemas - A filter of schemas to include in the document. Use * for all schemas.

ARGUMENTS

- doc - [optional] -The document to print. To use UIDs as references you should prefix them with 'doc:'.

cd

NAME

cd – Change the context document

SYNTAX

```
cd doc
```

ARGUMENTS

- doc - [required] - A reference to the new context document to use. To use UID references prefix them with 'doc:'.

connect

NAME

connect – Connect to a remote automation server

SYNTAX

```
connect [options] [url]
```

OPTIONS

- -p - The password
- -u - The username

ARGUMENTS

- url - [optional] -The URL of the automation server

cp

NAME

cp – Copy a document

SYNTAX

```
cp [options] src dst
```

OPTIONS

- -name - A new name for the copy. I not specified preserve the source name

ARGUMENTS

- src - [required] - The document to copy. To use UID references prefix them with 'doc:'.
- dst - [required] - The target parent. To use UID references prefix them with 'doc:'.

disconnect

NAME

disconnect – Close current connection to server. If not connected nothing is done.

SYNTAX

```
disconnect
```

fire

NAME

fire – Fire a core event in the context of the given document

SYNTAX

```
fire event [doc]
```

ARGUMENTS

- event - [required] - The event to fire. This parameter is required.
- doc - [optional] -The context document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

getfile

NAME

getfile – Get a document attached file

SYNTAX

```
getfile [options] [doc]
```

OPTIONS

- -todir - An optional target directory to save the file. The default is the current working directory.
- -xpath - The XPath of the blob property to get. Defaults to the one used by the File document type.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

getfiles

NAME

getfiles – Get all the files attached to a document

SYNTAX

```
getfiles [options] [doc]
```

OPTIONS

- -todir - An optional target directory to save the file. The default is the current working directory.
- -xpath - The XPath of the blob property to get. Defaults to the one used by the File document type.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document content is used. To use UID references prefix them with 'doc:'.

getp

NAME

getp – Get the value of a document property

SYNTAX

```
getp [options] [doc]
```

OPTIONS

- -xpath - The XPath of the property to get. This parameter is required.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

getrel

NAME

getrel – Get relations between two documents

SYNTAX

```
getrel [options] [doc]
```

OPTIONS

- -in - [flag] - Is the document the relation object?
- -predicate - The relation predicate - requested.
- -out - [flag] - Is the document the relation subject? This flag is by default on true.

ARGUMENTS

- doc - [optional] -The document involved in the relation

lcstate

NAME

lcstate – Set or view the current life cycle state of a document

SYNTAX

```
lcstate [options] [doc]
```

OPTIONS

- -set - If specified The new life cycle state. If not specified then in write mode the local ACL is used and in view mode all ACLs are printed.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

lock

NAME

lock – Lock a document

SYNTAX

```
lock [options] [doc]
```

OPTIONS

- -key - An optional lock key. If not specified the default one is used.

ARGUMENTS

- doc - [optional] -The document to lock. If not specified the current document is used. To use UID references prefix them with 'doc:'.

ls

NAME

ls – List children documents

SYNTAX

```
ls [options] [doc]
```

OPTIONS

- -uid - [flag] - If used the documents will be printed using the document UID.

ARGUMENTS

- doc - [optional] -A document to list its content. If not specified list the current document content. To use UID references prefix them with 'doc:'.

mkdir

NAME

mkdir – Create a document of the given type

SYNTAX

```
mkdir [options] type path
```

OPTIONS

- -title - An optional document title.

ARGUMENTS

- type - [required] - The document type
- path - [required] - The document path

mkrel

NAME

mkrel – Create a relation between two documents

SYNTAX

```
mkrel [options] subject object
```

OPTIONS

- -predicate - The relation predicate - requested.

ARGUMENTS

- subject - [required] - The subject of the relation
- object - [required] - The object of the relation

mv

NAME

mv – Move a document

SYNTAX

```
mv [options] src dst
```

OPTIONS

- -name - A new name for the document. If not specified preserve the source name

ARGUMENTS

- src - [required] - The document to move. To use UID references prefix them with 'doc:'.
- dst - [required] - The target parent. To use UID references prefix them with 'doc:'.

perms

NAME

perms – Set or view permissions on a document

SYNTAX

```
perms [options] [doc]
```

OPTIONS

- -grant - If used the ACL will be modified by granting the specified permission on the specified user. The grant value format is "user:permission".
- -deny - If used the ACL will be modified by denying the specified permission on the specified user. The deny value format is "user:permission".
- -remove - [flag] - Remove the given ACL.
- -acl - The ACL to view or modify. If not specified then in write mode the local ACL is used and in view mode all ACLs are printed.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

popd

NAME

popd – Change the context document and pop the document from the navigation stack.

SYNTAX

```
popd
```

publish

NAME

publish – Publish a document into a section

SYNTAX

```
publish [options] src section
```

OPTIONS

- -override - If set to false will not override an existing published document with same name. The default is "true".

ARGUMENTS

- src - [required] - The document to copy. To use UID references prefix them with 'doc:'.
- section - [required] - The target parent. To use UID references prefix them with 'doc:'.

pushd

NAME

pushd – Change the context document and push the document on the navigation stack.

SYNTAX

```
pushd doc
```

ARGUMENTS

- doc - [required] - A reference to the new context document to use. To use UID references prefix them with 'doc:'.

putfile

NAME

putfile – Attach a file to a document

SYNTAX

```
putfile [options] file [doc]
```

OPTIONS

- -xpath - The XPath of the blob property to set. Defaults to the one used by the File document type.

ARGUMENTS

- file - [required] - The file to upload
- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

pwd

NAME

pwd – Print the current context document

SYNTAX

```
pwd [options]
```

OPTIONS

- -s - [flag] - Use this flag to show the context documents stack

query

NAME

query – Query documents

SYNTAX

```
query [options] [query]
```

OPTIONS

- -uid - [flag] - If used the matching documents will be printed using the document UID.

ARGUMENTS

- query - [optional] -The document path

EXAMPLES

```
query "SELECT * FROM Document WHERE ecm:primaryType='Folder'"
```

```
query -uid "SELECT * FROM Document WHERE ecm:primaryType=\"Folder\""
```

rename

NAME

rename – Rename a document

SYNTAX

```
rename [options] [doc]
```

OPTIONS

- -name - A new name for the document. This parameter is required.

ARGUMENTS

- doc - [optional] -The document to rename. To use UID references prefix them with 'doc:'.

rm

NAME

rm – Remove a document

SYNTAX

```
rm [path]
```

ARGUMENTS

- path - [optional] -The document path. To use UID references prefix them with 'doc:'.

rmfile

NAME

rmfile – Remove an attached file from a document

SYNTAX

```
rmfile [options] [doc]
```

OPTIONS

- -xpath - The XPath of the blob property to remove. Defaults to the one used by the File document type.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

run**NAME**

run – Run a server automation chain that accepts a document or void input

SYNTAX

```
run [options] chain [doc]
```

OPTIONS

- -s - Use this to change the separator used in context variables. The default is ','
- -ctx - Use this to set execution context variables. Syntax is: k1=v1,k1=v2
- -void - [flag] - Use this to avoid having the server sending back the result.

ARGUMENTS

- chain - [required] - The chain to run
- doc - [optional] -A reference to the new context document to use. To use UID references prefix them with 'doc:'.

runonfile**NAME**

runonfile – Run a server automation chain that accepts a file as an input

SYNTAX

```
runonfile [options] chain file
```

OPTIONS

- -s - Use this to change the separator used in context variables. The default is ','
- -ctx - Use this to set execution context variables. Syntax is: k1=v1,k1=v2
- -void - [flag] - Use this to avoid having the server sending back the result.

ARGUMENTS

- chain - [required] - The chain to run
- file - [required] - A reference to the new context document to use. To use UID references prefix them with 'doc:'.

script**NAME**

script – Run a script on the server

SYNTAX

```
script [options] file
```

OPTIONS

- -s - Use this to change the separator used in context variables. The default is ','
- -ctx - Use this to set execution context variables. Syntax is: "k1=v1,k1=v2"

ARGUMENTS

- file - [required] - The script file. Must have a .mvel or .groovy extension

setup

NAME

setup – Set a property on a document

SYNTAX

```
setup [options] [doc]
```

OPTIONS

- -value - The property value. If not specified the current property value is removed.
- -xpath - The XPath of the property to set. This parameter is required.

ARGUMENTS

- doc - [optional] -The target document. If not specified the current document is used. To use UID references prefix them with 'doc:'.

tree

NAME

tree – List a subtree

SYNTAX

```
tree [doc]
```

ARGUMENTS

- doc - [optional] -A document to list its subtree. If not specified list the current document subtree. To use UID references prefix them with 'doc:'.

unlock

NAME

unlock – Unlock a document

SYNTAX

```
unlock [doc]
```

ARGUMENTS

- doc - [optional] -The document to unlock. If not specified the current document is used. To use UID references prefix them with 'doc:'.

update

NAME

update – Update document properties

SYNTAX

```
update [options] [properties] [path]
```

OPTIONS

- -s - Use this to change the separator used in properties. The default is ','

ARGUMENTS

- properties - [optional] -The properties to update.
- path - [optional] -The document path

Related pages

[Built-in Commands \(Nuxeo Installation and Administration - 5.5\)](#)

Nuxeo Automation Commands

Namespace: *automation*
Commands exposed by the Nuxeo Server through automation

Index:

- Audit.Log
- Audit.Query
- Auth.LoginAs
- Auth.Logout
- Blob.Attach
- Blob.Create
- Blob.CreateZip
- Blob.Get
- Blob.GetList
- Blob.Pop
- Blob.PopList
- Blob.Post
- Blob.Pull
- Blob.PullList
- Blob.Push
- Blob.PushList
- Blob.Remove

File System Commands (Nuxeo Installation and Administration - 5.5)

- [Nuxeo Server Commands \(Nuxeo Installation and Administration - 5.5\)](#)
- [Shell Batch Mode \(Nuxeo Installation and Administration - 5.5\)](#)
- [Shell Commands \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)
- [Shell Features \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)
- [Extending The Shell \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)
- [Shell Documentation \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)
- [Shell Namespaces \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)
- [Configuration Commands \(Nuxeo Installation and Administration - 5.5\)](#)
 - Blob.Set
 - Blob.SetFilename
 - BlobToFile
 - BlobToPDF
 - Context.FetchDocument
 - Context.RestoreBlobInput
 - Context.RestoreBlobsInput
 - Context.RestoreDocumentsInput
 - Context.RestoreDocumentsInput
 - Context.RunDocumentOperation
 - Context.RunInputScript
 - Context.RunOperation
 - Context.RunScript
 - Context.SetInputAsVar
 - Context.SetVar
 - Document.CheckIn
 - Document.CheckOut
 - Document.Copy
 - Document.Create
 - Document.CreateVersion
 - Document.Delete
 - Document.Fetch
 - Document.FetchByProperty
 - Document.Filter
 - Document.GetChild
 - Document.GetChildren
 - Document.GetParent
 - Document.GetPrincipalEmails
 - Document.GetUsersAndGroups
 - Document.Lock
 - Document.Move
 - Document.MultiPublish
 - Document.Pop
 - Document.PopList
 - Document.Publish
 - Document.Pull

- Document.PullList
- Document.Push
- Document.PushList
- Document.Query
- Document.Reload
- Document.RemoveACL
- Document.RemoveProperty
- Document.Save
- Document.SaveSession
- Document.SetACE
- Document.SetLifeCycle
- Document SetProperty
- Document.Unlock
- Document.Update
- Notification.SendEvent
- Notification.SendMail
- Relations.CreateRelation
- Relations.GetRelations
- Workflow.CreateTask
- Workflow.GetTask
- print

Audit.Log

NAME

Audit.Log – Log events into audit for each of the input document. The operation accept as input one ore more documents that are returned back as the output.

