HITACHI
Inspire the Next

Configuring Pentaho to Use Database-Based Security

HITACHI
Inspire the Next

Change log (if you want to use it):

Date	Version	Author	Changes
8/21/2017	1.0	Carlos Lopez	

Contents

Overview 1

Before You Begin 1

Terms You Should Know 1

Other Prerequisites 1

Use Case – Applying Pentaho to Existing Database-Based Security 1

Authentication and Authorization 2

Database Structure 3

Table Declarations 3

Table Population 3

`Users` Table 4

`Authorities` Table 4

`Granted_Authorities` Table 4

Configuring Pentaho to Use JDBC Security 5

Copy JDBC Driver 5

Change Pentaho's Default Security Provider 5

Connect Pentaho to Your Database 5

Map the Administrator Role 6

Understanding Queries Against Your JDBC Security 7

Spring Framework Queries 7

Pentaho Queries 8

Known Issues 11

Database and Table Structure are Different 11

Browse File Keeps Spinning with No Results 11

Passwords Stored in Cleartext 12

Related Information 12

Finalization Checklist 13

This page intentionally left blank.

Overview

This document covers some best practices on Java database connectivity (JDBC) authentication. In it, you will learn how to set up Pentaho to authenticate with a database-based authentication scheme.

Our intended audience is Pentaho administrators, or anyone with a background in authentication and authorization who is interested in applying JDBC.

The intention of this document is to speak about topics generally; however, these are the specific versions covered here:

Software	Version(s)
Pentaho	7.1

The Components Reference in Pentaho Documentation has a complete list of supported software and hardware.

Before You Begin

Before beginning, use the following information to prepare for the procedures described in the main section of the document.

Terms You Should Know

Here are some terms you should be familiar with:

- Spring Framework: A configurable access control method for Java-based enterprise applications

Other Prerequisites

This document assumes that you have some background in database administration and network authentication and that you have already installed Pentaho. More information about related topics outside of this document can be found at Installation and Security.

Use Case – Applying Pentaho to Existing Database-Based Security

Use cases employed in this document include the following:

> *In the case of internal applications that use their own authentication, rather than Microsoft Active Directory, users' passwords may be different from their Windows authentication information. JDBC security can be useful in these scenarios where you already have an existing database-based security and wish to plug Pentaho into it. The administrator of the database-based security can manage user access to Pentaho by manipulating their own tables and queries.*

You could use a hybrid configuration for this as well. Manual Hybrid Configuration in Pentaho documentation has information of doing this.

Authentication and Authorization

To configure Pentaho to use a database-based authentication scheme, you must first know how Pentaho uses authentication and authorization, and how Spring Framework fits in.

Authentication occurs when the user logs in with their credentials and the system checks to make sure the user is valid and active. Once the user is validated, the system checks to see what roles the user has, which will define what the user is authorized to do on the server. Roles are assigned only once authentication has occurred, and they handle operational permissions.

