

Deployit System Administration Manual

Version 3.6.1

Table of Content

Preface	3
Installing Deployit	3
Prerequisites	3
Server Requirements	3
Determining Harddisk Space Requirements	3
Unix Middleware Server Requirements	4
Windows Middleware Server Requirements	4
Extending Middleware Support	4
Client Requirements	4
GUI Clients	4
CLI Clients	5
Installation Procedure	5
Installing the Server	5
Installing the CLI	5
Running the Server Setup Wizard	6
Simple Setup	6
Manual Setup	7
Finishing the Setup Process	9
Changing the Admin Password	9
High Availability Setup	9
Upgrading Deployit	10
Upgrading the Server	10
Upgrading the CLI	10
Configuring Deployit	10
Configuring Security	10
Security in Deployit	10
Granting, Revoking and Denying	11
Configuring Repository Security	12
Inherited Permissions	12
Permissions on Deployed Applications	12
Security Configuration Example	12
Configuring LDAP Security	13
Configuring the Repository	14
Using a Database	14
Installing Plugins	16
Configuring Logging	16
Setting up Deployit	17
Starting and Stopping	17
Editing CIs	18
Maintaining Deployit	18
Creating Backups	19
Restoring Backups	19

Preface

This manual describes how to install and setup Deployit.

Installing Deployit

This section contains information on the installation of the Deployit server.

Prerequisites

Server Requirements

To install the Deployit server, the following prerequisites must be met:

- **Operating system:** Windows or Unix-family operating system running Java.
- **Java Runtime Environment:** JDK 1.6 (Oracle, IBM or Apple)
- **RAM:** At least 2GB of RAM available for Deployit.
- **Harddisk space:** Sufficient harddisk space to store the Deployit repository. This depends on your usage of Deployit. See section **Determining Harddisk Space Requirements**.

Depending on the environment, the following may also be required:

- **Database:** Deployit's Jackrabbit repository supports a number of different databases. For more information see section **Configuring the Repository**.
- **LDAP:** To enable group-based security, an LDAP x.509 compliant registry is needed. For more information see section **Configuring LDAP Security**.

Determining Harddisk Space Requirements

The Deployit server itself only uses about 70MB of disk space. The main harddisk space usage comes from the repository which stores your deployment packages and deployment history. The size of the repository will vary from installation to installation but depends mainly on:

- the size and storage mechanism used for artifacts
- the number of packages in the system
- the number of deployments performed (more specifically, the amount of logging information stored)

Follow this procedure to obtain an estimate of the total required disk space:

- Install and configure Deployit for your environment as described in this document. Make sure you correctly set up the database- or file-based repository.
- Estimate the number of packages to be imported (either the total number or the number per unit of time) (NumPackages)
- Estimate the number of deployments to be performed (either the total number or the number per unit of time) (NumDeployments)
- Record the amount of disk space used by Deployit (InitialSize).
- Import a few packages using the GUI or CLI.
- Record the amount of disk space used by Deployit (SizeAfterImport).
- Perform a few deployments.
- Record the amount of disk space used by Deployit (SizeAfterDeployments).

The needed amount of disk space in total is equal to:

$\text{Space Needed} = ((\text{SizeAfterImport} - \text{InitialSize}) * \text{NumPackages}) + ((\text{SizeAfterDeployments} - \text{SizeAfterImport}) * \text{NumDeployments})$

If `NumPackages` and `NumDeployments` are expressed per timeunit (e.g. the number of packages to be imported per month), then the end result represents the space needed per month as well.

Unix Middleware Server Requirements

Unix-based middleware servers that Deployit interacts with must meet the following requirements:

- **SSH Access:** The target systems should be accessible by SSH from the Deployit server, i.e. they should run an SSH2 server. It is also possible to handle key-based authorization. Notes:
 - The SSH daemon on AIX is known to hang with certain types of SSH traffic.
 - For security, the SSH account that is used to access a host should have limited rights.
 - A variety of Linux distributions have made SSH require a TTY by default. This setting is incompatible with Deployit and is controlled by the `Defaults requiretty` setting in the `sudoers` file.
- **Credentials:** Deployit should be able to log in to the target systems using a login/password combination that allows it to perform at least the following Unix commands:
 - `cp`
 - `ls`
 - `mv`
 - `rm`
 - `mkdir`
 - `rmdir`

If the login user cannot perform these actions, Deployit can also use a `sudo` user that can execute these commands.