SYNTAX

```
Audit.Log [options] [the input document(s)]
```

OPTIONS

- -event -
- -ctx - Can be used to inject context properties in Java properties format
- -category -
- -void - [flag] - If void the server will not return the result back
- -comment -

ARGUMENTS

- the input document(s) - [optional] -null

Audit.Query

NAME

Audit.Query – Execute a JPA query against the Audit Service. This is returning a blob with the query result. The result is a serialized JSON array. You can use the context to set query variables but you must prefix using 'audit.query.' the context variable keys that match the ones in the query.

SYNTAX

```
Audit.Query [options]
```

OPTIONS

- -maxResults -
- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -pageNo -
- -query -

Auth.LoginAs

NAME

Auth.LoginAs – Login As the given user. If no user is given a system login is performed. This is a void operations - the input will be returned back as the output.

SYNTAX

```
Auth.LoginAs [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Auth.Logout

NAME

Auth.Logout – Perform a logout. This should be used only after using the Login As operation to restore original login. This is a void operations - the input will be returned back as the output.

SYNTAX

```
Auth.Logout [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Blob.Attach

NAME

Blob.Attach – Attach the input file to the document given as a parameter. If the XPath points to a blob list then the blob is appended to the list, otherwise the XPath should point to a blob property. If the save parameter is set the document modification will be automatically saved. Return the blob.

SYNTAX

```
Blob.Attach [options] the input file(s)
```

OPTIONS

- -save -
- -ctx - Can be used to inject context properties in Java properties format
- -document -
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input file(s) - [required] - null

Blob.Create

NAME

Blob.Create – Creates a file from a given URL. The file parameter specifies how to retrieve the file content. It should be an URL to the file you want to use as the source. You can also use an expression to get an URL from the context. Returns the created file.

SYNTAX

```
Blob.Create [options]
```

OPTIONS

- -mime-type -
- -file -
- -ctx - Can be used to inject context properties in Java properties format

- -filename -
- -void - [flag] - If void the server will not return the result back
- -encoding -

Blob.CreateZip

NAME

Blob.CreateZip – Creates a zip file from the input file(s). If no file name is given, the first file name in the input will be used. Returns the zip file.

SYNTAX

```
Blob.CreateZip [options] the input file(s)
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -filename -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input file(s) - [required] - null

Blob.Get

NAME

Blob.Get – Gets a file attached to the input document. The file location is specified using an XPath to the blob property of the document. Returns the file.

SYNTAX

```
Blob.Get [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

Blob.GetList

NAME

Blob.GetList – Gets a list of files that are attached on the input document. The files location should be specified using the blob list property XPath. Returns a list of files.

SYNTAX

```
Blob.GetList [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

Blob.Pop

NAME

Blob.Pop – Restore the last saved input file in the context input stack. This operation must be used only if a PUSH operation was previously made. Return the last <i>pushed</i> file.

SYNTAX

```
Blob.Pop [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Blob.PopList

NAME

Blob.PopList – Restore the last saved input file list in the context input stack

SYNTAX

```
Blob.PopList [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Blob.Post

NAME

Blob.Post – Post the input file to a target HTTP URL. Returns back the input file.

SYNTAX

```
Blob.Post [options] the input file(s)
```

OPTIONS

- -url -
- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input file(s) - [required] - null

Blob.Pull

NAME

Blob.Pull – Restore the last saved input file in the context input stack. This operation must be used only if a PUSH operation was previously made. Return the first <i>pushed</i> file.

SYNTAX

```
Blob.Pull [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Blob.PullList

NAME

Blob.PullList – Restore the first saved input file list in the context input stack

SYNTAX

```
Blob.PullList [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Blob.Push

NAME

Blob.Push – Push the input file on the context stack. The file can be restored later as the input using the corresponding pop operation. Returns the input file.

SYNTAX

```
Blob.Push [options] the input file(s)
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input file(s) - [required] - null

Blob.PushList

NAME

Blob.PushList – Push the input file list on the context stack. The file list can be restored later as the input using the corresponding pop operation. Returns the input file list.

SYNTAX

```
Blob.PushList [options] the input file(s)
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input file(s) - [required] - null

Blob.Remove

NAME

Blob.Remove – Remove the file attached to the input document as specified by the 'xpath' parameter. If the 'xpath' point to a blob list then the list will be cleared. If the file to remove is part of a list it will be removed from the list otherwise the 'xpath' should point to a blob property that will be removed. If the save parameter is set the document modification will be automatically saved. Return the document.

SYNTAX

```
Blob.Remove [options] [the input document(s)]
```

OPTIONS

- -save -

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

Blob.Set

NAME

Blob.Set – Set the input file to the given property on the input document. If the XPath points to a blob list then the blob is appended to the list, otherwise the XPath should point to a blob property. If the save parameter is set the document modification will be automatically saved. Return the document.

SYNTAX

```
Blob.Set [options] [the input document(s)]
```

OPTIONS

- -save -
- -file -
- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

Blob.SetFilename

NAME

Blob.SetFilename – Modify the filename of a file stored in the input document. The file is found in the input document given its xpath specified through the 'xpath' parameter. Return back the input document.

SYNTAX

```
Blob.SetFilename [options] [the input document(s)]
```

OPTIONS

- -save -
- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

BlobToFile

NAME

BlobToFile – Save the input blob(s) as a file(s) into the given target directory. The blob(s) filename is used as the file name. You can specify an optional **prefix** string to prepend to the file name. Return back the blob(s).

SYNTAX

```
BlobToFile [options] the input file(s)
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format

- -prefix -
- -void - [flag] - If void the server will not return the result back
- -directory -

ARGUMENTS

- the input file(s) - [required] - null

Blob.ToPDF

NAME

Blob.ToPDF – Convert the input file to a PDF and return the new file.

SYNTAX

```
Blob.ToPDF [options] the input file(s)
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input file(s) - [required] - null

Context.FetchDocument

NAME

Context.FetchDocument – Fetch the input of the context as a document. The document will become the input for the next operation.

SYNTAX

```
Context.FetchDocument [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Context.RestoreBlobInput

NAME

Context.RestoreBlobInput – Restore the file input from a context variable given its name. Return the file.

SYNTAX

```
Context.RestoreBlobInput [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Context.RestoreBlobsInput

NAME

Context.RestoreBlobsInput – Restore the file list input from a context variable given its name. Return the files.

SYNTAX

```
Context.RestoreBlobsInput [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Context.RestoreDocumentInput

NAME

Context.RestoreDocumentInput – Restore the document input from a context variable given its name. Return the document.

SYNTAX

```
Context.RestoreDocumentInput [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Context.RestoreDocumentsInput

NAME

Context.RestoreDocumentsInput – Restore the document list input from a context variable given its name. Return the document list.

SYNTAX

```
Context.RestoreDocumentsInput [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Context.RunDocumentOperation

NAME

Context.RunDocumentOperation – Run an operation chain which is returning a document in the current context. The input for the chain to run is the current input of the operation. Return the output of the chain as a document.

SYNTAX

```
Context.RunDocumentOperation [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -id -

ARGUMENTS

- the input document(s) - [optional] -null

Context.RunInputScript

NAME

Context.RunInputScript – Run a script from the input blob. A blob containing script result is returned.

SYNTAX

```
Context.RunInputScript [options] the input file(s)
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -type -

ARGUMENTS

- the input file(s) - [required] - null

Context.RunOperation

NAME

Context.RunOperation – Run an operation chain in the current context

SYNTAX

```
Context.RunOperation [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -id -

Context.RunScript

NAME

Context.RunScript – Run a script which content is specified as text in the 'script' parameter

SYNTAX

```
Context.RunScript [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -script -
- -void - [flag] - If void the server will not return the result back

Context.SetInputAsVar

NAME

Context.SetInputAsVar – Set a context variable that points to the current input object. You must give a name for the variable. This operation works on any input type and return back the input as the output.

SYNTAX

```
Context.SetInputAsVar [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Context.SetVar

NAME

Context.SetVar – Set a context variable given a name and the value. To compute the value at runtime from the current context you should use an EL expression as the value. This operation works on any input type and return back the input as the output.

SYNTAX

```
Context.SetVar [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -value -
- -name -
- -void - [flag] - If void the server will not return the result back

Document.CheckIn

NAME

Document.CheckIn – Checks in the input document. Returns back the document.

SYNTAX

```
Document.CheckIn [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -version -
- -versionVarName -
- -void - [flag] - If void the server will not return the result back
- -comment -

ARGUMENTS

- the input document(s) - [optional] -null

Document.CheckOut

NAME

Document.CheckOut – Checks out the input document. Returns back the document.

SYNTAX

```
Document.CheckOut [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.Copy

NAME

Document.Copy – Copy the input document into the given folder. The name parameter will be used as the copy name otherwise if not specified the original name will be preserved. The target folder can be specified as an absolute or relative path (relative to the input document) as an UID or by using an EL expression. Return the newly created document (the copy).

SYNTAX

```
Document.Copy [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -target -
- -name -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.Create

NAME

Document.Create – Create a new document in the input folder. You can initialize the document properties using the 'properties' parameter. The properties are specified as *key=value* pairs separated by a new line. The key used for a property is the property XPath. To specify multi-line values you can use a \ character followed by a new line.

Example:

```
dc:title=The Document Title  
dc:description=foo bar
```

Returns the created document.

SYNTAX

```
Document.Create [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -properties -
- -name -
- -void - [flag] - If void the server will not return the result back
- -type -

ARGUMENTS

- the input document(s) - [optional] -null

Document.CreateVersion

NAME

Document.CreateVersion – Create a new version for the input document. Any modification made on the document by the chain will be automatically saved. Increment version if this was specified through the 'snapshot' parameter. Returns the live document (not the version).

SYNTAX

```
Document.CreateVersion [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -increment -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.Delete

NAME

Document.Delete – Delete the input document. The previous context input will be restored for the next operation.

SYNTAX

```
Document.Delete [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.Fetch

NAME

Document.Fetch – Fetch a document from the repository given its reference (path or UID). The document will become the input of the next operation.

SYNTAX

```
Document.Fetch [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -value -
- -void - [flag] - If void the server will not return the result back

Document.FetchByProperty

NAME

Document.FetchByProperty – For each specified string property value, fetch all documents that match the property and the optional where clause. Matching documents are collected into a list and the returned to the next operation. The operation has no input.

SYNTAX

```
Document.FetchByProperty [options]
```

OPTIONS

- -values -
- -ctx - Can be used to inject context properties in Java properties format
- -property -
- -void - [flag] - If void the server will not return the result back
- -query -

Document.Filter

NAME

Document.Filter – Filter the input list of documents given a condition. The condition can be expressed using 4 parameters: types, facets, lifecycle and condition. If more than one parameter is specified an AND will be used to group conditions.

The 'types' parameter can take a comma separated list of document type: File, Note.

The 'facet' parameter can take a single facet name.

The 'life cycle' parameter takes a name of a life cycle state the document should have.

The 'condition' parameter can take any EL expression.

Returns the list of documents that match the filter condition.

SYNTAX

```
Document.Filter [options] [the input document(s)]
```

OPTIONS

- -class -
- -types -
- -pathStartsWith -
- -ctx - Can be used to inject context properties in Java properties format
- -facet -
- -void - [flag] - If void the server will not return the result back
- -lifecycle -
- -condition -

ARGUMENTS

- the input document(s) - [optional] -null

Document.GetChild

NAME

Document.GetChild – Get a child document given its name. Take as input the parent document and return the child document.

SYNTAX

```
Document.GetChild [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.GetChildren

NAME

Document.GetChildren – Get the children of a document. The list of children will become the input for the next operation

SYNTAX

```
Document.GetChildren [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.GetParent

NAME

Document.GetParent – Get the parent document of the input document. The parent document will become the input for the next operation. You can use the 'type' parameter to specify which parent to select from the document ancestors

SYNTAX

```
Document.GetParent [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -type -

ARGUMENTS

- the input document(s) - [optional] -null

Document.GetPrincipalEmails

NAME

Document.GetPrincipalEmails – Fetch the principal emails that have a given permission on the input document and then set them in the context under the given key variable name. The operation returns the input document. You can later use the list of principals set by this operation on the context from another operation. The 'key' argument represents the variable name and the 'permission' argument the permission to check. If the 'ignore groups' argument is false then groups are recursively resolved, extracting user members of these groups. Be **warned** that this may be a very consuming operation.

