



Set Up DI Server and Tools

Copyright Page

This document supports Pentaho Business Analytics Suite 5.4 GA and Pentaho Data Integration 5.4 GA, documentation revision June 9th, 2015, copyright © 2015 Pentaho Corporation. No part may be reprinted without written permission from Pentaho Corporation. All trademarks are the property of their respective owners.

Help and Support Resources

To view the most up-to-date help content, visit <https://help.pentaho.com>.

If you do not find answers to your questions here, please contact your Pentaho technical support representative.

Support-related questions should be submitted through the Pentaho Customer Support Portal at <http://support.pentaho.com>.

For information about how to purchase support or enable an additional named support contact, please contact your sales representative, or send an email to sales@pentaho.com.

For information about instructor-led training, visit <http://www.pentaho.com/training>.

Liability Limits and Warranty Disclaimer

The author(s) of this document have used their best efforts in preparing the content and the programs contained in it. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, express or implied, with regard to these programs or the documentation contained in this book.

The author(s) and Pentaho shall not be liable in the event of incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of the programs, associated instructions, and/or claims.

Trademarks

The trademarks, logos, and service marks ("Marks") displayed on this website are the property of Pentaho Corporation or third party owners of such Marks. You are not permitted to use, copy, or imitate the Mark, in whole or in part, without the prior written consent of Pentaho Corporation or such third party. Trademarks of Pentaho Corporation include, but are not limited, to "Pentaho", its products, services and the Pentaho logo.

Trademarked names may appear throughout this website. Rather than list the names and entities that own the trademarks or inserting a trademark symbol with each mention of the trademarked name, Pentaho Corporation states that it is using the names for editorial purposes only and to the benefit of the trademark owner, with no intention of infringing upon that trademark.

Third-Party Open Source Software

For a listing of open source software used by each Pentaho component, navigate to the folder that contains the Pentaho component. Within that folder, locate a folder named licenses. The licenses folder contains HTML files that list the names of open source software, their licenses, and required attributions.

Contact Us

Global Headquarters Pentaho Corporation Citadel International, Suite 460

5950 Hazeltine National Drive Orlando, FL 32822

Phone: +1 407 812-OPEN (6736)

Fax: +1 407 517-4575

<http://www.pentaho.com>

Sales Inquiries: sales@pentaho.com

Introduction

In this section, you do the minimum configuration tasks for the server, design tool, and utilities, so you can get started creating ETL solutions.

Prerequisites

Before you begin, you must have [installed](#) Pentaho Data Integration software. If you chose to install the Pentaho Business Analytics software you must go through [a different configuration process](#).

Expertise

The topics in this section are written for IT administrators who know where data is stored, how to connect to it, details about the computing environment, and how to use the command line to issue commands for Microsoft Windows, Linux, or Microsoft OS.


Tools

We provide a design tool, Spoon, that you use to perform most configuration tasks.

Login Credentials

All of the configuration tasks that use Spoon require that you [login to Spoon](#) using a Pentaho administrator user name and password.


Configure the DI Server



Before you can use Pentaho Data Integration (PDI) to extract data, transform it into a useful format, and load it for analysis, you need to perform configuration tasks. Just follow the *guideposts* to make sure that you complete the entire process. It should take you approximately one hour or less.

- [Start and Stop the Data Integration Server](#)
- [Increase the DI Server Memory Limit](#)
- [Access Spoon](#)
- [Define Security for the DI Server](#)
- [Next Steps](#)

Start and Stop the Data Integration Server



To work with Pentaho Data Integration you must have the DI server running. There may be times when you need to manage Pentaho Data Integration components, stopping the DI server to perform system maintenance and restarting it after you are done. As part of the Installation Wizard and Archive Installation, Pentaho provides different ways to help you start and stop the components depending on the operating system you use and the method you used for installation.

For Microsoft Windows, these are your options for starting and stopping the components.

- [Windows Pentaho Installation Wizard](#)
- [Windows Pentaho Archive Installation](#)

For Linux and Macintosh OS, these are your options for starting and stopping the components.

- [Linux Pentaho Installation Wizard](#)
- [Linux Pentaho Archive Installation](#)
- [Linux Pentaho Wizard or Archive Installation—Starting on Boot](#)

When you are done, please go to the next stop on the *guidepost*.

Windows Pentaho Installation Wizard

When you ran the Installation Wizard on Windows, the BA Server installed in a Tomcat Application server. You can manage the Tomcat and DI servers by clicking **Start > Pentaho Enterprise Edition > Server Management** and then selecting from these menu items.

- Start the DI Server
- Stop the DI Server

The wizard also registered the DI Server and the PostgreSQL solution repository as services. These services are set to run automatically, enabling them to start and stop when the computer running them boots or shuts down. You can use the Windows Services applet found in the Control Panel to start and stop the Pentaho servers and the solution repository.

1. Click **Start > Control Panel > Administrative Tools > Services**.
2. In the **Services** window, right-click one of these services in the list and choose **Start** or **Stop** from the menu that appears.

- Data Integration
- Pentaho Solution Repository

Windows Pentaho Archive Installation

If you used the Archive Installation, we provide individual control scripts to start and stop the DI server, and DI repository. Here is where you can find the individual control scripts.

DI Repository

- The Archive Installation enables you to install PostgreSQL, MySQL, or Oracle as the solution repository. Consult the third-party documentation for the RDBMS to find more information about starting and stopping.

The solution repository must be started before the DI Server.