Windows Middleware Server Requirements

Windows-based middleware servers that Deployit interacts with must meet the following requirements:

- **File system access:** The target file system should be accessible via CIFS from the Deployit server.
- **Host access:** The target host should be accessible from the Deployit server via WinRM or Windows Telnet server running in *stream mode*.
- **Directory shares:** The account used to access a target system should have access to the host's administrative shares such as `C$`.
- **Ports:** For CIFS connectivity, port 445 on the target system should be accessible from the Deployit server. For Telnet connectivity, port 23 should be accessible from the Deployit server. For WinRM connectivity, port 5985 (HTTP) or port 5986 (HTTPS) should be accessible from the Deployit server.

Extending Middleware Support

It is possible to connect Deployit to middleware servers that do not support SSH, Telnet or WinRM. Using the Overthere remote execution framework, a custom *access method* can be created that connects to the server. See the **Customization Manual** for more details.

Client Requirements

GUI Clients

To use the Deployit GUI, you must meet the following requirements:

- **Web browser:** The following web browsers are supported:
 - IE 7.0 or up
 - Firefox 3.0 or up
 - Safari 3.0 or up
- **Flash Player:** A flash player is required, versions 9.0 and up are supported.

CLI Clients

To use the Deployit CLI, you must meet the following requirements:

- **Operating system:** Windows or Unix-family operating system running Java.
- **Java Runtime Environment:** JDK 1.6 (Oracle, IBM or Apple)

Installation Procedure

To begin installing Deployit, first unpack the distribution archive. The distribution archive contains the following:

- Release notes describing the changes made in this version of Deployit.
- A server archive.
- A CLI archive.

Installing the Server

Follow these steps to install the Deployit server application:

1. **Login to the server where the Deployit Server will be installed.** It is recommended to install Deployit Server as a non-root user, e.g. `deployit`.
2. **Create an installation directory**, e.g. `/opt/xebialabs/deployit`.
3. **Copy the Deployit Server archive to the directory.**
4. **Extract the archive into the directory.**

Deployit Server Directory Structure

Once the Deployit installation file is extracted, the following directory structure exists in the installation directory (in the remainder of the document this directory will be referred to as `DEPLOYIT_SERVER_HOME`):

- **bin**: contains the Server binaries
- **conf**: contains Server configuration files (this directory is only present once you have configured Deployit Server)
- **doc**: contains the Deployit product documentation
- **ext**: contains Server Java extensions. See the **Customization Manual** for more information.
- **hotfix**: contains hotfixes that fix issues with the Server software
- **importablePackages**: default location for importable packages
- **lib**: contains libraries that Server needs
- **log**: contains Server log files (this directory is only present once you have started Deployit Server)
- **plugins**: contains the Deployit middleware plugins
- **recovery.dat**: stores tasks that are in progress for recovery purposes (this file is only present once you have started Deployit Server)

Installing the CLI

Follow these steps to install the Deployit CLI application:

1. **Login to the server where the Deployit CLI will be installed.**
2. **Create an installation directory.**
3. **Copy the Deployit CLI archive to the directory.**
4. **Extract the archive into the directory.**

Deployit CLI Directory Structure

Once the Deployit installation file is extracted, the following directory structure exists in the installation directory:

- **bin**: contains the CLI binaries
- **ext**: contains CLI Python extension scripts
- **hotfix**: contains hotfixes that fix issues with the CLI software

- **lib**: contains necessary libraries
- **plugins**: contains the CLI plugins

Running the Server Setup Wizard

Run the Deployit Setup Wizard to start the Deployit server and prepare it for use. The command `server.sh -setup` starts the wizard. If you want to stop the Setup Wizard at any time, enter `exitsetup`. All changes to the configuration will be discarded.

The Setup Wizard displays the following welcome message:

```
Welcome to the Deployit setup.
You can always exit by typing 'exitsetup'.
To re-run this setup and make changes to the Deployit server configuration
you can run server.cmd -setup on Windows or server.sh -setup on Unix.

Do you want to use the simple setup?
Default values are used for all properties. To make changes to the default
properties, please answer no.
Options are yes or no.
[yes]:
```

Answer **yes** (or press Enter) to use the simple setup. Simple setup makes it easy to quickly get started with Deployit and to use the product's default configuration. See **Simple Setup** for more information.

Answer **no** to use the manual setup. Manual setup provides explicit control over all Deployit settings. See **Manual Setup** for more information.

Note: if you installed Deployit in the same location before, the Setup Wizard will ask you whether you want to edit the existing configuration or create a new one. Answer **yes** (or press Enter) to edit the existing configuration. The Setup Wizard will load all settings from the existing configuration and allow you to choose simple or manual setup. Answer **no** to start over with an empty configuration.

Simple Setup

Using simple setup, the Setup Wizard will assume default values for all configuration parameters. Specifically, the following defaults will be used:

- The server will run with security enabled.
- The server will **not** use secure communication between the Deployit GUI and the Deployit server.
- The server will listen on Deployit's standard HTTP port (4516).
- The server will use a minimum of 3 and a maximum of 24 threads.
- Applications can be imported from the `importablePackages` directory.