Note that:

- groups are not included,
- the list pushed into the context is a string list of emails.

SYNTAX

```
Document.GetPrincipalEmails [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -ignore groups -
- -variable name -
- -void - [flag] - If void the server will not return the result back
- -permission -

ARGUMENTS

- the input document(s) - [optional] -null

Document.GetUsersAndGroups

NAME

Document.GetUsersAndGroups – Fetch the users and groups that have a given permission on the input document and then set them in the context under the given key variable name. The operation returns the input document. You can later use the list of identifiers set by this operation on the context from another operation. The 'key' argument represents the variable name and the 'permission' argument the permission to check. If the 'ignore groups' argument is false then groups will be part of the result. If the 'resolve groups' argument is true then groups are recursively resolved, adding user members of these groups in place of them. Be **warned** that this may be a very consuming operation. If the 'prefix identifiers' argument is true, then user identifiers are prefixed by 'user:' and groups identifiers are prefixed by 'group:'.

SYNTAX

```
Document.GetUsersAndGroups [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -ignore groups -
- -resolve groups -
- -variable name -
- -void - [flag] - If void the server will not return the result back
- -permission -
- -prefix identifiers -

ARGUMENTS

- the input document(s) - [optional] -null

Document.Lock

NAME

Document.Lock – Lock the input document in the name of the given 'owner'. The lock owner is an username and identifies the user that owns the lock on the document. If the owner is not specified, the current user will be used as the owner. Returns back the locked document.

SYNTAX

```
Document.Lock [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -owner -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.Move

NAME

Document.Move – Move the input document into the target folder.

SYNTAX

```
Document.Move [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -target -
- -name -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.MultiPublish

NAME

Document.MultiPublish – Publish the input document(s) into several target sections. The target is evaluated to a document list (can be a path, UID or EL expression). Existing proxy is overridden if the override attribute is set. Returns a list with the created proxies.

SYNTAX

```
Document.MultiPublish [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -target -
- -void - [flag] - If void the server will not return the result back
- -override -

ARGUMENTS

- the input document(s) - [optional] -null

Document.Pop

NAME

Document.Pop – Restore the last saved input document in the context input stack. This operation must be used only if a PUSH operation was previously made. Return the last *pushed* document.

SYNTAX

```
Document.Pop [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Document.PopList

NAME

Document.PopList – Restore the last saved input document list in the context input stack

SYNTAX

```
Document.PopList [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Document.Publish

NAME

Document.Publish – Publish the input document into the target section. Existing proxy is overridden if the override attribute is set. Return the created proxy.

SYNTAX

```
Document.Publish [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -target -
- -void - [flag] - If void the server will not return the result back
- -override -

ARGUMENTS

- the input document(s) - [optional] -null

Document.Pull

NAME

Document.Pull – Restore the first saved input document in the context input stack. This operation must be used only if a PUSH operation was previously made. Return the first *pushed* document.

SYNTAX

```
Document.Pull [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Document.PullList

NAME

Document.PullList – Restore the first saved input document list in the context input stack

SYNTAX

```
Document.PullList [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Document.Push

NAME

Document.Push – Push the input document on the context stack. The document can be restored later as the input using the corresponding pop operation. Returns the input document.

SYNTAX

```
Document.Push [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.PushList

NAME

Document.PushList – Push the input document list on the context stack. The document list can be restored later as the input using the corresponding pop operation. Returns the input document list.

SYNTAX

```
Document.PushList [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.Query

NAME

Document.Query – Perform a query on the repository. The query result will become the input for the next operation.

SYNTAX

```
Document.Query [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -language -
- -void - [flag] - If void the server will not return the result back
- -query -

Document.Reload

NAME

Document.Reload – Reload the input document from the repository. Any previous modification made by the chain on this document will be lost if these modifications were not saved. Return the reloaded document.

SYNTAX

```
Document.Reload [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.RemoveACL**NAME**

Document.RemoveACL – Remove a named Acces Control List from the input document(s). Returns the document(s).

SYNTAX

```
Document.RemoveACL [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -acl -

ARGUMENTS

- the input document(s) - [optional] -null

Document.RemoveProperty**NAME**

Document.RemoveProperty – Remove the given property of the input document(s) as specified by the 'xpath' parameter. If the property points to a list then clear the list. Removing a property means setting it to null. Return the document(s).

SYNTAX

```
Document.RemoveProperty [options] [the input document(s)]
```

OPTIONS

- -save -
- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

Document.Save**NAME**

Document.Save – Save in the repository any modification that was done on the input document. Returns the saved document.

SYNTAX

```
Document.Save [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document.SaveSession

NAME

Document.SaveSession – Commit any changes made by the operation on the documents. This can be used to explicitly commit changes. This operation can be executed on any type of input. The input of this operation will be preserved as the input for the next operation in the chain.

SYNTAX

```
Document.SaveSession [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

Document.SetACE

NAME

Document.SetACE – Set Access Control Entry on the input document(s). Returns the document(s).

SYNTAX

```
Document.SetACE [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -grant -
- -overwrite -
- -void - [flag] - If void the server will not return the result back
- -user -
- -acl -
- -permission -

ARGUMENTS

- the input document(s) - [optional] -null

Document.SetLifeCycle

NAME

Document.SetLifeCycle – Follow the given transition on the input document life cycle state

SYNTAX

```
Document.SetLifeCycle [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -value -

- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document SetProperty

NAME

Document SetProperty – Set a single property value on the input document. The property is specified using its XPath. The document is automatically saved if 'save' parameter is true. If you unset the 'save' you need to save it later using Save Document operation. Return the modified document.

SYNTAX

```
Document SetProperty [options] [the input document(s)]
```

OPTIONS

- -save -
- -ctx - Can be used to inject context properties in Java properties format
- -value -
- -void - [flag] - If void the server will not return the result back
- -xpath -

ARGUMENTS

- the input document(s) - [optional] -null

Document Unlock

NAME

Document Unlock – Unlock the input document. The unlock will be executed in the name of the current user. An user can unlock a document only if has the UNLOCK permission granted on the document or if it the same user as the one that locked the document. Return the unlocked document

SYNTAX

```
Document Unlock [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Document Update

NAME

Document Update – Set multiple properties on the input document. The properties are specified as <i>key=value</i> pairs separated by a new line. The key used for a property is the property xpath. To specify multi-line values you can use a \ charcater followed by a new line.

Example:

```
dc:title=The Document Title
```

```
dc:description=foo bar.
```

Returns back the updated document.

SYNTAX

```
Document Update [options] [the input document(s)]
```

OPTIONS

- -save -

- -ctx - Can be used to inject context properties in Java properties format
- -properties -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Notification.SendEvent

NAME

Notification.SendEvent – Send a Nuxeo event.

SYNTAX

```
Notification.SendEvent [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -name -
- -void - [flag] - If void the server will not return the result back

Notification.SendMail

NAME

Notification.SendMail – Send an email using the input document to the specified recipients. You can use the asHTML parameter to specify whether your message is in HTML format or in plain text. Also you can attach any blob on the current document to the message by using the comma separated list of XPath expressions 'files'. If you XPath points to a blob list all blobs in the list will be attached. Return back the input document(s).

SYNTAX

```
Notification.SendMail [options] [the input document(s)]
```

OPTIONS

- -viewId -
- -message -
- -ctx - Can be used to inject context properties in Java properties format
- -from -
- -files -
- -void - [flag] - If void the server will not return the result back
- -to -
- -subject -
- -asHTML -

ARGUMENTS

- the input document(s) - [optional] -null

Relations.CreateRelation

NAME

Relations.CreateRelation – Create a relation between 2 documents. The subject of the relation will be the input of the operation and the object of the relation will be retrieved from the context using the 'object' field. The 'predicate' field specifies the relation predicate. Return back the subject document.

SYNTAX

```
Relations.CreateRelation [options] [the input document(s)]
```

OPTIONS

- -object -

- -ctx - Can be used to inject context properties in Java properties format
- -predicate -
- -void - [flag] - If void the server will not return the result back

ARGUMENTS

- the input document(s) - [optional] -null

Relations.GetRelations

NAME

Relations.GetRelations – Get the relations for the input document. The 'outgoing' parameter can be used to specify whether outgoing or incoming relations should be returned. Returns a document list.

SYNTAX

```
Relations.GetRelations [options] [the input document(s)]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -predicate -
- -void - [flag] - If void the server will not return the result back
- -outgoing -

ARGUMENTS

- the input document(s) - [optional] -null

Workflow.CreateTask

NAME

Workflow.CreateTask – Enable to create a task bound to the document.

Directive, **comment** and **due date** will be displayed in the task list of the user. In **accept operation chain** and **reject operation chain** fields, you can put the operation chain ID of your choice among the one you contributed. Those operations will be executed when the user validates the task, depending on whether he accepts or rejects the task. You have to specify a variable name (the **key for ...** parameter) to resolve target users and groups to which the task will be assigned. You can use Get Users and Groups to update a context variable with some users and groups. If you check **create one task per actor**, each of the actors will have a task to achieve, versus "the first who achieves the task makes it disappear for the others".

SYNTAX

```
Workflow.CreateTask [options] [the input document(s)]
```

OPTIONS

- -variable name for actors prefixed ids -
- -reject operation chain -
- -ctx - Can be used to inject context properties in Java properties format
- -directive -
- -create one task per actor -
- -accept operation chain -
- -additional list of actors prefixed ids -
- -due date -
- -void - [flag] - If void the server will not return the result back
- -comment -
- -task name -

ARGUMENTS

- the input document(s) - [optional] -null

Workflow.GetTask

NAME

Workflow.GetTask – List tasks assigned to this user or one of its group. Task properties are serialized using JSON and returned in a Blob.

SYNTAX

```
Workflow.GetTask [options]
```

OPTIONS

- -ctx - Can be used to inject context properties in Java properties format
- -void - [flag] - If void the server will not return the result back

print

NAME

print – Print operation(s) definition

SYNTAX

```
print [options] [operation]
```

OPTIONS

- -p - The password if any.
- -u - The username if any.
- -out - An optional file to save the operation definition into. If not used the definition will be printed on stdout.

ARGUMENTS

- operation - [optional] -The operation to print.

Related pages



[Built-in Commands \(Nuxeo Installation and Administration - 5.5\)](#)

• Configuration Commands

Namespace: *config*

Commands for configuring the shell.



[File System Commands \(Nuxeo Installation and Administration - 5.5\)](#)



[Nuxeo Server Commands \(Nuxeo Installation and Administration - 5.5\)](#)



[Shell Batch Mode \(Nuxeo Installation and Administration - 5.5\)](#)



[Shell Commands \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Shell Features \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Extending The Shell \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Shell Documentation \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[Shell Namespaces \(Nuxeo Enterprise Platform \(EP\) - 5.5\)](#)



[background](#)

[Configuration Commands \(Nuxeo Installation and Administration - 5.5\)](#)

Index:

- background
- color
- font
- settings
- theme

background

NAME

background – Modify the background color used by the shell. This command is available only in UI mode.

SYNTAX

color

NAME

color – Modify the foreground color used by the shell. This command is available only in UI mode.

SYNTAX

```
color
```

font

NAME

font – Modify the font used by the shell. This command is available only in UI mode.

SYNTAX

```
font
```

settings

NAME

settings – Print or modify the shell settings.

SYNTAX

```
settings [options] [name] [value]
```

OPTIONS

- -reset - [flag] - Reset settings to their defaults. Need to restart shell.