DI Server

- `/pentaho/server/data-integration-server`

Linux and Macintosh OS Pentaho Installation Wizard

When you ran the Installation Wizard on Linux, the DI Server was deployed in an included Apache Tomcat application server. You can control the Tomcat server using the start and stop scripts that come with the Pentaho installation. This script is also used as an easy way to start and stop the DI Server and the PostgreSQL repository. You can find this script at `/pentaho/ctlscript.sh`.

Here is a list of the script arguments you can use with the `data-integration-server` service.

Arguments

- start
- stop
- restart
- status
- help

```
./ctlscript.sh start data-integration-server
```

```
./ctlscript.sh status data-integration-server
```

```
./ctlscript.sh help
```

Linux and Macintosh OS Pentaho Archive Installation

If you used the Archive Installation, Pentaho provides individual control scripts to start and stop the Pentaho Tomcat application server, the DI Server, and the PostgreSQL server. Here is where you can find the individual control scripts.

DI Repository

- The Archive Installation enables you to install either PostgreSQL, MySQL, or Oracle as the solution repository. Consult the documentation for the RDBMS you selected for information about starting and stopping.

The solution repository must be started before the DI Server.

DI Server

- `pentaho/server/data-integration-server/start-pentaho.sh` and `stop-pentaho.sh`

Linux and Macintosh OS Pentaho Wizard or Archive Installation—Starting on Boot

Because the Installation Wizard and the Archive Installation do not provide a way to start the DI repository and DI Server automatically on boot or shut-down, here are examples of how you might approach creating a script to do this.

This procedure assumes that you are running the DI Server under the `pentaho` local user account. If you are using a different account to start these services, substitute it in the script in step 2. This script also assumes you are using the PostgreSQL solution repository. Where `postgres` appears in this script, change it to reflect the RDBMS you are using as a solution repository, either MySQL or Oracle. This script was tested on Red Hat Enterprise Linux. You may have to modify the details of the script if you use a different distribution of Linux or other Unix-like operating system, different shells, or different `init` systems.

1. With root permissions, create a file in `/etc/init.d/` named `pdi`.
2. Using a text editor, copy the following content into the new `pentaho` script, changing `postgres` to the name of the `init` script for your database, if it is running on the remote machine, or remove `postgres` entirely if you are using a remote database. You may also have to adjust the paths to the DI Server scripts to match your situation.

```
#!/bin/sh
### BEGIN INIT INFO
# Provides: start-data-integration-server stop-data-integration-server
# Required-Start: networking postgresql
# Required-Stop: postgresql
# Default-Start: 2 3 4 5
# Default-Stop: 0 1 6
# Description: Pentaho DI Server
### END INIT INFO

case "$1" in
"start")
su - pentaho -c "/home/pentaho/pentaho/server/data-integration-server/start-pentaho.sh"
;;
"stop")
su - pentaho -c "/home/pentaho/pentaho/server/data-integration-server/stop-
```



```
pentaho.sh"
;;
*)
echo "Usage: $0 { start | stop }"
;;
esac
exit 0
```

3. Save the file and close the text editor.
4. Open `/home/pentaho/pentaho/server/data-integration-server/start-pentaho.sh` with a text editor.
5. Change the last `if` statement to match the this example.

```
if [ "$?" = 0 ]; then
    cd "$DIR/tomcat/bin"
    export CATALINA_OPTS="-Xms2048m -Xmx2048m -XX:MaxPermSize=2048m -Dsun.rmi.
dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000"
    env JAVA_HOME=$_PENTAHO_JAVA_HOME sh ./startup.sh
fi
```


6. Save the file and close the text editor.
7. Make the `init` script executable.

```
chmod +x /etc/init.d/pentaho
```

8. Add the `pentaho` `init` script to the standard run levels by using the `update-rc.d` command, so that it runs when the system starts, and stops when the system is shut down or rebooted. This command may not exist on your computer if it is not Debian-based. If that is the case, consult your distribution documentation or contact your distribution's support department to determine how to add `init` scripts to the default run levels.

```
update-rc.d pentaho defaults
```

Increase the DI Server Memory Limit



We recommend that you increase PDI's memory limit so the DI Server and Data Integration Design Tool (Spoon) can perform memory-intensive tasks, like sort or group large datasets or run complex transformations and jobs. You will need to increase the memory limit for both the DI Server and Spoon. If you choose to not increase the memory limit, PDI uses the default memory settings in the PDI startup scripts.

The way that you increase the memory settings for the DI Server depends on *how you installed* PDI and *what operating system* you are running it on.

- [Increase DI Server Memory Limit if You Used the Wizard to Install PDI on Windows](#)
- [Increase DI Server Memory Limit if You Used the Wizard to Install PDI on Linux](#)
- [Increase DI Server Memory Limit if Used a Custom Installation Method to Install PDI on Windows or Linux](#)

Increase DI Server Memory Limit for Wizard Installations on Windows

If you used the Installation Wizard to install the DI Server on a Windows machine, you can increase the DI Server's memory limits by editing the Java memory settings for Tomcat. (Tomcat, which is installed with the installation wizard, is the web application server that the DI Server runs on.) If you didn't use the Wizard or you are not running PDI on a Windows machine, check out the [other ways to increase memory](#).