The Setup Wizard will ask one more question:

```
Do you want Deployit to initialize the JCR repository?
Options are yes or no.
[yes]:
```

Answer **yes** (or press Enter) if you want the Deployit repository to be recreated. The Setup Wizard must have write access to the repository directory. Answer **no** to leave the repository intact. This option is useful if you already have an existing repository that you want to reuse.

See **Finishing the Setup Wizard** for completing the setup process.

Warning: if you choose to recreate the Deployit repository and you have installed Deployit in the same location before,

any information stored in the repository will be lost.

Manual Setup

The manual setup procedure contains the following steps:

Secure Communication Configuration

The Setup Wizard will show the following message:

```
Would you like to enable SSL?
Options are yes or no.
[yes]:
```

Answer **no** to use regular unsecured communication between the GUI and the server. Continue with the **port configuration** section.

Answer **yes** (or press Enter) if you want to use a secure connection from the GUI to the server.

If you answer **yes**, the Setup Wizard will ask the following question to help you configure secure communication:

```
Would you like Deployit to generate a keystore with a self-signed
certificate for you?
N.B.: Self-signed certificates do not work correctly with some versions
of the Flash Player and some browsers!
Options are yes or no.
[yes]:
```

Answer **yes** (or press Enter) if you want the Setup Wizard to generate a digital certificate automatically. The digital certificate is required to secure communication and is normally signed by a Certificate Authority (CA). The Setup Wizard can generate a *self-signed* certificate if there is no official certificate available. Beware that using a self-signed certificate may trigger security warnings in some Flash players and browsers. Continue with the **port configuration** section.

Answer **no** if you want to use your own keystore. Deployit uses the built-in Jetty webserver to communicate with the GUI. Jetty requires a certificate with the name `Jetty` to be present in the keystore.

The Setup Wizard prompts you for the following keystore information:

```
What is the path to the keystore?
[]:

What is the password to the keystore?
[]:

What is the password to the key in the keystore?
[]:
```

Enter the filesystem location of the keystore (for example, *mykeystore.jks*), the password to unlock the keystore and the password for the `Jetty` certificate in the keystore.

Port Configuration

The Setup Wizard shows the following question:

```
What http port number would you like the server to listen on?
[4516]:
```

Note: if you chose to enable secure communication, the default port will be 4517 instead of 4516.

Enter the port number that the Deployit server listens on for connections.

Thread Configuration

The Setup Wizard shows the following questions:

```
Enter the minimum number of threads the HTTP server should use (recommended:
  3 per client, so 3 for single user usage)
[3]:
```

Enter the minimum number of threads that the Deployit server uses to handle incoming connections. The recommended minimum number of threads is 3 per Deployit application client.

```
Enter the maximum number of threads the HTTP server should use (recommended :
  3 per client, so 24 for 8 concurrent users)
[24]:
```

Enter the maximum number of threads that the Deployit server uses to handle incoming connections. The recommended maximum number of threads is 3 per Deployit application client.

Repository Configuration

The Setup Wizard shows the following questions:

```
Where would you like Deployit to store the JCR repository?
[repository]:
```

Enter the filesystem path to a directory where Deployit will create the repository. If the directory does not exist, the Setup Wizard will create it.

```
Do you want Deployit to initialize the JCR repository?
Options are yes or no.
[yes]:
```

Answer **yes** (or press Enter) if you want the Deployit repository to be recreated. The Setup Wizard must have write access to the repository directory.

Answer **no** to leave the repository intact.

Warning: if you choose to recreate the Deployit repository and you have installed Deployit in the same location before, any information stored in the repository will be lost.

Importable Packages Configuration

The Setup Wizard shows the following question:

```
Where would you like Deployit to import packages from?
[importablePackages]:
```


Enter the filesystem path to a directory from which Deployit will import packages. The Setup Wizard assumes that this directory exists once the Deployit server starts and will not create it.

Finishing the Setup Process

Once you have completed configuration of the setup process, the Setup Wizard displays an overview of all selected options. The following text is an example:

```

Do you agree with the following settings for Deployit and would you like
to save them?
Changes will be saved in deployit.conf
  SSL will be disabled
  HTTP port is 4516
  HTTP server will use a minimum of 3 and a maximum of 24 threads
  JCR repository home is at repository
  JCR repository will be initialized.
  Task recovery file will be deleted
  Application import location is importablePackages
[yes]:

```

Answer **yes** (or press Enter) to store the configuration settings and end the Setup Wizard. If you selected the option to initialize the repository, this will be done now.

Answer **no** to abort the Setup Wizard.