ARGUMENTS

- name - [optional] -The variable to print or set.
- value - [optional] -The variable value to set.

theme

NAME

theme – Modify the theme used by the shell. This command is available only in UI mode.

SYNTAX

```
theme [name]
```

ARGUMENTS

- name - [optional] -The theme name to set. If not specified the current theme is printed.

Related pages

-  [Built-in Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [Nuxeo Automation Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [File System Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [Nuxeo Server Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [Shell Batch Mode](#) (Nuxeo Installation and Administration - 5.5)
-  [Shell Commands](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Shell Features](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Extending The Shell](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Shell Documentation](#) (Nuxeo Enterprise Platform (EP) - 5.5)

• Shell Batch Mode

As I said in the overview section there are three modes to execute batch commands:

1. Run commands from a file
2. Run commands from standard input.
3. Run commands specified on the command line - this is a convenient way to run a short list of commands.

Running Commands From a File

To run commands from a file you should use the **-f** parameter to specify a file containing commands when launching Nuxeo Shell.

Example:

Shell Namespaces (Nuxeo Enterprise Platform (EP) - 5.5)

 Configuration Commands (Nuxeo Installation and Administration - 5.5)
java -cp nuxeo-shell.jar org.nuxeo.shell.Main -f my_batch_file

Where **my_batch_file** is a file containing the commands to execute - each command on one line. Empty lines and lines beginning with # are ignored. The # character can be used to add comments to a batch file.

Here is an example of a batch file:

```
# connect to local server using the Administrator account.  
connect -u Administrator -p Administrator http://localhost:8080/nuxeo/site/automation  
  
# we are now in the repository root. Go to /default-domain/workspaces  
cd /default-domain/workspaces  
  
# list the content of the workspaces root - as document UIDs  
ls -uid
```

If you want to span a command on multiple lines (you may want this for improved readability in case of long commands) you can end the line with a * character (***make sure** you don't have a space after **). In that case the command will continue on the next line, and so on until no more line ending ** is found or the end of file is reached.

Example:

```
# connect to local server using the Administrator account.  
connect -u Administrator -p Administrator http://localhost:8080/nuxeo/site/automation  
  
# get all workspaces in the repository  
query "SELECT * FROM Document WHERE ecm:primaryType='Workspace' \  
      ORDER BY dc:title DESC"
```

Running Commands From Standard Input

If you want to run batch commands from the terminal standard input you can use the - option when launching the Nuxeo shell. The format of the commands is the same as the one described when running commands from a file.

Here is an example which will run the commands from my_batch_file file by using the Unix **cat** application and pipes:

```
cat my_batch_file | java -cp nuxeo-shell.jar org.nuxeo.shell.Main -
```

Running Batch Commands from the Command Line

If you just run a few short commands you can specify them directly in the command line of the Nuxeo Shell.

Example:

```
java -cp nuxeo-shell.jar org.nuxeo.shell.Main -e "connect -u Administrator -p  
Administrator http://localhost:8080/nuxeo/site/automation; ls"
```

Note that commands are separated using a semicolon character.

⚠ Limitations

You cannot run that way commands that contains illegal characters and needs to be escaped.

Related pages

-  [Built-in Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [Nuxeo Automation Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [File System Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [Nuxeo Server Commands](#) (Nuxeo Installation and Administration - 5.5)
-  [Shell Batch Mode](#) (Nuxeo Installation and Administration - 5.5)
-  [Shell Commands](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Shell Features](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Extending The Shell](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Shell Documentation](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Shell Namespaces](#) (Nuxeo Enterprise Platform (EP) - 5.5)
-  [Configuration Commands](#) (Nuxeo Installation and Administration - 5.5)

(look at the data directory value in the Admin Center).

For Tomcat, it should be \$TOMCAT/nxserver/data. For JBoss, it should be \$JBoss/server/default/data.

• Backup and restore

Backup

Nuxeo supports hot backup of your data.

If you have followed the recommendations, then you have configured Nuxeo to use a "true" database (instead of the default Derby) and have set a path for nuxeo.data.dir in your nuxeo.conf. In that case:

1. simply first backup your database (make a SQL dump),
2. then backup your data on filesystem.

[Performing the backup in that order \(the database first, then the filesystem\) will ensure backup consistency \(more details\).](#)

If you didn't configure Nuxeo to use a database, then the default database is embedded in the data directory but you will have to stop the server before backup.

If you didn't configure Nuxeo data directory (nuxeo.data.dir in nuxeo.conf), then you have to find the default path which depends on the server (Tomcat/JBoss) and the Nuxeo version

Restore

1. Restore the database and data filesystem you had previously backed up.
2. [Configure Nuxeo](#) to use those database and data directory.
3. Start Nuxeo.

Import a file system

Multi threaded document importer from server file system

Build/install

From Nuxeo Marketplace

The Bulk document importer v. 1.0.0 is available from [the Nuxeo Marketplace](#).

By hand

Its source code lives at <https://github.com/nuxeo/nuxeo-platform-importer>

Build with: mvn package -Dmaven.test.skip=true

And deploy the two jars from the target subfolder into the nuxeo.ear/plugins or nxserver/bundles folder of your nuxeo server (and restart).

Usage

The file importer comes as a Java library (with nuxeo runtime service) and a sample JAX-RS interface to launch, monitor and abort import jobs.

To import the folder '/home/ogrissel/Documents' into the workspace /default-domain/workspaces/my-workspace while monitoring the import logs from a REST client use the following HTTP GET queries:

```
GET http://localhost:8080/nuxeo/site/fileImporter/logActivate  
  
GET  
http://localhost:8080/nuxeo/site/fileImporter/run?targetPath=/default-domain/workspace  
s/my-workspace&inputPath=/home/ogrisel/Documents&batchSize=10&interactive=false&nbThre  
ads=4  
  
GET http://localhost:8080/nuxeo/site/fileImporter/log
```

To execute those HTTP queries you can either use a browser with an active Nuxeo session (JSESSIONID cookie) or use a third party stateless HTTP client with HTTP Basic Authentication, eg with the curl command line client:

```
curl --basic -u 'Administrator:Administrator'  
"http://localhost:8080/nuxeo/site/fileImporter/log"
```

Don't forget to quote the URL if it includes special shell characters such as '&'.

You can also use the generic HTTP GUI client from [the rest-client java project](#).

Don't forget to fill in the 'Auth' tab with your user credentials.

For developers

If you want to programmatically reuse or extend multi-threaded import on a Nuxeo repository, look at [Nuxeo Bulk Document Importer](#).

Upgrading your Nuxeo Version

Each version of Nuxeo comes with bug fixes, new features and improvements.

This means upgrading to the latest public release is a smart move.

In order to make upgrade easy we are very careful not to break anything.

Remember that we provide support and use Nuxeo EP in a lot of project, so if we break something, we have to fix it :)

So basically:

- we don't break the APIs between two versions: we add new APIs and deprecate old ones,
- there are several minor versions between the deprecation and the removal so you have time to adapt your code,
- if we completely replace a service (that was the case of SearchService and EventService for example), we provide compatibility packages so you can continue using the old API (even if migrating to the new API is highly recommended).

In terms of Data Migration we are also very careful not to break anything.

Once again, we run Nuxeo EP for a lot of our internal needs and we upgrade them very frequently and don't want to have data migration issue.

Anyway, when some changes or optimizations impact the storage layer we either:

- make the upgrade automatically,
- or provide guidelines for the upgrade.

Upgrading is usually a simple and painless process.

Using the template system also allows to easily transpose your configuration from one version to an other.

In the worse cases, in case of problem, Nuxeo Support is there to help you 😊

For upgrade steps, here are the upgrade procedures per Nuxeo version:

- [Upgrade from 5.4.2 to 5.5](#)
- [Upgrade to 5.4.2 with Oracle](#)
- [Upgrade from 5.3.2 to 5.4.0](#)
- [Upgrade from 5.3.1 to 5.3.2](#)
- [Upgrade from 5.3.1 with MySQL to 5.3.2](#)
- [Upgrade from 5.3.0 to 5.3.1](#)
- [Upgrade from 5.1.6 with JCR + PostgreSQL to 5.2.0](#)
- [Upgrade from 5.1.2 to 5.1.3](#)

Upgrade from 5.4.2 to 5.5

This chapter presents the detailed process to upgrade from Nuxeo 5.4.2 to Nuxeo 5.5. Most of it is useful information you need to have to fully understand what has changed in this release.

You can have a look at this interesting use case of a "Mostly painless Nuxeo upgrade from 5.4.2 to 5.5 under Windows / PostgreSQL"!

Installation & Configuration

Follow [Installation Instruction](#).

Under all OS:

- H2 Embedded database is not supported for data detection.
- After the installation, uncomment the line `custom.target` in your `template/custom/nuxeo.default` and set it to `".` to indicate the path of your custom templates.
- `"session.timeout"` has been added to `nuxeo.conf` and can be overridden for defining the web session timeout which is then integrated into the `web.xml` file.

Under Linux:

- The new installation changes the `opt/nuxeo` content location to `/var/lib/nuxeo/server` - templates and conf still in `etc/nuxeo` folder.
- The package Debian autoconfigure
 - detects your data,
 - detects/adds your marketplace addons (including DM,CMF,DAM).

Under Windows:

- Windows installer:
 - detects your data,
 - detects/adds your marketplace addons (including DM,CMF,DAM).

Packaging

The DM, DAM, SC, and CMF distributions are now available as Marketplace packages.

This new packaging system is used in the Setup Wizard to allow to choose between different profiles at installation time.

You can also use the Admin Center or the [nuxeoctl commands](#) to add or remove these packages. For projects having a custom distribution based on one of ours, no problem, we provide presets for automatically transforming the new unique Tomcat distribution into a DM, DAM or CMF.

Also, the "EAR" (zip) assemblies do still exist.

Using the wizard is just an additional option.

Digital Asset & Case Management

For now, it is not possible to install CMF with other packages like DM and DAM because there are some content model incompatibilities.

Distribution

Regarding to custom distributions and related to the new 5.5 packaging (DM, CMF, DAM are now addons), Ant assembly script (`assembly.xml`) has to be modified:

Deploy the Nuxeo CAP distribution (only `nuxeo-cap` classifier still exists):

```
<!-- Deploy CAP distribution -->
<unzip dest="${stagedir}">
    <artifact:resolveFile
key="org.nuxeo.ecm.distribution:nuxeo-distribution-tomcat:${nuxeo.version}:zip"
                                classifier="nuxeo-cap" />
</unzip>

${stagedir}: distribution parent folder
```

Define type distribution:

```
<!-- Set the addon deploying in distribution -->
<copy file="${app.path}/nxserver/data/installAfterRestart-?.log"
      tofile="${app.path}/nxserver/data/installAfterRestart.log"
      overwrite="true" />

${app.path}: define your distribution path (ie ./stage/nuxeo-custom-server)
?: DM,DAM,CMF,SC
```

Optional: choose the wizard distribution type by setting wizard addon preset ([NXP-8031](#)):

```
<!-- Set the wizard.preset by default -->
<echo file="${app.path}/setupWizardDownloads/packages-default-selection.properties"
      message="preset=nuxeo-?" />

${app.path}: define your distribution path (ie ./stage/nuxeo-custom-tomcat)
?: dm,cmf,dam
```

Third party libraries upgrades

- **GWT** - Nuxeo is now using **GWT 2.4.0**.
- **JAX-WS** - Libraries have been upgraded to 2.2.5 in order to fix some compatibilities issues.
- **OpenCMIS** - Nuxeo Platform is now aligned on **OpenCMIS 0.6** that comes with experimental support for the CMIS Browser binding (JS compliant API).
- **JEXL** - Location is changed from Nuxeo Runtime to Nuxeo Platform Action.