1. [Stop the Pentaho DI Server](#) if it is running.
2. Type `services.msc` into the **Windows Search Box**.
3. Find the Pentaho Server name (**Data Integration Server**) and open it so you can find the **service name**. The **service name** should appear at the top of the first tab (**General**). It will be **pentahoDataIntegrationServer**.
4. Go into the bin file (`C:\Program Files\pentaho\server\data-integration-server\tomcat\bin\`) and rename the **tomcat6w.exe** file to match the **service name** (**pentahoDataIntegrationServerw.exe**). This will ensure that the server starts with the software.
5. After you have renamed the file, open it by double-clicking on it. This will not open the file, it will allow you to edit it. You may need to right-click and select **Run as Administrator**. This depends on your user permission settings. The **Properties Window** will open.

6. Select the **Java** tab.
7. Set the memory setting to a minimum of **4096 M** and a maximum of **6144 M**, depending on your computer's memory capabilities.
8. [Start the Tomcat server or service](#). **Note:** Make sure to also [increase the Spoon memory limit](#).

Your Tomcat server now has increased minimum and maximum memory limits. You can adjust the `JvmMx` parameter, which is a parameter that specifies the maximum limit, to a higher number if you prefer. However, if the Java virtual machine refuses to start with increased limits, then you will have to add more RAM to your system, stop some memory-intensive services, or reduce the maximum memory limit to a lower number. This problem occurs when there is not enough contiguous memory available to assign to the JVM.

Increase DI Server Memory Limit for Wizard Installations on Linux

If you used the Wizard to install PDI on your Linux machine, you can increase memory limits by editing a variable in one of the Pentaho-supplied scripts. If you didn't use the Wizard or you are not running PDI on a Linux machine, see the [other ways to increase memory](#).

1. Go to `/pentaho/server/data-integration-server/tomcat/bin/` directory and run the `./shutdown.sh` command to stop the appropriate server.
2. Change the directory to `data-integration-server/tomcat/scripts`.
3. Edit the `ctl.sh` file.
4. Locate the line under **start tomcat**, which looks like this: `export JAVA_OPTS="-Dpentaho.installed.licenses.file=/opt/pentaho/.installedLicenses.xml -Xms128m Xmx768m -XX-MaxPermSize=256m -Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000"`
5. Set the memory to a minimum of **4096 M** and a maximum of **6144 M**, depending on your computer's memory capabilities.
6. Start the Tomcat server or service. **Note:** Make sure to also [increase the Spoon memory limit](#).

Your Tomcat server now has increased minimum and maximum memory limits. You can adjust the `JvmMx` parameter to specify a higher maximum limit if you prefer. However, if the Java virtual machine refuses to start with increased limits, then you will have to add more RAM to your system, stop some memory-intensive services, or reduce the maximum memory limit to a lower number. This problem occurs when there is not enough contiguous memory available to assign to the JVM.


Increase DI Server Memory Limit for Custom Installations on Windows or Linux

If you used a custom method to install PDI on your Windows or Linux machines, you can increase memory limits by editing a text file. If you used the Wizard to install PDI, then check out the other the [other ways to increase memory](#).

- [Install with Your Own DI Repository: Modify Tomcat Startup Script](#)
- [DI Server Manual Installation: Modify Tomcat and JBoss Startup Scripts](#)

Note: Make sure to also [increase the Spoon memory limit](#).

Access Spoon



Spoon is a desktop application that you will use primarily as a graphical interface and editor for transformations and jobs. With Spoon you can author, edit, run, and debug transformations and jobs. You can also use Spoon to enter license keys, add data connections, and define security.

- [Start Spoon](#)
- [Enter Licenses](#)
- [Connect to the DI Repository](#)
- [Tour Spoon](#)
- [Change Administrator Password](#)

When you are done, please go on to the next stop on the *Guide Post* graphic.

Start Spoon

There are a few different ways to start Spoon. The method that you should use depends on the way you installed Pentaho Data Integration (PDI).