If the Setup Wizard is successfully completed, it will display the following message:

```

You can now start your Deployit server by executing the command server.cmd
on Windows or server.sh on Unix.
Note: If your Deployit server is running please restart it.
Finished setup.

```

Changing the Admin Password

By default, Deployit is installed with a special user with administrative permissions. This user has the username `admin` and password `admin`. As the last step in the installation, the admin password should be changed to something more secure. Issue the following commands in the CLI to do this:

```

adminUser = security.readUser('admin')
adminUser.password = 'newpassword'
security.modifyUser(adminUser)

# Test whether the change is successful
security.logout()
security.login('admin', 'newpassword')

```

High Availability Setup

Deployit can be configured to ensure maximum uptime of the application. In such a high availability setup, two instances of Deployit are running in an *active/passive* configuration. At any one time, only one Deployit instance is active but as soon as a failure is detected, the passive Deployit instance is activated and the failed instance is taken down for repair.

The easiest way to achieve such a configuration is by using the same configuration for each Deployit instance.

Warning: unpredictable results may occur when running two Deployit instances against the same repository *at the same time*.

When switching from one Deployit instance to another, any running tasks will be recovered by the new instance (see the **Deployit Reference Manual** for more information).

To configure such a setup, a router with *active/passive* support must be used.

Upgrading Deployit

To begin upgrading Deployit, first unpack the distribution archive. The distribution archive contains the following:

- Release notes describing the changes made in this version of Deployit.
- A server archive.
- A CLI archive.

Upgrading the Server

To upgrade an existing Deployit server installation, do the following:

1. Create a directory for the new Deployit server installation, including the new Deployit server version number in the directory name.
2. Extract the server archive in this directory.
3. Copy the contents of the **conf**, **repository** and **hotfix** directories from the previous installation into the new installation directory.
4. Copy the contents of the **importablePackages** directory from the previous installation into the new installation directory.
5. Copy the contents of the **plugins** directory from the previous installation into the new installation directory.
6. Copy the contents of the **ext** directory from the previous installation into the new installation directory.
7. If you have made any changes to the Deployit server startup scripts (*server.sh* or *server.cmd*), copy these from the **bin** directory in the previous installation into the new installation directory.

Note: please make sure that the plugins and extensions in the old Deployit installation are compatible with the new Deployit server version.

This completes upgrading of the Deployit server.

Upgrading the CLI

To upgrade an existing Deployit CLI installation, do the following:

1. Create a directory for the new Deployit CLI installation, including the new Deployit CLI version number in the directory name.
2. Extract the CLI archive in this directory.
3. Copy the contents of the **plugins**, **ext** and **hotfix** directories from the previous installation into the new installation directory.
4. If you have made any changes to the Deployit CLI startup scripts (*cli.sh* or *cli.cmd*), copy these from the **bin** directory in the previous installation into the new installation directory.

This completes upgrading of the Deployit CLI.

Configuring Deployit

This section contains information on the configuration of the Deployit server.

Configuring Security

Security in Deployit

Deployit supports a fine-grained access control scheme to ensure the security of your middleware and deployments. Deployit's security mechanism is based on the concepts of *principals* and *permissions*.

Principals

A (security) principal is an entity that can be authenticated and that can be assigned rights over resources in Deployit. Out of the box, Deployit supports only users as principals -- users are authenticated by means of a username and password and rights within Deployit are assigned to the user itself. When using an LDAP repository, groups in LDAP are also treated as principals. See below for more information about LDAP.

There is one special user, `admin`, who has special rights in Deployit. This user is allowed to grant and revoke security permissions.

Permissions

Permissions are rights to execute particular actions in Deployit. Some permissions also imply the rights to make modifications to (certain parts of) the repository. Permissions can also be restricted to a subset of the resources in Deployit if fine-grained access control is required. See the **Reference Manual** for more information about the repository structure.

Deployit supports the following permissions:

- **read**. The right to see particular CIs.
- **login**. The right to log into the Deployit application. This permission does not allow the user has to access nodes in the repository.
- **import#initial**. The right to import a package for which the application does not yet exist in the repository and for which a new application will be created. This implies read and write access to the **Applications** node. The permission can also be given for specific *Application* CIs.
- **import#upgrade**. The right to import a package for which the application already exists in the repository. This implies read and write access to the **Applications** node. The permission can also be given for specific *Application* CIs.
- **deploy#initial**. The right to perform an initial deployment of a package to an environment. This implies read and write access to the **Environments** and **Infrastructure** nodes. The permission can also be given for specific *Environment* CIs.
- **deploy#upgrade**. The right to perform an upgrade of a deployment on an environment. Note that this does *not* allow deploying items from your package to new targets. This implies read and write access to the **Environments** and **Infrastructure** nodes. The permission can also be given for specific *Environment* CIs.
- **discovery**. The right to perform discovery. This implies read and write access to the **Infrastructure** node.
- **repo#edit**. The right to edit (create and modify) CIs in the repository. This implies write access to the **Applications**, **Environments** and **Infrastructure** nodes. The user must also have read access to CIs to be able to edit them. The permission can also be given for specific CIs.
- **deploy#undeploy**. The right to undeploy an application. This implies read and write access to the **Environments** and **Infrastructure** nodes. The permission can also be given for specific *Environment* CIs.
- **task#skip_step**. The right to skip steps in the generated steplist before starting a deployment.
- **task#move_step**. The right to move steps in the generated steplist before starting a deployment.
- **controltask#execute**. The right to execute control tasks on configuration items.