Data Migration

VCS

Only fews column additions were done between 5.4.2 and 5.5 (no alter). So you can migrate and retrieve all your 5.4.2 data after 5.5 installation.

For relations, these attributes are added to the "relation" schema ([NXP-7962](#)):

```
<xss:element name="predicate" type="xs:string" />
<xss:element name="sourceUri" type="xs:string" />
<xss:element name="targetUri" type="xs:string" />
<xss:element name="targetString" type="xs:string" />
```

A new metadata has been added to follow the legacy definition of dublincore schema ([NXP-7884](#)) :

```
<xss:element name="publisher" type="xs:string"/>
```

A new schema has been added: task.xsd (related to the nuxeo-platform-task feature - [NXP-7852](#)):

```
<xss:element name="actors" type="nxt:stringList" /> (Task actors list)
<xss:element name="task_variables" type="nxt:task_variables" /> (tasks vars list)
<xss:element name="taskComments" type="nxt:taskComments" /> (Task comments list)
```

(Four new tables due to complex types added.)

Fulltext

Partially missing fulltext index for the title field

Old versions of Nuxeo DM might have existing documents present before the introduction of the "fulltext_title" index. This is visible on the Nuxeo 5.5 release thanks to the new search suggestion widget that might be missing some suggestions on old documents.

To update the title fulltext index, just perform the following SQL query on your PostgreSQL server:

```
UPDATE fulltext SET simpletext_title = NX_TO_TSVECTOR("dublincore"."title") FROM dublincore WHERE "fulltext"."id" = "dublincore"."id";
```

PostgreSQL fulltext phrase search

In Nuxeo 5.5 for PostgreSQL we've added a better way to store fulltext information that enables the use of phrase search. If you want to use phrase search, you should follow the upgrade notes of [NXP-5689](#). If you do this fulltext upgrade, you may want to check the (unsupported for now) [nuxeo-reindex-fulltext plugin](#) to get more accurate phrase search results.

If you don't do the upgrade described in [NXP-5689](#), you'll get the following error message:

```
Cannot use phrase search in fulltext compatibility mode. Please upgrade the fulltext table: ...
```

Directories

Directories with auto-incremented columns must be upgraded, as the mechanism for auto-increment has been changed to be more robust. Please follow the [NXP-7124](#) upgrade notes if you have auto-incremented columns (there aren't any in a default Nuxeo installation).

Code Migration

Nuxeo 5.5 is mainly backward compatible with Nuxeo 5.4.2. If you have any problems, you can contact Nuxeo Support.

Automation

Changes in Nuxeo Automation: there was a Java package renaming from `org.nuxeo.ecm.automation.client.jaxrs.model` to `org.nuxeo.ecm.automation.client.model`.

Nuxeo Theme

Nuxeo Theme service has been extended so that you can now contribute page styles in a plain CSS stylesheet.

The page layouts are still managed by the Theme engine using an XML description, but all CSS information is now externalized to CSS stylesheets that can manage flavors (pretty much as with LessCSS).

- [Theme documentation page](#)
- [Migrating a custom theme](#)

Tasks

Until 5.5, the Task system was directly bound to JBPM.

Starting with 5.5, a new TaskService is available and uses VCS to store tasks.

This new TaskService is a first step towards the integration of Content Routing as the default Workflow engine in DM.

Migration should be 100% transparent:

- the Task Operations have not changed,
- REST APIs are maintained,
- Tasks created in JBPM and not directly associated to a process will be automatically migrated upon first access,
- JBPM Tasks are still accessible via the new TaskService,
- the JBPM task API is maintained.

Relations

The new default configuration takes care about compatibility so that if you have existing relations in Jena graph you will still be able to transparently access them.

On this page

- Installation & Configuration
 - Packaging
 - Digital Asset & Case Management
 - Distribution
 - Third party libraries upgrades
 - Data Migration
 - VCS
 - Fulltext
 - Partially missing fulltext index for the title field
 - PostgreSQL fulltext phrase search
 - Directories
 - Code Migration
 - Automation
 - Nuxeo Theme
 - Tasks
 - Relations

Upgrade to 5.4.2 with Oracle

If you were using Nuxeo DM 5.3.2 or later with Oracle and ACL optimizations enabled, you need to update an index before plugging your database to Nuxeo DM 5.4.2.

```
SELECT index_name FROM USER_INDEXES WHERE table_name LIKE 'READ_ACLS';
```

Get the index name, result of the previous query, which should look like "SYS_C00XXXX"

Then run the following statement:

```
ALTER INDEX "SYS_C00XXXX" RENAME TO "ACLR_ACL_ID_IDX";
```

Now you can start Nuxeo DM 5.4.2 which will rename the table READ_ACLS to ACLR and the previous index will be ready to work with this new configuration.

Upgrade from 5.3.2 to 5.4.0

Installation & Configuration

While Nuxeo EP 5.3.2 was based on the JBoss distribution, Nuxeo EP 5.4.0 is based on Tomcat. There is also a distribution based on JBoss 5, though we strongly recommend you to use the Tomcat distribution: http://cdn.nuxeo.com/nuxeo-5.4.0/nuxeo-dm-5.4.0_01-tomcat.zip.

JBoss 5

If using the JBoss distribution, see [Upgrade to 5.4 and JBoss 5](#).

Data Migration

On JBoss, data used to be in `$NUXEHOME/server/default/data/NXRuntim`e, where you can find the `binaries` folder. Now, with the Tomcat distribution, it stands in `$NUXEHOME/nxserver/data`.

- If you have moved your data outside Nuxeo using the `nuxeo.data.dir`, as recommended, you need to move the `binaries` folder one level up and then remove the `NXRuntim`e folder.
- If you haven't moved your data outside Nuxeo, you need to move the `binaries` folder from `$NUXEHOME/server/default/data/NXRuntim`e to `$NUXEHOME/nxserver/data`.

Code Migration

Workflow

The workflow implementation has changed, see [From the old workflow system to the new 5.4 workflow system](#).

From the old workflow system to the new 5.4 workflow system

Even though the 5.4 jBPM service doesn't implement backward compatibility with the old workflow service, is it possible to deploy both the old workflow framework AND the jBPM service so that the migration can be gradual?

It should be possible to use the m3 version of jBPM and still run the new workflow.

If you create an action for the new workflow tab different from the one for the old workflow, you could use both at the same time. (I assume you are not using publishing tab and forum).

However, I think it would be much easier to move your old workflow to the new one. That would be:

- create the new workflow with handler,
- define what variable you need in the new workflow (most probably docId, repo name),
- get the variable from the unfinished process instance in the old jBPM table (object are just serialized into it).

See [NXP-2850](#) for technical details.

Upgrade to 5.4 and JBoss 5

During the migration from JBoss 4.2 to 5.1 we had to do some small changes in Nuxeo bundles.

You will have to do the same for your custom bundles.

Nevertheless, the needed changes are very small and it should only take a few minutes (it should be really quick when you don't have 200 bundles and you know what to do).

Impact on sources and resources

Web resources

Bundles contribution resources to the WAR used to have these resources stored in: `src/main/resources/nuxeo.war`.

This does break JBoss 5 deployment, because it tries to deploy the `nuxeo.war` as a nested WAR, but it fails (because the WAR is not complete and because there are several bundles containing the same `nuxeo.war`).

Rather than changing JBoss deployer config we chose to change our packaging to avoid the problem.

The solution is simple: `nuxeo.war` tree should not be at root of the JAR.

Sample directory structure:

Before (JBoss 4.2)	After (JBoss 5.1)
<pre>. `-- src `-- main `-- java `-- resources -- META-INF `-- nuxeo.war `-- my_page.xhtml `-- OSGI-INF</pre>	<pre>. `-- src `-- main `-- java `-- resources -- META-INF `-- OSGI-INF `-- web `-- nuxeo.war `-- my_page.xhtml</pre>

This implies a small change in the deployment-fragment:

Before (JBoss 4.2)	After (JBoss 5.1)
<pre><install> <unzip from="\${bundle.fileName}" to="/" > <include>nuxeo.war/**</include> </unzip> </install></pre>	<pre><install> <unzip from="\${bundle.fileName}" to="/" prefix="web" > <include>web/nuxeo.war/**</include> </unzip> </install></pre>

If this is not done, or if there is still a "nuxeo.war" directory present in your jar, you will get an exception like:

```
DEPLOYMENTS IN ERROR:
Deployment "vfsfile:/opt/jboss/server/default/deploy/nuxeo.ear/" is in error due
to the following reason(s):
java.lang.IllegalStateException: jboss.web.deployment:war=/nuxeo is already
installed.
```

Module declaration

From within the deployment fragment you can declare contributions to the application.xml. Be sure that if you declare your module as a EJB module it does really contains EJB3.

```
<extension target="application#MODULE">
<module>
<ejb>${bundle.fileName}</ejb>
</module>
</extension>
```

Otherwise, you should remove the contribution to `application.xml` (the Java declaration is not needed and will in fact be ignored by the pre-deployer).

Typical deprecated contribution:

```
<extension target="application#MODULE">
<module>
<java>${bundle.fileName}</java>
</module>
</extension>
```

Web services binding

The previous versions of Nuxeo were using SUN-JAX-WS (Metro) to handle WebService deployment.
In order to avoid having to modify default JBoss 5.1/EAP 5.0.1 config, we now support to deploy on JBoss native stack (JBoss WS).

This involves doing some small changes in the way the WebService are implemented and deployed.

No more EJB3

Because of some limitation of the JBossWS EJB3 deployer we can not deploy WebServices on top of EJB3 if we want to keep endpoint URLs consistent.

So this simply means you should remove the @Stateless annotation in the Beans providing WebService.

Before (Jboss 4.2)	After (Jboss 5)
<pre>@Local(WSAuditLocal.class) @Stateless @Remote(WSAudit.class) @WebService(name = "WSAuditInterface", serviceName = "WSAuditService") @SOAPBinding(style = Style.DOCUMENT) public class WSAuditBean extends AbstractNuxeoWebService implements</pre>	<pre>@Local(WSAuditLocal.class) @Remote(WSAudit.class) @WebService(name = "WSAuditInterface", serviceName = "WSAuditService") @SOAPBinding(style = Style.DOCUMENT) public class WSAuditBean extends AbstractNuxeoWebService implements</pre>

SUN-JAX-WS vs JBossWS binding

Both WS framework need to have a declaration for endpoints, but they (of course) don't use the same way to do it.
Basically, SUN-JAX-WS uses a dedicated sun-jax-ws.xml file and JBossWS uses the web.xml.

The Nuxeo template system are configured so that you can declare both bindings in your bundle and Nuxeo will deploy the right one depending on the target deployment host.

JBossWS binding	SUN-JAX-WS binding
<pre><!-- JBossWS Native EndPoint Declaration --> <extension target="web#JBOSWSWS"> <servlet> <description>NuxeoAudit WS EndPoint</description> <display-name>NuxeoAudit EndPoint</display-name> <servlet-name>NuxeoAuditEndPoint</servlet-name> <servlet-class>org.nuxeo.ecm.platform.audit.ws.WS AuditBean</servlet-class> </servlet> <servlet-mapping> <servlet-name>NuxeoAuditEndPoint</servlet-name> <url-pattern>/webservices/nuxeoaudit</url-pattern > </servlet-mapping> </extension> NB: Yes, you declare your Bean as a servlet even if it does not implement the needed interface!</pre>	<pre><!-- SUN JAXWS / Metro EndPoint Declaration --> <extension target="jaxws#ENDPOINT"> <endpoint name="nuxeoaudit" implementation="org.nuxeo.ecm.platform.audit.ws.W SAuditBean" url-pattern="/webservices/nuxeoaudit" /> </extension></pre>

Requirements declaration

In JBoss deployed, as it was already the case for tomcat deployment, the Require-Bundle is no longer used by the deployed and should be reserved for OSGi deployment.

The Nuxeo-Require should not longer be used either. Nuxeo-RequiredBy will be deprecated but should still be used to override XHTML

pages.