Start Spoon for Windows Pentaho Graphical Installation

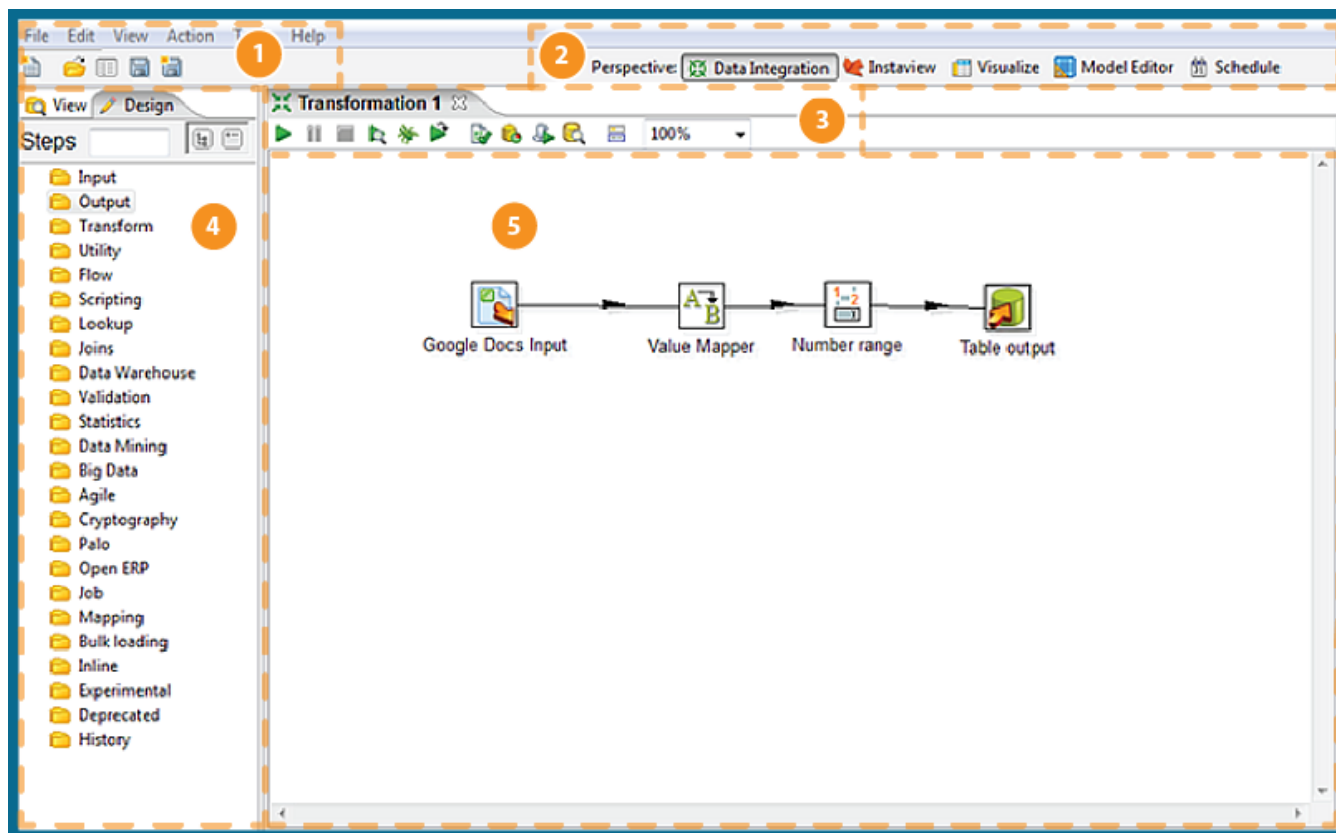
1. [Start the DI server.](#)
2. From the Windows Start menu, select **All Programs > Pentaho Enterprise Edition > Design Tools > Data Integration**. The PDI Design Tool, Spoon, appears.

Start Spoon for Archive or Manual Installation

1. [Start the DI server.](#)
2. Navigate to the folder where you have installed PDI. For example `...\pentaho\design-tools\data-integration`.
3. Launch Spoon in the best way for your operating system.
 - a. **For Windows:** double-click **Spoon.bat**.
 - b. **For Linux:** double-click **spoon.sh**
 - c. **For Macintosh:** go to `.../pdi-ee/data-integration` and double click on the **Data Integration 32-bit** or **Data Integration 64-bit** icon depending on your system.

The **Pentaho License Manager** dialog appears. You must [enter licenses](#).

Tour Spoon



Component Name	Name	Function
1	Toolbar	Single-click access to common actions such as create a new file, opening existing documents, save and save as.
2	Perspectives Toolbar	<p>Switch between the different perspectives.</p> <ul style="list-style-type: none"> • Data Integration — Create ETL transformations and jobs • Instaview — Use pre-made templates to create visualizations from PDI transformations • Visualize — Test reporting and OLAP metadata models created in the Model perspective using the Report Design Wizard and Analyzer clients • Model Editor — Design reporting and OLAP metadata models which can be tested right from within the Visualization perspective or




















Component Name	Name	Function
		<p>published to the Pentaho BA Server</p> <ul style="list-style-type: none"> • Schedule — Manage scheduled ETL activities on the Data Integration Server
	Sub-toolbar	Provides buttons for quick access to common actions specific to the transformation or job such as Run , Preview , and Debug .
	Design and View Tabs	<p>The Design tab of the Explore pane provides an organized list of transformation steps or job entries used to build transformations and jobs. Transformations are created by simply dragging transformation steps from the Design tab onto the canvas and connecting them with hops to describe the flow of data.</p> <p>The View tab of the Explore pane shows information for each job or transformation. This includes information such as available database connections and which steps and hops are used.</p> <p>In the image, the Design tab is selected.</p>
	Canvas	Main design area for building transformations and jobs describing the ETL activities you want to perform

Table 1. Spoon Icon Descriptions

Icon	Description
	Create a new job or transformation

Icon	Description
	Open transformation/job from file if you are not connected to a repository or from the repository if you are connected to one
	Explore the repository
	Save the transformation/job to a file or to the repository
	Save the transformation/job under a different name or file name (Save as)
	Run transformation/job; runs the current transformation from XML file or repository
	Pause transformation
	Stop transformation
	Preview transformation: runs the current transformation from memory. You can preview the rows that are produced by selected steps.
	Run the transformation in debug mode; allows you to troubleshoot execution errors
	Replay the processing of a transformation
	Verify transformation
	Run an impact analysis on the database
	Generate the SQL that is needed to run the loaded transformation.
	Launch the database explorer allowing you to preview data, run SQL queries, generate DDL and more
	Hide execution results pane
	Lock transformation

Enter Licenses

You must [start Spoon](#) before adding licenses.

1. Copy your license files to a convenient location.
2. Log into Spoon. The **Pentaho License Manager** dialog appears.

3. From within the **Pentaho License Manager**, click on the **Add** button.
4. Double-click on the license key to open it. A green check appears in the Status column to show that the license key installed correctly.
5. Click **Close** to close the dialog box.
6. Close any windows that appear and restart Spoon so that the license key will take effect.

Alternatively, you can [install licenses using the command line interface](#).

Connect to the DI Repository

The DI Repository provides you a place to centrally store ETL jobs and transformations. You may choose to set up common connections for your users now or you can close the dialog box and allow your users to connect to the DI Repository themselves.