Granting, Revoking and Denying

Access rights in Deployit can be *granted* to a principal or *revoked* from a principal. When rights are granted, the principal is allowed to perform some action or access repository entities. Rights once granted can be revoked again to prevent the action in the future. Rights can also be explicitly *denied*. Denying permission acts as a negative grant -- the right is explicitly disallowed.

Access rights can be used stand-alone or in combination with one or more CIs. In the former case, the principal will have access to **all** CIs associated with the permission. In the latter case, the access rights will be restricted to the particular CIs. For example, granting *import#initial* to a principal allows the principal to import any application. Granting *import#initial* on *Applications/PetClinic* allows the principal to import only *PetClinic* packages.

Configuring Repository Security

Security in the Deployit repository can be configured using the command line interface. See the **Deployit Command Line Interface (CLI) Manual** for more information.

Creating Users

Deployit can only create users in it's own repository, even if it is configured to use an LDAP repository for authentication and authorization. To do this, use a statement such as the following:

```
deployer = security.createUser("john", "secret")
```

Granting Permissions

To grant a particular permission to a principal, use a statement such as the following:

```
security.grant("import#initial", "john")
```

To grant a particular permission to a principal on a CI, use a statement such as the following:

```
security.grant("read", "mary", ['Environments/Dev'])
```

Revoking Permissions and Privileges

To revoke a particular permission from a principal, use a statement such as the following:

```
security.revoke("read", "john", ['Environments/Dev'])
```

Inherited Permissions

Permissions in Deployit are *inherited* for all CIs that are contained in the CI node you specified the permissions for. For example, if you have read permission on *Environments/Dev*, you will also have read permission on all of the CIs under this node such as deployed applications. This is what happens when you use stand-alone permissions. If you use permissions on specific CIs, these permissions are set only on the specified CIs and are inherited from this CI onwards.

Permissions on Deployed Applications

Giving users permissions on deployed applications requires some knowledge of how these CIs are stored in the repository. As described in the **Deployit Reference Manual**, deployed applications are stored under both the **/Environments** as well as **/Infrastructure** nodes. Giving users read permission on a deployed application involves giving them permissions under both nodes.

Security Configuration Example

Let's illustrate the security setup with an example.

In a typical medium to large size company, there several different groups of people that perform tasks related to deployments. There are administrators that install, test and maintain hardware, there are deployers that deploy applications to development, test, acceptance and production environments. And finally there are the developers who build these applications.

Translating this into Deployit terms:

- **administrators:** permission to create, update and delete infrastructure as well as permission to create, update and delete environments. (all handled by *repo#edit* permission)
- **deployers:** permission to import new applications (*import#initial*), deploy to DEV, TEST and view PROD (*deploy#initial* and *deploy#upgrade* on environments DEV and TEST, *read* rights on environment PROD).
- **senior deployers:** permission to import new applications (*import#initial*), deploy to DEV, TEST and PROD (*deploy#initial* and *deploy#upgrade* on environments DEV, TEST and PROD).
- **developers:** permission to import new versions of existing applications (*import#upgrade*) and to upgrade existing deployments (*deploy#upgrade*).

See the **Deployit Command Line Interface (CLI) Manual** for the exact commands to implement this.

Configuring LDAP Security

By default, Deployit authenticates users and retrieves authorization information from its repository. Deployit can also be configured to use an LDAP repository to authenticate users and to retrieve role (group) membership. In this scenario, the LDAP users and groups are used as principals in Deployit. Rights can be assigned to both users and groups. The rights assigned to a principal are always stored in the JCR repository.

Deployit treats the LDAP repository as **read-only**. This means that Deployit will use the information from the LDAP repository, but can not make changes to that information.

When authenticating a user, Deployit first tries to locate the user in the LDAP repository. If this fails, Deployit will check its own repository as a backup.