In the JBoss 5 deployment (as it was the case in Tomcat), all the Nuxeo bundles are put in the classloader from the start, so there is no risk of ClassNotFound during the loading.

The only dependencies you have to worry about are Runtime dependencies (server availability, contribution override ...): all this should be managed by using the <require> tag in the XML contribution.

NB: the Require-Bundle is still used when Nuxeo is deployed on an OSGi container, because in this case the class loading issues remains.

Other checks you may want to do

ejb-jar.xml files

Unless you know what you are doing, you should remove any ejb-jar.xml file from the META-INF folder.

Servlets and filters initialization

Servlets and Filters should not assume that when activated (call to init method by the servlet container) Nuxeo Runtime is ready. It won't be always the case.

So if you need to do some Nuxeo Runtime calls at init time, you should rely on the Framework initialization event, rather on the servlet container init.

EJB3 declaration

JBoss 5 is more strict on the JEE spec, so your EJB3 can not use the same Java interface for @Local and @Remote

Web.xml ordering

JBoss 5 validate the web.xml against the DTD and checks order on the tags.

The Nuxeo template is OK and respect the standard, so if you contribute your filters in the rights section and the servlets in the right sections there should be no problem.

But be aware that if you took some shortcuts and contributed several kind of objects (Filters, Servlets, Context-params) in side the same slot, it may have worked OK with JBoss 4 but it won't with JBoss 5.

CoreSession usage

In JBoss environment, CoreSession is delivered via DocumentManagerBean that is a Stateful Session Bean.

JEE spec does not allow concurrent calls on a SFSB : this applies to both JBoss 4 and JBoss 5.

But in the case of JBoss 5 the check is more strict and also prohibits reentrant calls from the same thread.

The typical use case if you create a DocumentModel and you have a Listener that will run and use the core session to do some work.

In JBoss 4 it runs without problem but in JBoss 5 this is detected as a concurrent call.

DocumentManagerBean and DocumentModel implementations have been modified to avoid that, but there are still cases where you may have the problem.

Typical unsafe code is taking the CoreSession via:

```
CoreInstance.getInstance().getSession(doc.getSessionId()).
```

This can be replaced by

```
doc.getCoreSession()
```

that contains the needed logic to detect the reentrant call and return the correct Session (Local or EJB) depending on the context.

So if you have errors like this one:

```
no concurrent calls on stateful bean
' jboss.j2ee:service=EJB3,name=DocumentManagerBean' (EJB3 4.3.13)
```

first check that you don't access the CoreSession from inside a Listener using the DocumentModel sessionId.

Dependencies

Hibernate dependencies in Nuxeo's root `pom.xml` has changed. The core artifact for Hibernate is named `hibernate-core` now instead of `hibernate`. If you were using this dependency, you need to change from:

```
<dependency>
  <groupId>org.hibernate</groupId>
  <artifactId>hibernate</artifactId>
</dependency>
```

to:

```
<dependency>
  <groupId>org.hibernate</groupId>
  <artifactId>hibernate-core</artifactId>
</dependency>
```

Impact on the Packaging

Templates

Following structure changes in `nuxeo.ear`, templates structure has changed a little.

If you override a template, check the directories. For instance, here are the default template changes:

Before (Jboss 4.2)	After (Jboss 5.1)
<pre>. -- config -- default-repository-config.xml `-- sql.properties -- datasources -- default-repository-ds.xml `-- unified-nuxeo-ds.xml `-- nuxeo.defaults</pre>	<pre>. -- nuxeo-ds.xml -- nuxeo.defaults `-- nuxeo.ear -- META-INF '-- default-repository-ds.xml `-- config `-- default-repository-config.xml `-- sql.properties</pre>

As defined in `nuxeo.defaults ("default.target=server/default/deploy")`, the target directory of this template is now `server/default/deploy` instead of `server/default/deploy/nuxeo.ear`.

If you defined your own template, you only have to follow the `nuxeo.ear` structure:

- `default-repository-ds.xml` has moved from `nuxeo.ear/datasources/` to `nuxeo.ear/META-INF/`
- `unified-nuxeo-ds.xml` has moved to `nuxeo-ds.xml` and its content now includes mbeans declarations
- `nuxeo.ear/system/` has been renamed to `nuxeo.ear/bundles/`

Assemblies

Deprecated compliance artifacts with old resources were removed.

`nuxeo-platform-ear` and `nuxeo-distribution-dm` now only contain:

```
nuxeo.ear/
|-- bundles
`-- lib
```

All resources files are generated from templates directories.

Also, the default EAR archives contain in `lib/` all third-party libraries from the dependency tree (only duplicates are removed). The filtering of provided libraries is done when building the server distribution (JBoss, Tomcat, ...).

In this section

- Impact on sources and resources
 - Web resources
 - Module declaration
 - Web services binding
 - No more EJB 3
 - SUN-JAX-WS vs JBoss WS bin ding
 - Requirements declaration
 - Other checks you may want to do
 - ejb-jar.xml files
 - Servlets and filters initialization
 - EJB 3 declaration
 - Web.xml ordering
 - CoreSession usage
 - Dependencies
- Impact on the Packaging
 - Templates
 - Assemblies

Upgrade from 5.3.1 to 5.3.2

Code migration

5.3.2 is fully backward compatible with 5.3.1 (no compat package is needed).

So, you should have no issues with running your custom code against 5.3.2. If you have any problems, you can contact Nuxeo Support.

Packaging

The packaging system is basically the same as the one used in 5.3.1.

The only change that may have an impact involves resources that are now managed by the new template system.

This means that resources are no longer embedded inside the EAR but handled in a separated templates directory.

This makes changing configurations easier (like switching from H2 to PostgreSQL) and will also allow for upgrades without having to redo all custom system configurations.

The documentation has been updated accordingly:

- [Installation Guide](#)
- [Description of the new configuration system](#)

Data

The only changes done between 5.3.1 and 5.3.2 are the way tags are stored.

Because the Tag Service is now directly part of VCS, some small changes have been done.

Nevertheless, migration should be automatic and transparent.

If you have any problems, you can contact Nuxeo Support.

Configuration

We have changed the way Nuxeo starts OpenOffice.

This is an intermediate solution before we upgrade to JODConverter 3.

The new OOolauncher (that replaces OOodeamon) should:

- be more stable,
- be easier to set up (removed the dependencies on JNI UNO libs).

Upgrade from 5.3.1 with MySQL to 5.3.2

Why a migration script is needed

A database structure change was introduced with Nuxeo 5.3.2 to fix query with operator IN (for more details, see <https://jira.nuxeo.com/browse/NXP-5183>).

Below is an example of structure change:

Sample structure in Nuxeo DM <= 5.3.1

```
CREATE TABLE `common` (
  `id` varchar(36) NOT NULL,
  `icon` varchar(4000) DEFAULT NULL,
  `icon-expanded` varchar(4000) DEFAULT NULL,
  `size` bigint(20) DEFAULT NULL,
  PRIMARY KEY (`id`),
  KEY `common_id_hierarchy_fk` (`id`),
  CONSTRAINT `common_id_hierarchy_fk` FOREIGN KEY (`id`) REFERENCES `hierarchy` (`id`)
ON DELETE CASCADE
) ENGINE=InnoDB DEFAULT CHARSET=latin1
```

Sample structure in Nuxeo DM >= 5.3.2

```
CREATE TABLE `common` (
  `id` varchar(36) CHARACTER SET latin1 COLLATE latin1_bin NOT NULL DEFAULT '',
  `icon` varchar(4000) DEFAULT NULL,
  `icon-expanded` varchar(4000) DEFAULT NULL,
  `size` bigint(20) DEFAULT NULL,
  PRIMARY KEY (`id`),
  KEY `common_id_hierarchy_fk` (`id`),
  CONSTRAINT `common_id_hierarchy_fk` FOREIGN KEY (`id`) REFERENCES `hierarchy` (`id`)
ON DELETE CASCADE
) ENGINE=InnoDB DEFAULT CHARSET=latin1
```

See the part CHARACTER SET latin1COLLATE latin1_bin that was added for the column id.

As a consequence, the columns (mainly id) used for foreign key don't have the same definition and it is no more possible to add new constraints, needed by the features of Nuxeo DM. trying to upgrade to Nuxeo DM 5.4.x will raise this error

```
java.sql.SQLException: Can't create table 'nuxeodb.#sql-5f31_37d' (errno: 150)
```

Steps for the migration

You need to follow the steps below to migrate your database structure:

- download the script upgradeMySQL.sh attached to this page
- edit these properties in the file upgradeMySQL.sh

```
DB_HOST=localhost
DB_PORT=3306
DB_NAME=nuxeo
DB_USER=user
DB_PWD=password
```

- change permission for the script to run it

```
$ chmod u+x upgradeMySQL.sh
```

- launch the script

```
$ ./updateMySQL.sh
```

- if everything is fine, you'll have a message to confirm the upgrade was done

```
Database structure upgraded successfully
```

Upgrade to Nuxeo DM 5.4.1

Now your database is upgraded, you can test it against Nuxeo DM 5.4.1

Once you've downloaded and unzipped it,

- start it without editing anything: it's needed because you'll need to get the lastest hot fixes to make it work

```
$ cd $NUXEO_HOME/bin  
$ ./nuxeoctl start
```

- follow the wizard steps and register on Nuxeo Connect (needed to download Hot Fixes) if you don't have already an account
- restart
- log in using the default credentials
- navigate to Nuxeo Admin Center > Update Center > Software updates tab
- download and install all availables hot fixes (actually at least the three first ones)
- stop the server

```
$ ./nuxeoctl stop
```

- edit the configuration file `nuxeo.conf` and set the database parameters

```
nuxeo.templates=mysql  
nuxeo.db.name=nuxeodb  
nuxeo.db.user=user  
nuxeo.db.password=password  
nuxeo.db.host=localhost  
nuxeo.db.port=3306
```

- restart the server

```
$ ./nuxeoctl start
```

- contemplate that all data are present 😊

Upgrade from 5.3.0 to 5.3.1

Nuxeo DM 5.3.1 is fully backward compatible with Nuxeo DM 5.3.0 GA, hence upgrade is painless and requires no data migration or code change.

Follow these steps to upgrade.

1. Get the differences between a vanilla Nuxeo DM 5.3.0 and your custom Nuxeo

If you have specific configuration, you'll need to know what files were changed, in order to apply them onto the default one. Here are the folders and files you have to watch:

- \$JBOSS/server/default/deploy/nuxeo.ear/config/ => main configuration elements of Nuxeo
- \$JBOSS/server/default/deploy/nuxeo.ear/datasources/ => configuration of data sources
- \$JBOSS/server/default/conf/ => specific configuration of Jboss
- \$JBOSS/server/default/lib/ => specific libraries used by your project (JDBC drivers for instance)
- \$JBOSS/server/default/deploy/mail-service.xml => configuration of the mail service
- \$JBOSS/server/default/deploy/nuxeo.ear/OSGI-INF/templates/web.xml

2. Backup your data

- Follow this [documentation](#)

3. Apply the differences

- From the differences you got at step 1, apply them on the Nuxeo DM 5.3.1 you've downloaded
- Copy your specific plugins into \$JBOSS/server/default/deploy/nuxeo.ear/plugins/

4. Restore data

- Copy the data folder (server/default/data) to Nuxeo DM 5.3.1

Upgrades notes

Groups stored in SQL directory

If you were using groups stored in the SQL directory, you have to consider that the "group2group" table must be fixed as its columns were inverted. childGroupId should be populated with the content of parentGroupId and vice-versa. It is related to [NXP-4401](#)

Before applying the command below, you have to check that your SQL configuration has changed. This will be the case if you get the new **default-sql-directories-bundle.xml** or if your patch doesn't change the **tableReference** attribute defined in the new **default-sql-directories-bundle.xml**. The attribute should look like:

```
<tableReference field="subGroups" directory="groupDirectory"
    table="group2group" sourceColumn="parentGroupId"
    targetColumn="childGroupId" schema="group2group" />
```

In that case, run the following query for PostgreSQL to update the table:

```
UPDATE group2group SET "childGroupId" = "parentGroupId", "parentGroupId" =
"childGroupId";
```

OpenSocial

- The `opensocial.properties` file format has been changed in 5.3.1, so you may need to use the one provided in 5.3.1 rather than trying to upgrade the one used in 5.3.
- For a dashboard initialized in 5.3.0, existing OpenSocial gadgets need to be migrated: see [NXP-4923](#) for script and procedure.