1. Start Spoon.
2. Select **Tools > Repository > Connect**.
3. In the **Repository Connection** dialog box, click the add button.
4. Select **DI Repository: DI Repository** and click **OK**. The **Repository Configuration** dialog box appears.
5. Keep the default URL, which is <http://localhost:9080/pentaho-di>.
6. Click **Test** to ensure your connection is properly configured. If you get an error, make sure you have [started the DI Server](#).
7. Click **OK** to exit the **Success** dialog box.
8. Enter an **ID** and **Name** for your repository.
9. Click **OK** to exit the Repository Configuration dialog box. Your new connection appears in the list of available repositories.
10. Log on to the repository by entering the following credentials: user name = **admin**, password = **password**.


NOTE:

If you want the **Repository Connection** window to automatically appear when Spoon starts, go to **Tools > Options** and click **Show repository dialog at startup**.

Change the Administrator Password

After you have [logged into Spoon](#) for the first time, it is a best practice to [change the default administrator password](#).

Define Security for the DI Server



Using security is a best practice, but is not required. If you want to get started quickly or do not have information about your user community, skip this for now and go on to the next stop on the *Guide Post* graphic. You can always come back to it later.

We support two different security options: Pentaho Security or advanced security providers, such as LDAP, Single Sign-On, or Microsoft Active Directory. This table can help you choose the option that is best for you. If you are using an advanced security provider, see [Implementing Advanced Security for the DI Server](#) before continuing the configuration process.

Table 1. Security Decision Table

Explore Considerations	Choose Options	
	Pentaho Security	Advanced Security Providers—LDAP, Single Sign-On, or Microsoft Active Directory
Summary	<p>Pentaho Security is the easiest way to configure security quickly. Spoon enables you to define and manage users and roles. The DI Server controls which users and roles can access resources in the DI repository.</p> <p>Pentaho Security works well if you do not have a security provider or if you have a user community with less than 100 users.</p>	<p>If you are already using a security provider, such as LDAP, Single Sign-On, or Microsoft Active Directory, you can use the users and roles you have already defined with Pentaho. Your security provider controls which users and roles can access the DI repository.</p> <p>Advanced security scales well for production and enterprise user communities.</p>
Expertise	<p>Knowledge of your user community and which users should have which roles in the Pentaho system.</p> <p>Knowledge about security in general is <i>not</i> required.</p>	<p>Knowledge of your user community and which users should have which roles in the Pentaho system.</p> <p>Knowledge about your particular security provider and its options is required.</p>

Explore Considerations	Choose Options	
	Pentaho Security	Advanced Security Providers—LDAP, Single Sign-On, or Microsoft Active Directory
Time	It takes approximately 5 minutes per user and role to configure Pentaho Security.	It takes approximately 1 hour to configure the DI Server to use your existing security provider.
Recommendation	Recommended for the Pentaho Trial Download, evaluating, and rapid development.	Recommended for production.

- [Use Pentaho Security on the DI Server](#)

Use Pentaho Security on the DI Server

You must log into Spoon as an administrator (or be assigned to a role that has *Administer Security* permission) to manage users and roles for Pentaho Security. This section provides an overview of the out-of-box users and roles, along with the permissions that are included with each role.

Here is how you can manage users.

- [Add Users](#)
- [Change Passwords Using Spoon](#)
- [Delete Users](#)
- [Assign Users to Roles](#)
- [Edit User Information](#)

Here is how you can manage roles.

- [Add Roles](#)
- [Edit Roles](#)
- [Delete Roles](#)
- [Make Changes to the Admin Role](#)
- [Assign Permissions in the Repository Using Spoon](#)
- [Enable System Role Permissions](#)

Before changing security settings, play it safe and back up these relevant files.

- If you installed PDI using the wizard or custom methods, back up all Data Integration directories.
- If you installed PDI using the manual method, back up the `pentaho-di.war` file and solutions.

Control users and roles in the DI Repository with a point-and-click user interface. The users and roles radio buttons allow you to switch between user and role settings. You can add, delete, and edit users and roles from this page.

Default Users, Roles, and Permissions

Viewing default users and roles gives you an idea of how you can define your specific users and roles. To view the default users and roles, log into Spoon, click **Tools > Repository > Explore** and select the **Security** tab. Highlighting a user in the users list shows which roles are available for that user, as well as which role is currently defined for that user.

Out-of-Box Role	Out-of-Box User	Permissions
Administrator	admin	<ul style="list-style-type: none">• Administer Security• Read Content

Out-of-Box Role	Out-of-Box User	Permissions
		<ul style="list-style-type: none"> • Execute • Create content
Power User	suzy	<ul style="list-style-type: none"> • Read Content • Execute • Create Content
Report Author	tiffany	<ul style="list-style-type: none"> • None configured by default
Business Analyst	pat	<ul style="list-style-type: none"> • None configured by default


Each operation permission gives a specific set of permissions for Pentaho tools and the DI Server.

Table 1. Operation Permissions Defined

Operation Permission	Definition
Administer Security	<p>The default Administrator role automatically conveys all operation permissions to users assigned to that role, even if the check box next to it is cleared. This includes the Read and Create Content permissions, which are required for accessing the Administration perspective.</p> <ul style="list-style-type: none"> • Allows access to and the ability to manage all content in each perspective. • Allows the ability to view and work with all user schedules in the Schedules perspective.
Read Content	<ul style="list-style-type: none"> • Gives the user the ability to view content in each perspective.
Create Content	<ul style="list-style-type: none"> • Allows the user to create, import, delete, and save jobs and transformations to the repository. • Gives the user the ability to see the data sources that are used to create jobs and transformations. • When the user is also granted the Execute permission, users can export jobs and transformations, copy and paste, and save the file in a VFS.
Execute	<ul style="list-style-type: none"> • Allows the user to run, preview, debug, replay, verify, and schedule. • When the user is also granted the Create permission, users can export jobs and transformations, copy and paste, and save the file in a VFS.