To configure Deployit to use an LDAP repository, the built-in JCR repository, Jackrabbit, must defer to the LDAP server for authentication. This requires modification of the default Deployit Jackrabbit configuration. Follow these steps:

1. Edit `conf/jackrabbit-jaas.config` file.

The file must be accessible for the Deployit server. The sample file can be adapted to suit your needs. The `LdapLoginModule` takes the following arguments:

- **userProvider:** The LDAP url to connect to, including the bind path (example: `userProvider="ldap://localhost:389/ou=system"`)
- **userFilter:** The LDAP filter to determine the LDAP dn for the user that's logging in, `{USERNAME}` will be replaced with the username that's logging in (example: `userFilter="(&(uid={USERNAME})(objectClass=inetOrgPerson))"`)
- **useSSL:** Whether or not we need to use SSL to connect to LDAP (example: `useSSL=false`)
- **groupFilter:** The LDAP filter to determine the group memberships for the user, `{DN}` will be replaced with the DN of the user (example: `groupFilter="(member={DN})"`)
- **principalProvider:** The LDAP PrincipalProvider that will provide principals to Jackrabbit (example: `principalProvider=com.xebialabs.deployit.security.LdapPrincipalProvider`)
- **java.naming.security.authentication:** The security scheme to use to connect to LDAP (example: `java.naming.security.authentication=simple`)
- **java.naming.security.principal:** The DN of the user that you need to do the initial connect to search LDAP (example: `java.naming.security.principal="cn=admin,ou=users,ou=system"`)
- **java.naming.security.credentials:** The password matching the DN needed to connect to LDAP (example: `java.naming.security.credentials=secret`)

2. Modify the Deployit server startup command.

Notify the Deployit server of the new configuration file by specifying the following JVM variable in the startup script:

```
-Djava.security.auth.login.config="$DEPLOYIT_SERVER_HOME/conf/jackrabbit-jaas.config"
```

3. Edit the `conf/jackrabbit-repository.xml`

Change the `jackrabbit-repository.xml` by looking up the `Security`-block, and changing it to look like this:

```
<Security appName="Jackrabbit">
  <SecurityManager class="org.apache.jackrabbit.core.DefaultSecurityManager"
    workspaceName="security" />
  <AccessManager class="org.apache.jackrabbit.core.security.DefaultAccessManager" />
</Security>
```

Configuring the Repository

Using a Database

Deployit can also use a database to store its repository. The built-in Jackrabbit JCR implementation must be configured to make this possible. The two relevant components are Jackrabbit's `PersistenceManager` (used for storing nodes and revisions) and `DataStore` (optionally used for storing large binary objects). In the default Deployit configuration, the [Derby database](#) and file system `DataStore` are used.

Here are some examples of configuring Deployit to use a database for various database vendors. The XML snippets below must be put into the `conf/jackrabbit-repository.xml` file.

Note: Deployit **must** initialize the repository before it can be used. Run Deployit's setup wizard after making any changes to the Jackrabbit configuration.

For more information about using a database with Jackrabbit, see its [PersistenceManager FAQ](#) and [DataStore FAQ](#).

Using Deployit with **MySQL**

```
<Workspace name="${wsp.name}">
  <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
    <param name="path" value="${wsp.home}" />
  </FileSystem>

  <PersistenceManager class="org.apache.jackrabbit.core.persistence.pool.MySqlPersistenceManager">
    <param name="url" value="jdbc:mysql://localhost:3306/deployit" />
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="schemaObjectPrefix" value="${wsp.name}_" />
  </PersistenceManager>

  <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
    <param name="path" value="${wsp.home}/index" />
    <param name="supportHighlighting" value="true" />
  </SearchIndex>
</Workspace>

<Versioning rootPath="${rep.home}/version">
  <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
    <param name="path" value="${rep.home}/version" />
  </FileSystem>

  <PersistenceManager class="org.apache.jackrabbit.core.persistence.pool.MySqlPersistenceManager">
    <param name="url" value="jdbc:mysql://localhost:3306/deployit" />
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="schemaObjectPrefix" value="version_" />
  </PersistenceManager>
</Versioning>
```

Note: The MySQL database is not suited for storage of large binary objects, see [the MySQL bug tracker](#).