Indexing

There are no impacting changes on the storage structure.

If you want to leverage the new default indexing configuration (separated full-text index for title and description), you will have to update your repository configuration (or use the one provided with 5.3.1) and build the new indexes.

For developers

If you use a custom Nuxeo assembly to package your Nuxeo distribution with your plugins, you will need to modify your existing assembly.

The new `nuxeo-distribution` system is simpler to configure and extend than the previous one.

See [here](#) and [here](#) for more details.

Upgrade from 5.1.6 with JCR + PostgreSQL to 5.2.0

This article will help you to migrate your data from Nuxeo 5.1.6 to Nuxeo 5.2 in the case you are using JackRabbit with PostgreSQL as backend.

We assume that your Nuxeo 5.1.6 is installed in \$JBOSS_516 directory and Nuxeo 5.2 in \$JBOSS_52 and you have well configured your Nuxeo 5.2 to work with Jackrabbit/PSQL. Otherwise, let's see this [article](#).

The steps to migrate are:

- Start an empty nuxeo 5.2 configured in JCR
 - Customize a nuxeo 5.2 JCR with an **empty** database, created for the occasion.
 - Start nuxeo 5.2 and log in
 - shutdown nuxeo 5.2
- copy the file \$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/nodetypes/custom_nodetypes.xml and keep it in a temporary location
- keep either the directory \$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/namespaces/
- remove \$JBOSS_52/server/default/data
- copy the data folder from \$JBOSS_516/server/default/data to \$JBOSS_52/server/default/data
- copy the custom_nodetypes.xml file you kept to \$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/nodetypes/
- change searchIndex class to org.nuxeo.ecm.core.repository.jcr.jackrabbit.SearchIndex in \$JBOSS_52/server/default/data/NXRuntime/repos/default/worksheets/default/workspace.xml
- remove the \$JBOSS_52/server/default/data/NXRuntime/repos/default/worksheets/default/index folder to force JackRabbit to rebuild the indexes
- update discriminator column in NXP_LOGS table to allow this value to be null
 - alter table NXP_LOGS alter discriminator DROP not null
- Here is the tricky part, customize the ns_idx.properties in the directory namespaces that you kept:

```
Compare the file  
$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/namespaces/ns_idx.properties  
with the one you kept from the namespaces directory.  
They contain uri and an identifier, example :  
http://www.nuxeo.org/ecm/schemas/common/=19  
Each identifier is unique!  
You need to adapt the ns_idx.properties kepted in order that each uri keep its old  
identifier unchanged .
```

simple example :

\$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/namespaces/ns_idx.properties

```
http://www.nuxeo.org/ecm/schemas/common/=21  
http://www.nuxeo.org/ecm/schemas/dublincore/=18  
http://project.nuxeo.org/schemas/webengine/site/blog/post=19
```

\$JBOSS_516/server/default/data/NXRuntime/repos/default/repository/namespaces/ns_idx.properties

```
http://www.nuxeo.org/ecm/schemas/common/=19  
http://www.nuxeo.org/ecm/schemas/dublincore/=18
```

As you see in the 516 file, <http://www.nuxeo.org/ecm/schemas/common/> was identified by 19, so we need to keep this identifier, but <http://project.nuxeo.org/schemas/webengine/site/blog/post> is already identified by 19 so we will just switch the two identifiers.

Here is the new file :

```
http://www.nuxeo.org/ecm/schemas/common/=19  
http://www.nuxeo.org/ecm/schemas/dublincore/=18  
http://project.nuxeo.org/schemas/webengine/site/blog/post=21
```

- remove the directory \$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/namespaces/

- copy your namespace directory (With the customized ns_idx.properties) in \$JBOSS_52/server/default/data/NXRuntime/repos/default/repository/namespaces/
- finally adapt nuxeo 5.2 to use the 5.1.6 database.

Just for information, below are the changes you can make manually to update your custom_nodetypes.xml

Two main problems occurs are present in the node type definitions from Nuxeo 5.1.6:

- the whole versioning features are not working : no document modification, no version increase, no reading of the previous versions, ...: this is due to the fact that `ecm:version` and `ecm:versionHistory` are not mixin type any more. Manually you can change these nodes and chose `isMixin="true"` to `isMixin="false"`
- some document definitions have changed :
 - Workspace type has two new supertypes: `ecmst:publish_ergo` and `ecmst:webcontainer`
 - Forum, Thread and post types use now "`ecmdt:Document`" as supertype instead of "`ecmnt:document`"
 - WikiPage and BlogPost types use "`ecmmix:versionable`" as supertype instead of "`mix:versionable`"

Editing the custom_nodetypes file is not easy because you have to format this file (`tidy -xml ...`) to edit it. So we recommend to replace the old custom_nodetypes.xml by the new one, generated from a fresh Nuxeo 5.2 installation.

Upgrade from 5.1.2 to 5.1.3

Follow [Upgrade Nuxeo](#) and apply the following procedure _before_ starting Nuxeo.

While upgrading from 5.1.2 to 5.1.3, you may have to manage with a Blob format issue : that means to patch \$JBOSS/server/default/data/NXRuntime/repos/default/repository/nodetypes/custom_nodetypes.xml

Take this file, format it with tidy (`tidy -wrap 999 -indent -xml`) and apply this patch (manually as it can't guarantee any line numbers; add the lines beginning with a "+" if not already present):

```
+ <nodeType hasOrderableChildNodes="false" isMixin="true" name="ecmmix:content"
primaryItemName=" " >
+   <propertyDefinition autoCreated="false" mandatory="false" multiple="false"
name="digest" onParentVersion="COPY" protected="false" requiredType="String" />
+   <propertyDefinition autoCreated="false" mandatory="false" multiple="false"
name="length" onParentVersion="COPY" protected="false" requiredType="Long" />
+   <propertyDefinition autoCreated="false" mandatory="false" multiple="false"
name="filename" onParentVersion="COPY" protected="false" requiredType="String" />
+ </nodeType>
<nodeType hasOrderableChildNodes="false" isMixin="false" name="ecmft:content"
primaryItemName=" " >
  <supertypes>
    <supertype>ecmnt:property</supertype>
+   <supertype>ecmmix:content</supertype>
+   <supertype>nt:resource</supertype>
  </supertypes>
  <propertyDefinition autoCreated="false" mandatory="false" multiple="false"
name="mime-type" onParentVersion="COPY" protected="false" requiredType="String" />
    <propertyDefinition autoCreated="false" mandatory="false" multiple="false"
name="data" onParentVersion="COPY" protected="false" requiredType="Binary" />
    <propertyDefinition autoCreated="false" mandatory="false" multiple="false"
name="encoding" onParentVersion="COPY" protected="false" requiredType="String" />
  </nodeType>
```

Then, you need to re-index your data. Using [nuxeo-shell](#) (versus web function in advanced search) is recommended.

Another solution (than patching custom_nodetypes.xml file) is to export then re-import data before and after the upgrade (using nuxeo-shell too); but this method will make you loose versioning information.

Marketplace addons

The Nuxeo Marketplace is Nuxeo's ECM application store. It offers plugins and packages that enable you to easily add features to your Nuxeo application. Packages are listed by module (Nuxeo DM, Nuxeo DAM), and by categories (workflow, collaborative tools...). The list of available packages is available to everyone, but some packages require a Nuxeo Connect account to be able to install packages.

The Marketplace offers packages aimed at developers and other that provide new features to end-users. Most of the packages can be [installed from the Update Center](#) very easily and don't require any additional installation or configuration step. However, some other addons, typically connectors with other systems, involve some additional configuration.

Each package has a dedicated page on the Marketplace, that describes the feature the package enables, if there are prerequisites, etc. Here is the information available about the packages from the Marketplace:

- **Production state:** Indicates if the package is approved by Nuxeo or is still in testing phase.
- **Certification status:** Indicates if the packages has been certified by Nuxeo or not.
- **Vendor support:** Indicates if the package is covered by Nuxeo Connect support contracts.
- **Type:** Possible types are: addons will provide new features, hot-fixes provide corrections, and studio packages install your Studio customizations in your instance.
- **Last version:** Most recent version number of the plugin.
- **Updated:** Date on which the package was last updated.
- **Target platforms:** Nuxeo applications on which you can install the package.
- **License:** License applied to the package.
- **Categories:** List of categories the package belongs to.
- **Rating:** Comments on the package.
- **Vendor:** Name of the person or company who developed the package.
- **Package dependencies:** indicates if there are some requirements for the package to be correctly installed.
- **Hot-reload support:** Indicates if the plugin is immediately functional (ie, no server reboot required).

Although most packages are installed in a few clicks from the Update Center, some of them require specific installation or configuration steps. These packages are listed below.

Package name	Public / Registered access	Comment
Amazon S3 Online Storage	Registered access	

The user guide holds a Marketplace section in which you will find the user documentation for Marketplace addons.

Amazon S3 Online Storage

The Amazon S3 Online Storage is a Nuxeo Binary Manager for S3. It stores Nuxeo's binaries (the attached documents) in an [Amazon S3](#) bucket.

Before you start

The S3 Binary Manager requires Nuxeo at least DM 5.4.1.

You should be familiar with Amazon S3 and be in possession of your credentials.

Installing the package

Use the Update Center to [install the package](#) from the Nuxeo Marketplace.

In this section:

- Before you start
- Installing the package
- Configuring the package
- Activate the custom template
- Checking your configuration

Configuring the package

In order to configure the package, you will need to change a few Nuxeo templates, and provide values for the configuration variables that define your S3 credentials, bucket and encryption choices.

Create a new default repository config

1. Find the XML file `default-repository-config.xml` from a reference template. For instance, if in `nuxeo.conf` you use the template `nuxeo.templates=postgresql`, then the XML file is `templates/postgresql/nxserver/config/default-repository-config.xml`.
2. Copy this XML file into `templates/custom/nxserver/config/default-repository-config.xml` (you may need to create the intermediate `nxserver` and `config` folders).
3. Then modify this file to add a `<binaryManager>` section like follows:

```
...
<repository name="default" factory="...">>
  <repository name="default">
    <binaryManager class="org.nuxeo.ecm.core.storage.sql.S3BinaryManager" />
    ...
  </repository>
</repository>
...
```

Activate the custom template

In the `nuxeo.templates` property of `nuxeo.conf`, you must then activate the S3 repository configuration by adding `custom`; for instance:

```
nuxeo.templates=postgresql,custom
```

Check that the target for the `custom` template is set to `'.'`: the `custom` template folder should have a `nuxeo.defaults` file with the line:

```
custom.target=.
```

Specify your Amazon S3 parameters

In `nuxeo.conf`, add the following lines:

```
nuxeo.s3storage.bucket=your_s3_bucket_name
nuxeo.s3storage.awsid=your_AWS_ACCESS_KEY_ID
nuxeo.s3storage.awssecret=your_AWS_SECRET_ACCESS_KEY
```

If you installed the bundle JAR manually instead of using the marketplace package you will also need:

```
nuxeo.core.binarymanager=org.nuxeo.ecm.core.storage.sql.S3BinaryManager
```



The bucket name is unique across all of Amazon, you should find something original and specific.



The file `nuxeo.conf` now contains S3 secret access keys, you should protect it from prying eyes.