Add Users

1. In Spoon, go to **Tools > Repository > Explore**. The Repository Explorer opens.
2. Click the **Security** tab.
Note: The **Users** radio button is selected by default.
3. Next to **Available**, click the round green plus button, **Add**. The **Add User** dialog box appears.
4. Type the **User Name** and **Password** associated with your new user account in the appropriate fields.
Note: An entry in the **Description** field is optional.
5. If you have available roles that can be assigned to the new user, under **Member**, select a role and click **OK**.



The role you assigned to the user appears in the right pane under **Assigned**.

6. Click **OK** to save your new user account and exit the Add Users dialog box.

The name of the user you added appears in the list of Available users.

Change Passwords Using Spoon

1. Launch Spoon as described in [Start Spoon](#).
2. Click on **Tools > Repository > Explore**.
3. Click on **Security**.
4. Select **Users**, **Roles**, or **System Roles** from the option button.
5. Select the role for which you want to change the password and click the **Edit** icon.
6. In the **Password** field, type the new password. Click **OK**.

Delete Users

You must be logged into the DI Repository as an administrative user.

1. In Spoon, go to **Tools > Repository > Explore**. The Repository Explorer opens.
2. Click the **Security** tab.
3. Select the user you want to delete from the list of available users.
4. Next to **Users**, click **Remove**. A confirmation message appears.
5. Click **Yes** to delete the user.

If a user or role is deleted in the DI Repository, content that refers to the deleted user, either by way of owning the content or having an ACL that mentions the user or role, is left unchanged. This makes it possible to create a new user or role using an identical name. In this scenario, content ownership and access control entries referring to the deleted user or role now apply to the new user or role.

To avoid this problem, we recommend that you disable a user or role instead of deleting it. This prevents a user or role with an identical name from ever being created again. Use these alternatives rather than deleting the user or role.

IF	THEN
You are dealing with a role	Unassign all current members associated with the role
You are dealing with a user	Reset the password to a password that is so cryptic that it is impossible to guess and is unknown to any users

Assign Users to Roles

You must be logged into the DI Repository as an administrative user.

You can assign users to roles and roles to users when you add a new user or role. You can also assign users to roles as a separate task.

1. In Spoon, go to **Tools > Repository > Explore**. The Repository Explorer opens.
2. Click the **Security** tab.
3. Click the **Roles** radio button. The list of available roles appear.
4. Select the role to which you want to assign one or more users.
Note: If the role has users currently assigned to it, the names of the users appear in the table on the right under **Members**. You can assign or unassign any users to a role. You can select a single item or multiple items from the list of members. Click **Remove** to remove the user assignment.
5. Next to **Members**, click **Add**. The **Add User to Role** dialog box appears.
6. Select the users you want assigned to the role and click **Add**. The users assigned to the role appear in the right pane.
7. Click **OK** to save your entries and to exit the Add User to Role dialog box.

The specified users are assigned to the specified role.

Edit User Information

You must be logged into the DI Repository as an administrative user.

1. In Spoon, go to **Tools > Repository > Explore**. The Repository Explorer opens.
2. Click the **Security** tab.
Note: The **Users** radio button is selected by default.
3. Select the user whose details you want to edit from the list of available users.
4. Click **Edit**. The **Edit User** dialog box appears.
5. Make the appropriate changes to the user information.
6. Click **OK** to save changes and exit the **Edit User** dialog box.

Add Roles

You must be logged into the DI Repository as an administrative user.

1. In Spoon, go to **Tools > Repository > Explore**. The **Repository Explorer** opens.
2. Click the **Security** tab.
3. Click the **Roles** radio button. The list of available roles appear.
4. Click **Add**. The **Add Role** dialog box appears.
5. Enter the **Role Name** in the appropriate field.
Note: An entry in the **Description** field is optional.
6. If you have users to assign to the new role, select them (using the <SHIFT> or <CTRL> keys) from the list of available users and click the yellow arrow to move it from the left pane to the right pane. The user(s) assigned to your new role appear in the right pane.
7. Click **OK** to save your entries and exit the **Add Role** dialog box.

The specified role is created and is ready to be assigned to user accounts.

Edit Roles

You must be logged into the DI Repository as an administrative user.

1. In Spoon, go to **Tools > Repository > Explore**. The **Repository Explorer** opens.
2. Click the **Security** tab.
3. Click the **Roles** radio button. The list of available roles appear.
4. Select the role you want to edit and click **Edit**. The **Edit Role** dialog box appears.
5. Make the appropriate changes.
6. Click **OK** to save your changes and exit the **Edit Role** dialog box.

Delete Roles

You must be logged into the DI Repository as an administrative user.

1. In Spoon, go to **Tools > Repository > Explore**. The **Repository Explorer** opens.
2. Click the **Security** tab.
3. Select the role you want to delete from the list of available roles.
4. Click **Remove**. A confirmation message appears.
5. Click **Yes** to delete the role.

The specified role is deleted.

Make Changes to the Administrator Role

The assigning of action-based permissions, (read, create, execute, and administrate), associated with the administrator role in the DI Repository cannot be edited in the user interface. The administrator role is the only role that is assigned the *Administer Security* permission; the Administer Security permission controls user access to the **Security** tab.