Using Deployit with **DB2**

```

<Workspace name="${wsp.name}">
  <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
    <param name="path" value="${wsp.home}" />
  </FileSystem>

  <PersistenceManager class="org.apache.jackrabbit.core.persistence.pool.BundleDbPersistenceManager">
    <param name="driver" value="com.ibm.db2.jcc.DB2Driver"/>
    <param name="url" value="jdbc:db2://localhost:50002/deployit" />
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="databaseType" value="db2" />
    <param name="schemaObjectPrefix" value="${wsp.name}_" />
  </PersistenceManager>

  <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
    <param name="path" value="${wsp.home}/index" />
    <param name="supportHighlighting" value="true" />
  </SearchIndex>
</Workspace>

<Versioning rootPath="${rep.home}/version">
  <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
    <param name="path" value="${rep.home}/version" />
  </FileSystem>

  <PersistenceManager class="org.apache.jackrabbit.core.persistence.pool.BundleDbPersistenceManager">
    <param name="driver" value="com.ibm.db2.jcc.DB2Driver"/>
    <param name="url" value="jdbc:db2://localhost:50002/deployit" />
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="databaseType" value="db2" />
    <param name="schemaObjectPrefix" value="version_" />
  </PersistenceManager>
</Versioning>

```

Using Deployit with Oracle

```

<Workspace name="${wsp.name}">
  <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
    <param name="path" value="${wsp.home}" />
  </FileSystem>

  <PersistenceManager
    class="org.apache.jackrabbit.core.persistence.bundle.OraclePersistenceManager">
    <param name="driver" value="oracle.jdbc.driver.OracleDriver"/>
    <param name="url" value="jdbc:oracle:thin:@localhost:1521:orcl"/>
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="databaseType" value="oracle" />
    <param name="schemaObjectPrefix" value="${wsp.name}_" />
  </PersistenceManager>

  <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
    <param name="path" value="${wsp.home}/index" />
    <param name="supportHighlighting" value="true" />
  </SearchIndex>
</Workspace>

<Versioning rootPath="${rep.home}/version">
  <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
    <param name="path" value="${rep.home}/version" />
  </FileSystem>

  <PersistenceManager
    class="org.apache.jackrabbit.core.persistence.bundle.OraclePersistenceManager">
    <param name="driver" value="oracle.jdbc.driver.OracleDriver"/>
    <param name="url" value="jdbc:oracle:thin:@localhost:1521:orcl"/>
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="databaseType" value="oracle" />
    <param name="schemaObjectPrefix" value="version_" />
  </PersistenceManager>
</Versioning>

```

If you use the TNSNames Alias syntax to connect to Oracle, you may need to inform the driver where to find the TNSNAMES file. See the Oracle documentation for more details.

Clustering Jackrabbit

It is also possible to run Deployit server with its repository shared with other Deployit server instances. For this to happen, the jackrabbit JCR must be configured to run in a [clustered mode](#). This needs a cluster configuration to be present in the jackrabbit-repository.xml file.

The following XML snippet shows a sample clustering configuration for a JCR using Oracle as its repository.

```
<Cluster id="101" syncDelay="2000">
  <Journal class="org.apache.jackrabbit.core.journal.OracleDatabaseJournal">
    <param name="revision" value="{rep.home}/revision" />
    <param name="driver" value="oracle.jdbc.driver.OracleDriver" />
    <param name="url" value="jdbc:oracle:thin:@localhost:1521:orcl" />
    <param name="user" value="deployit" />
    <param name="password" value="deployit" />
    <param name="schemaObjectPrefix" value="JOURNAL_" />
  </Journal>
</Cluster>
```

Note that each jackrabbit cluster node should have a unique value for `id`. For more information on JCR clustering, or ways to configure clustering using other databases, please refer to the jackrabbit [clustering documentation](#).

Installing Plugins

Deployit Server supports various plugins that add functionality to the system. When it starts, the Deployit server scans the `plugins` directory and loads all plugins it finds. The additional functionality they provide is immediately available. Any plugins added or removed when Deployit server is running will not take effect until the server is restarted.

Installing a Plugin

To install a new plugin, stop the Deployit server and copy the plugin JAR archive into the `plugins` directory, then restart the Deployit server.

Uninstalling a Plugin

To uninstall a plugin, stop the Deployit server and remove the plugin JAR archive from the `plugins` directory, then restart the Deployit server.

Configuring Logging

Out of the box, Deployit server writes informational, warning and error log messages to standard output as well as `log/deployit.log` when running. It is possible to change this behavior to write log output to a file or to log output from a specific source.

Deployit uses the Logback logging framework for it's logging. To change it's behavior, edit the file `logback.xml` in the `conf` directory of the Deployit server installation directory.

The following is an example `logback.xml` file:

```
<configuration>
  <appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
    <!-- encoders are assigned the type
         ch.qos.logback.classic.encoder.PatternLayoutEncoder by default -->
    <encoder>
      <pattern>
        %d{yyyy-MM-dd HH:mm:ss.SSS} [%thread] %-5level %logger{36} - %msg%n
      </pattern>
    </encoder>
```



```

</appender>

<!-- Create a file appender that writes log messages to a file -->
<appender name="FILE" class="ch.qos.logback.core.FileAppender">
  <layout class="ch.qos.logback.classic.PatternLayout">
    <pattern>%-4relative [%thread] %-5level %class - %msg%n</pattern>
  </layout>
  <File>log/my.log</File>
</appender>

<!-- Set logging of classes in com.xebialabs to DEBUG level -->
<logger name="com.xebialabs" level="debug"/>

<!-- Set logging of class HttpClient to DEBUG level -->
<logger name="HttpClient" level="debug"/>

<!-- Set the logging of all other classes to INFO -->
<root level="info">
  <!-- Write logging to STDOUT and FILE appenders -->
  <appender-ref ref="STDOUT" />
  <appender-ref ref="FILE" />
</root>

</configuration>

```

For more information see the [Logback website](#).