You can also add the following optional parameters:

```
nuxeo.s3storage.region=us-west-1
nuxeo.s3storage.cachesize=100MB
```

The region code can be:

- for `us-east-1` (the default), don't specify this parameter
- for `us-west-1` (Northern California), use `us-west-1`
- for `eu-west-1` (Ireland), use `EU`
- for `ap-southeast-1` (Singapore), use `ap-southeast-1`

Since 5.6, you can also use:

- for us-west-2 (Oregon), us-west-2
- for ap-southeast-2 (Tokyo), use ap-southeast-2
- for sa-east-1 (Sao Paulo), use sa-east-1

Crypto options

With S3 you have the option of storing your data encrypted (note that the local cache will *not* be encrypted).

The S3 Binary Manager can use a keystore containing a keypair, but there are a few caveats to be aware of:

- The Sun/Oracle JDK doesn't always allow the AES256 cipher which the AWS SDK uses internally. Depending on the US export restrictions for your country, you may be able to modify your JDK to use AES256 by installing the "Java Cryptography Extension Unlimited Strength Jurisdiction Policy Files". See the following link to download the files and installation instructions:
<http://www.oracle.com/technetwork/java/javase/downloads/index.html>
- Don't forget to specify the key algorithm if you create your keypair with the `keytool` command, as this won't work with the default (DSA). The S3 Binary Manager has been tested with a keystore generated with this command:

```
keytool -genkeypair -keystore </path/to/keystore/file> -alias <key alias>
-storepass <keystore password> -keypass <key password> -dname <key distinguished
name> -keyalg RSA
```

If you get `keytool` error: `java.io.IOException: Incorrect AVA format`, then ensure that the distinguished name parameter has a form such as: `-dname "CN=AWS S3 Key , O=example, DC=com"`.

! Don't forget to **make backups of the `/path/to/keystore/file` file along with the store password, key alias and key password**. If you lose them (for instance if the EC2 machine hosting the Nuxeo instance with the original keystore is lost) you will lose the ability to recover any encrypted blob from the S3 bucket.

With all that above in mind, here are the crypto options that you can add to `nuxeo.conf` (they are all mandatory once you specify a keystore):

```
nuxeo.s3storage.crypt.keystore.file=/absolute/path/to/the/keystore/file
nuxeo.s3storage.crypt.keystore.password=the_keystore_password
nuxeo.s3storage.crypt.key.alias=the_key_alias
nuxeo.s3storage.crypt.key.password=the_key_password
```

i The Nuxeo `S3BinaryManager` class is using **S3 Client-Side Encryption** instead of **S3 Server-Side Encryption**. CSE is safer than SSE. With CSE an attacker need both access to the **AWS credentials and the key** to be able to access the unencrypted data while SSE will only require the potential attacker to provide the **AWS credentials**.

Checking your configuration

To check that installation went well, you can check your startup logs and look for a line like:

```
INFO [S3BinaryManager] Repository 'default' using S3BinaryManager
```

Don't forget to enable the `INFO` level for the group `org.nuxeo` in `$NUXEO_HOME/lib/log4j.xml` to see `INFO` level messages from Nuxeo classes.

If your configuration is incorrect, this line will be followed by some error messages describing the problems encountered.

Related topics

No content found for label(s) amazon-s3 nuxeo_conf templates.

Nuxeo GSA Connector

The Nuxeo GSA Connector is a Nuxeo DM addon that enables you to have your Nuxeo application indexed by your Google Search Appliance and so search its content directly from the Google website.

The package includes:

- a Google Enterprise Connector Manager,
- a Google Search Appliance Nuxeo connector type.

Before you start

The Nuxeo GSA Connector requires Nuxeo DM 5.4.0.1 or Nuxeo DM 5.4.1.

Installing the package

1. In the Update Center, download the Nuxeo GSA Connector by clicking on the **Download** button.
2. Click on the **Install** button.
3. Click on the **Start** button to start the installation.
A configuration form is displayed to help you set up the Nuxeo GSA Connector with your GSA configuration.
4. Fill in the form (see below for the descriptions of the fields).
5. Click on the **Submit query** button.
6. When the installation is done, click on the **Finish** button.
You now need to restart your Nuxeo DM server to finish the installation process.

Package configuration properties

Field	Description
Nuxeo Base URL	The URL used to access Nuxeo DM.
GSA host name FQDN	Complete host name of the Google Search Appliance (FQDN stands for "Full Qualified Domain Name").
GSA feed port	GSA port used to submit documents for indexing. Default value is 19900 and should work in most cases.

Configuring the GSA

After Nuxeo DM is restarted to complete the installation of the Nuxeo GSA Connector, you need to configure the GSA so it indexes the Nuxeo DM documents and returns them in search results.

1. Log in the GSA admin GUI.

2. Set up the URL pattern for the connector:
 - a. Go to: Crawl and Index/Crawl URLs.
 - b. In the Follow and Crawl Only URLs with the Following Patterns **textarea**, add: ^googleconnecto
r://nuxeo_connector.localhost/.

The screenshot shows the 'Google Search Appliance >Crawl and Index >Crawl URLs' page. The left sidebar contains a navigation menu with sections like Home, Crawl and Index (selected), Crawl URLs, Databases, Feeds, Crawl Schedule, Crawler Access, Proxy Servers, Forms Authentication, Case-Insensitive Patterns, HTTP Headers, Duplicate Hosts, Document Dates, Host Load Schedule, Freshness Tuning, Collections, Composite Collections, Serving, Status and Reports, Connector Administration, Social Connect, Cloud Connect, GSA Unification, GSA^n, and Administration. The main content area has three sections: 'Start Crawling from the Following URLs:', 'Follow and Crawl Only URLs with the Following Patterns:', and 'Do Not Crawl URLs with the Following Patterns:'. The 'Follow and Crawl Only URLs with the Following Patterns:' section contains the URL '^googleconnector://nuxeo_connector.localhost/'. A note at the bottom right of this section says '*required'. Below the patterns section is a note: 'example: http://www.myorganization.mycompany.com/' followed by another '*required' note. The 'Do Not Crawl URLs with the Following Patterns:' section is currently empty. At the bottom right of the main content area is a 'Save URLs to Crawl' button. The footer of the page includes the Google logo and the text '© Copyright Google Inc. - [About](#)'.

- c. Save the URLs.

3. Add a Nuxeo connector manager:

- a. Go to: Connector Manager Administration/Connector Managers.
- b. Go to: Register a New Connector Manager:.
- c. Enter:
 - Manager name: nuxeo_mycompany
 - Description: The Nuxeo DM of my-company
 - Service URL: <http://nuxeo.in.my-company.com/nuxeo/gsa-connector-manager>

The screenshot shows the 'Connector Managers' page of the Google Search Appliance administration interface. The left sidebar has a 'Connector Administration' section expanded, showing 'Connector Managers' selected. The main content area displays a table of connector managers:

Status	Name	Description	Service URL	Version	Action
Internal	ConnectorManager0	Connector Manager	http://ent1:7886/connector-manager/	Google Search Appliance Connector Manager 2.6.4 (build 2623 September 21 2010)	Logs Export

Below the table, there is a form for 'Register a New Connector Manager':

Manager name:

Description:

Service URL:

Buttons: Save, Cancel

- d. Save.

4. Add a Nuxeo connector:

- Go to: Connector Manager Administration/Connectors.

- Select your newly created Connector manager (in our example `nuxeo_mycompany`).
- Click on the button Add New Connector.
- Type the Connector Name: `nuxeo_connector`.
- Click on the button Get Configuration Form.

⚠ The name must match a part of the URL crawl pattern (^googleconnector://[nuxeo_connector].localhost/).

- Configure the Traversal Rate and Retry Delay.
- Save the configuration.

Status	Name	Type	Version	Actions
●	nuxeo_connector	nuxeo-connector	1.1.2-SNAPSHOT	Edit Reset Delete

The indexing of Nuxeo DM documents starts immediately. You can check how it goes on by visiting the Crawl and Index/Feeds page or Status and Reports/Crawl Status.

The Nuxeo DM documents will then be displayed in the Google search results (see the #Using the Nuxeo GSA Connector section).

Using the Nuxeo GSA Connector

The use of the Nuxeo GSA Connector is transparent to users. After it's been installed by an administrator from the Admin Center, users will find Nuxeo DM documents in the search results of their local Google website. They just need to click on them to be directed on the document in Nuxeo DM.

The access rights of Nuxeo DM are taken into account in the search results. The first time search results include Nuxeo DM documents, users are displayed a login pop-up window. They need to log in to be able to see Nuxeo DM results and they can only see documents they can access in Nuxeo DM.



User's credentials are then saved for the next searches.

Purge audit logs

Depending on usage (lots of updates, lots of workflows, lots of logins ...), audit tables can grow very quickly.

You can configure the audit to select what must be recorded, but there is no API or UI to do cleanup inside audit tables. Actually, this is not something we forgot, we simply considered that it was safer like that: the Audit Service is here to record activity in the platform ; it makes sense that a component cannot easily delete its trail.

This means the cleanup should be done at SQL level. Since the table structure is really obvious, this is an easy job. However, you can find below the source of a PostgreSQL function that can be used to purge Audit entries older than a given date.

You can easily adapt it:

- to change the filtering done on audit record to filter
- to match other DB

nx_audit-purge

```
CREATE OR REPLACE FUNCTION nx_audit_purge(olderThan character varying)
RETURNS int AS
$BODY$
DECLARE
-- INPUT format is 'YYYY-MM-DD'
maxDate varchar(11) := olderThan;
nblines int;
total int;
BEGIN
-- Because nxp_logs_mapextinfos has 2 FK on external tables
-- we must remove records from this table first
-- so we need to store the values in a tmp table before
CREATE TEMP TABLE audit_purge_tmp ON COMMIT DROP AS
SELECT nxp_logs_mapextinfos.log_fk, nxp_logs_mapextinfos.info_fk
FROM nxp_logs, nxp_logs_extinfo, nxp_logs_mapextinfos
WHERE
nxp_logs.log_event_date < maxDate::date
AND nxp_logs_mapextinfos.log_fk = nxp_logs.log_id
AND nxp_logs_mapextinfos.info_fk=nxp_logs_extinfo.log_extinfo_id;
-- CLEANUP on nxp_logs_mapextinfos bridge table first to drop constraints
RAISE INFO 'run cleanup on nxp_logs_mapextinfos (level 2)';
DELETE
FROM nxp_logs_mapextinfos
WHERE nxp_logs_mapextinfos.log_fk IN (
SELECT log_fk FROM audit_purge_tmp);
GET DIAGNOSTICS nblines = ROW_COUNT;
SELECT nblines INTO total;
RAISE INFO '% lines cleanup on table nxp_logs_mapextinfos' ,nblines;
-- LEVEL 3 cleanup

RAISE INFO 'run cleanup on nxp_logs_extinfo (level 3)';

DELETE
FROM nxp_logs_extinfo
WHERE nxp_logs_extinfo.log_extinfo_id IN (
SELECT info_fk FROM audit_purge_tmp);
GET DIAGNOSTICS nblines = ROW_COUNT;
SELECT nblines+total INTO total;
RAISE INFO '% lines cleanup on table nxp_logs_extinfo' ,nblines;
-- LEVEL1 cleanup
RAISE INFO 'run cleanup on nxp_logs (level 1)';

DELETE
FROM nxp_logs
WHERE nxp_logs.log_id IN (SELECT log_fk FROM audit_purge_tmp);
GET DIAGNOSTICS nblines = ROW_COUNT;
SELECT nblines+total INTO total;

RAISE INFO '% lines cleanup on table nxp_logs' ,nblines;
RAISE INFO '% lines total cleanup ' ,total;

RETURN total;
END $BODY$
LANGUAGE plpgsql VOLATILE
COST 100;
ALTER FUNCTION nx_audit_purge(character varying)
OWNER TO nuxeo;
```