Deleting the administrator role prevents *all users* from accessing the **Security** tab, unless another role is assigned the administrator permission.

These are the scenarios that require a configuration change not available through Spoon:

- You want to delete the administrator role
- You want to unassign the administrator permission from the administrator role
- You want to configure LDAP

Follow these instructions to change the administrator role:

1. Shut down the DI Server.
2. Open the **repository.spring.xml** file located at `\pentaho\server\data-integration-server\pentaho-solutions\system\`.
3. Locate the element with an ID of **immutableRoleBindingMap**.
4. Replace the entire node with the XML shown below. Make sure you change **yourAdminRole** to the role that will have **Administrate** permission.

```
<util:map id="immutableRoleBindingMap">
  <entry key="yourAdminRole">
    <util:list>
      <value>org.pentaho.di.reader</value>
      <value>org.pentaho.di.creator</value>
      <value>org.pentaho.di.securityAdministrator</value>
    </util:list>
  </entry>
</util:map>
```

5. Restart the DI Server.

The administrator role changes according to your requirements.

Assign User Permissions in the Repository using Spoon

You must be logged into the repository as an administrative user (or be assigned to a role that has Administer Security permission).

You can restrict what users see by assigning roles to users. For example, you can create administrative groups who are allowed to administer security and create new content.

To assign permissions in the repository, follow these instructions.

1. In Spoon, go to **Tools > Repository > Explore**. The **Repository Explorer** opens.
2. Click the **Security** tab.
3. Click the **Roles** radio button. The list of available roles appear.
4. Select the role to which you want to assign permissions.
5. Enable the appropriate permissions and click **Apply**.

The permissions you enabled for the role take effect the next time the specified users log in.

Enable System Role Permissions

When users log into the DI Repository, they are automatically assigned the **Authenticated** system role in addition to the role you assigned them. Pentaho requires the **Authenticated** system role for users to log into the DI Repository. This includes administrative users. By default, the **Authenticated** system role provides **Read**

Content and **Execute** permissions to all users who are logged in. You can change these permissions as needed.


Note: **Important!** The Anonymous system role is not being used at this time.

Follow the steps below to change permissions for the Authenticated system role.

1. In Spoon, go to **Tools > Repository > Explore**. The **Repository Explorer** opens.
2. Click the **Security** tab.
3. Click the **System Roles** radio button. The list of available system roles appear.
Note: The Anonymous role is not functional.
4. Select the **Authenticated** role from the list of available roles.
5. Under **Permissions**, enable the appropriate permissions for this role.
6. Click **Apply** to save your changes.


The specified permissions are enabled for the **Authenticated** system role.

Next Steps



After you have performed these tasks, [configure the PDI tool and utilities](#).

Configure the PDI Design Tools and Utilities



Before you can use the Pentaho Data Integration tools, you need to perform configuration tasks for each workstation that runs the tools and utilities. Just follow the *Guide Post* graphics to ensure you complete the entire process. It should take you less than 1 hour for each workstation running the design tools.


If you used the **Pentaho Trial Download** and the Installation Wizard, the DI Server and its tools and utilities reside on the same workstation. If you used any of the other installation techniques, the DI Server and its tools and utilities may reside on separate workstations.

Here is the list of the Pentaho Data Integration tools and utilities and what they do.

Tool or Utility	What It Does
Carte	<ul style="list-style-type: none">• Enables remote execution of jobs and transformations continuously in the background• Enables distribution and coordinated job execution across a collection of computers, a process known as clustering
Kitchen	<ul style="list-style-type: none">• Provides a command-line method to run jobs immediately• Typically used after the ETL development phase for production environments
Pan	<ul style="list-style-type: none">• Provides a command-line method to run transformations immediately• Typically used after the ETL development phase for production environments
Spoon	Allows users to design and manage complex ETL workloads with a graphical user interface

- [Start and Stop PDI Design Tools and Utilities](#)
- [Increase the Spoon Memory Limit](#)
- [Enter Licenses and Connect to the DI Repository](#)
- [Next Steps](#)

Start and Stop PDI Design Tools and Utilities




Each of the DI tools has a [specific function](#), so you start and stop each one individually. As part of the Installation Wizard and Archive Installation, we provide different ways to help you start each tool and utility, depending on the operating system you use. This table shows you options. After you have decided how you want to start and stop the tool and utilities, please go on to the next stop on the *Guide Post* graphic.

Operating System	Installation Method	Start the Tools and Utilities
Microsoft Windows	Pentaho Installation Wizard	<p>Spoon</p> <ul style="list-style-type: none"> Click Start > Programs > Pentaho Enterprise Edition > Design Tools and double-click on Data Integration the item to launch Spoon. <p>For these tools and utilities, you must navigate to a directory to launch the .bat file.</p> <p>Carte</p> <ul style="list-style-type: none"> <code>\pentaho\design-tools\data-integration\carte.bat</code> <p>Kitchen</p> <ul style="list-style-type: none"> <code>\pentaho\design-tools\data-integration\kitchen.bat</code> <p>Pan</p> <ul style="list-style-type: none"> <code>\pentaho\design-tools\data-integration\pan.bat</code>
Microsoft Windows	Pentaho Archive or Manual Installation	<p>Create shortcuts to the .bat file or each design tool and add them to the Start menu. Or, navigate to the one of these directories and run the design tool's .bat file to launch it.</p> <p>Carte</p>