Setting up Deployit

This section describes how to setup Deployit server in your environment.

Deployit must be setup for your environment before it can be used to execute deployments. This entails the following steps:

1. **Start the Deployit server.**
2. **Discover your middleware.** Deployit can inspect your environment and automatically create CIs for your middleware. Alternatively, you can use a bulk-import to import your middleware or create them by hand using the Deployit GUI.
3. **Add the discovered middleware to an environment.** CIs must be grouped in an environment to use them for deployment.

Setup of Deployit is performed using the Deployit Command Line Interface (CLI). For more information about the CLI, see the **Deployit Command Line Interface (CLI) Manual**.

Starting and Stopping

Starting the Server

Open a terminal window and change to the DEPLOYIT_SERVER_HOME directory. Start the Deployit server with the command:

```
bin/server.sh
```

on Unix and

```
bin/server.cmd
```

on Windows.

By starting the server with the `-h` flag, a message is printed that shows the possible options it supports:

```
server.sh arguments...
  -reinitialize : Reinitialize the repository, only useful with -setup
  -setup       : (Re-)run the setup for Deployit
```

The command line options are:

- `-reinitialize` -- tells Deployit to reinitialize the repository. Used only in conjunction with `-setup`. **N.B.** This flag only works if Deployit is running on the filesystem repository, not when you've configured Deployit to run against a database.
- `-setup` -- runs the Deployit Setup Wizard.
- `-test-mode` -- enables Deployit server test-mode. When test-mode is enabled, WebDAV access to the JCR repository is installed. The repository can be accessed using a URL like `http://localhost:4516/repository/default/`.

Server Options

Any options you want to give the Deployit Server when it starts can be specified in the `DEPLOYIT_SERVER_OPTS` environment variable.

Starting Deployit in the Background

By running the `server.sh` or `server.cmd` command, the Deployit server is started in the foreground. To run the server as a background process, use:

```
nohup bin/server.sh &
```

on Unix or run Deployit as a [service on Windows](#).

Java Properties

Deployit server also responds to certain Java properties that influence its behavior. These properties may be set in the environment (e.g. by executing `export jetty.host=127.0.0.1` in the terminal used to start the server) or by passing them to Java on the command line (for instance, by adding the flag `-Djetty.host=127.0.0.1` to the command that starts the server). The following option is supported:

- `jetty.host` -- sets the host that Deployit's embedded HTTP server binds to. Default value is `localhost`.

Stopping the Server

It is possible to stop the Deployit server using a REST API call. The following is an example of a command to generate such a call (replace `admin:admin` with your own credentials):

```
curl -X POST --basic -u admin:admin
http://admin:admin@localhost:4516/deployit/server/shutdown
```

This requires the external `curl` command, available for both Unix and Windows.

Editing CIs

The CIs in the Deployit repository can also be edited, either using the command line interface (CLI) or graphical user interface (GUI). See the respective manuals for more details.

Maintaining Deployit

This section describes how to maintain the Deployit server in your environment.

Creating Backups

To create a backup of Deployit, several components may need to be backed up depending on your configuration:

- **JCR repository.**
 - Built-in repository: Create a backup of the built-in JCR repository by backing up the files in the repository directory.
 - Database repository: Create a backup of the database using the tools provided by your vendor.
- **Configuration.** Create a backup of the Deployit configuration by backing up the files in the `conf` directory in the installation directory.
- **Customization.** Create a backup of the Deployit customizations by backing up the files in the `ext` and `plugins` directory in the installation directory.

Note: Deployit **must not** be running when you are making a backup. Schedule backups outside planned deployment hours to ensure the server is not being used.

Restoring Backups

To restore a backup of Deployit, restore one of the following components:

- **JCR repository.**
 - Built-in repository: Remove the repository directory and replace it with the backup.
 - Database repository: Restore a backup of the database using the tools provided by your vendor.
- **Configuration.** Remove the `conf` directory in the `DEPLOYIT_SERVER_HOME` directory and replace it with the backup.
- **Customization.** Remove the `ext` and `plugins` directories in the `DEPLOYIT_SERVER_HOME` directory and replace them with the backups.

Note: Deployit **must not** be running when you are restoring a backup.