Operating System	Installation Method	Start the Tools and Utilities
		<ul style="list-style-type: none"> • \pentaho\design-tools\data-integration\carte.bat Kitchen <ul style="list-style-type: none"> • \pentaho\design-tools\data-integration\kitchen.bat Pan <ul style="list-style-type: none"> • \pentaho\design-tools\data-integration\pan.bat Spoon <ul style="list-style-type: none"> • \pentaho\design-tools\data-integration\spoon.bat
Linux	Pentaho Wizard, Archive, or Manual Installation	<p>Navigate to one of these directories and run the design tool .sh file to launch it.</p> Carte <ul style="list-style-type: none"> • pentaho/design-tools/data-integration/carte.sh Kitchen <ul style="list-style-type: none"> • pentaho/design-tools/data-integration\kitchen.sh Pan <ul style="list-style-type: none"> • /pentaho/design-tools/data-integration/pan.sh Spoon <ul style="list-style-type: none"> • /pentaho/design-tools/data-integration/spoon.sh
Mac OS	Pentaho Wizard, Archive, or Manual Installation	<p>Navigate to one of these directories and run the design tool .app file to launch it.</p> Carte <ul style="list-style-type: none"> • pentaho/design-tools/data-integration/carte.app Kitchen <ul style="list-style-type: none"> • pentaho/design-tools/data-integration\kitchen.app Pan <ul style="list-style-type: none"> • /pentaho/design-tools/data-integration/pan.app Spoon <ul style="list-style-type: none"> • /pentaho/design-tools/data-integration/spoon.app

Increase the Spoon Memory Limit



We recommend that you increase PDI's memory limit so the DI Server and Data Integration Design Tool (Spoon) can perform memory-intensive tasks, like process or sort large datasets or run complex transformations and jobs. You will need to increase the memory limit for both the DI Server and Spoon. If you choose to not increase the memory limit, PDI uses the default memory settings that are defined in your startup scripts.

Note: Instead of modifying the Spoon startup script, you can also set the environment variable `PENTAHO_DI_JAVA_OPTIONS` equal to `-Xmx2g -XX:MaxPermSize=256m` on your client.

Modify the Spoon Startup Script

1. Exit from Spoon if it is currently running.
2. Open the Spoon startup script with a text editor. The name of startup script depends on your operating system.

Windows Name and Location of Startup Script: `pentaho/design-tools/data-integration/Spoon.bat`

Linux Name and Location of Startup Script: `pentaho/design-tools/data-integration/Spoon.sh`

3. Modify the `-Xmx` value so that it specifies a larger upper memory limit. In this example, 2g (two gigabytes) of heap space has been allocated.

```
PENTAHO_DI_JAVA_OPTIONS="-Xmx2g -XX:MaxPermSize=256m"
```

4. Save and close the startup file.
5. Start Spoon and ensure that there are no memory-related errors. If you see an error, repeat these steps to increase the memory again. **Note:** You need to [increase the DI Server memory limit](#) as well.

Spoon now has higher memory limits.

Enter Licenses and Connect to the DI Repository



You must also use Spoon to set up each workstation.

- [Start Spoon](#)
- [Enter Licenses](#)
- [Make a Connection to the DI Repository](#)

When you are done, please go on to the next stop on the *Guide Post* graphic.

Start Spoon

1. Start the DI server.
2. Navigate to the folder where you have installed Pentaho Data Integration; for example `... \pentaho \ design-tools\data-integration`.
3. Launch Spoon in the best way for your operating system.
 - a. **For Windows:** double-click **Spoon.bat**.
 - b. **For Linux:** double-click **spoon.sh**
 - c. **For Macintosh:** go to `... /pdi-ee/data-integration` and double click on the **Data Integration 64-bit** icon depending on your system.

The **Pentaho License Manager** dialog appears. You must [enter licenses](#).

Enter Licenses

You must [start Spoon](#) before adding licenses.

1. Copy your license files to a convenient location.
2. Log into Spoon. The **Pentaho License Manager** dialog appears.
3. From within the **Pentaho License Manager**, click on the **Add** button.
4. Double-click on the license key to open it. A green check appears in the Status column to show that the license key installed correctly.
5. Click **Close** to close the dialog box.

6. Close any windows that appear and restart Spoon so that the license key will take effect.

Connect to the DI Repository

The DI Repository provides you a place to centrally store ETL jobs and transformations. You may choose to set up common connections for your users now or you can close the dialog box and allow your users to connect to the DI Repository themselves.

1. Start Spoon.
2. Select **Tools > Repository > Connect**.
3. In the **Repository Connection** dialog box, click the **Add** button.
4. Select **DI Repository: DI Repository** and click **OK**. The **Repository Configuration** dialog box appears.
5. Keep the default URL, which is <http://localhost:9080/pentaho-di>.
6. Click **Test** to ensure your connection is properly configured. If you get an error, make sure you have [started the DI Server](#).
7. Click **OK** to exit the **Success** dialog box.
8. Enter an **ID** and **Name** for your repository.
9. Click **OK** to exit the **Repository Configuration** dialog box. Your new connection appears in the list of available repositories.
10. Log on to the DI Repository by entering the following credentials: user name = **admin**, password = **password**. If you have changed the credentials, then those are the credentials you must use.

NOTE:

If you want the **Repository Connection** window to automatically appear when Spoon starts, go to **Tools > Options** and click **Show repository dialog at startup**.

Next Steps



After you have performed these tasks, you can choose your next steps.

- If you want to create jobs and transformations with your data, you must [connect to your data](#) and then [connect to the DI Repository](#) so you have somewhere to save your jobs and transformations.