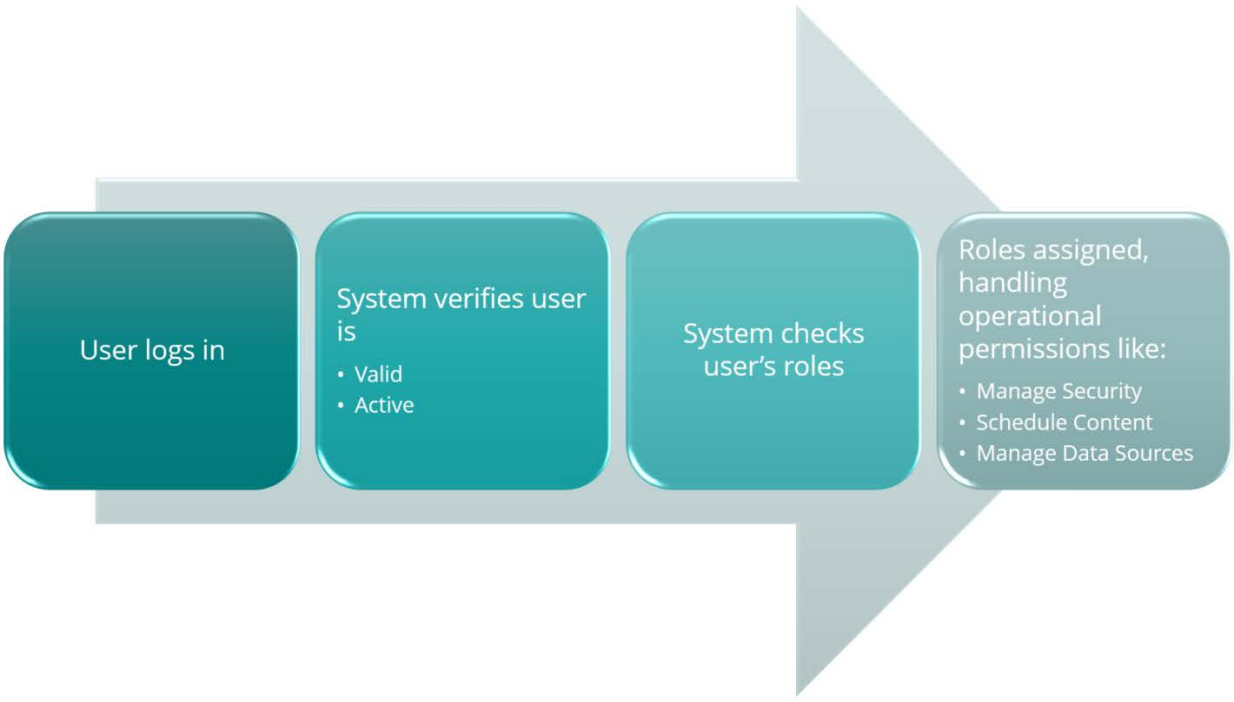

Figure 1: Authentication

Seeing the contents of a report is controlled by Mondrian roles in Analyzer reports and platform roles in Metadata models used in Interactive Reports and Pentaho Report Designer Reports, and should not be confused with authorization, which allows a user to open a report. Seeing the contents of a report is security-constrained access and beyond the scope of this document.

Database Structure

To work with Pentaho and Spring Framework, you can use any database structure as long as it has, at a minimum, the equivalent of the following three tables. You can find information for Setting Up User Security in the Pentaho documentation.

1. **Users**: containing user information
2. **Authorities**: containing role information
3. **Granted_Authorities**: combining users and roles granted each user. In practice, there should be a one-to-one relationship between user and role in the table, so if a user has four roles, there should be four entries for the user, one for each role.

Table Declarations

These table declarations are minimal settings for JDBC security:

1. Create table `users`:

```
CREATE TABLE USERS (
USERNAME VARCHAR2(50) NOT NULL PRIMARY KEY,
PASSWORD VARCHAR2(50) NOT NULL,
ENABLED INTEGER DEFAULT 1 NOT NULL,
DESCRIPTION VARCHAR2(100);
```

2. Create table `authorities`:

```
CREATE TABLE AUTHORITIES(
AUTHORITY VARCHAR(50) NOT NULL PRIMARY KEY,
DESCRIPTION VARCHAR(100));
```

3. Create table `granted_authorities`:

```
CREATE TABLE GRANTED_AUTHORITIES(
USERNAME VARCHAR(50) NOT NULL,
AUTHORITY VARCHAR(50) NOT NULL,
CONSTRAINT FK_GRANTED_AUTHORITIES_USERS FOREIGN KEY(USERNAME) REFERENCES
USERS(USERNAME),
CONSTRAINT FK_GRANTED_AUTHORITIES_AUTHORITIES FOREIGN KEY(AUTHORITY)
REFERENCES AUTHORITIES(AUTHORITY))
```

Table Population

Skip these steps if you already have these tables built and used in other applications.

`Users` Table

By default, we use the PlaintextPasswordEncoder, which reads the password as it is in the database. You can use a different password encoder if you want. Here is example code to illustrate the `users` table with plaintext passwords:

```
INSERT INTO USERS VALUES('gabellard','Password1',1,NULL);
INSERT INTO USERS VALUES('wfaulkner','Password1',1,NULL);
INSERT INTO USERS VALUES('clopez','Password1',1,NULL);
INSERT INTO USERS VALUES('skemparaju','mypassword',1,NULL);
```

`Authorities` Table

Example code to illustrate the `authorities` table:

```
INSERT INTO AUTHORITIES VALUES('DBPentAdmins','Super User');
INSERT INTO AUTHORITIES VALUES('DBPentHR','HR Users');
INSERT INTO AUTHORITIES VALUES('DBPentFinance','Finance Users');
INSERT INTO AUTHORITIES VALUES('DBPentUsers','User has not logged in');
INSERT INTO AUTHORITIES VALUES('DBPentSales','Sales Users');
```

`Granted_Authorities` Table

Example code to illustrate the `granted_authorities` table with a one-to-one relationship between users and roles:

```
INSERT INTO GRANTED_AUTHORITIES VALUES('gabellard','DBPentAdmins');
INSERT INTO GRANTED_AUTHORITIES VALUES('gabellard','DBPentUsers');
INSERT INTO GRANTED_AUTHORITIES VALUES('clopez','DBPentUsers');
INSERT INTO GRANTED_AUTHORITIES VALUES('clopez','DBPentFinance');
INSERT INTO GRANTED_AUTHORITIES VALUES('wfaulkner','DBPentUsers');
INSERT INTO GRANTED_AUTHORITIES VALUES('wfaulkner','DBPentHR');
INSERT INTO GRANTED_AUTHORITIES VALUES('skemparaju','DBPentSales');
INSERT INTO GRANTED_AUTHORITIES VALUES('skemparaju','DBPentUsers');
```

Configuring Pentaho to Use JDBC Security

The following steps assume that Pentaho has already been installed. Make sure you follow these steps after you have stopped the Pentaho Server.

Copy JDBC Driver

Filling in your own installation path, copy the JDBC driver to:

```
<installation path>/pentaho-server/tomcat/lib/
```

Change Pentaho's Default Security Provider

Pentaho's default security provider is Jackrabbit. To change this to JDBC, follow these steps:

1. Locate the file <installation path>/pentaho-server/Pentaho-solutions/system/security.properties.
2. Change the provider from provider=jackrabbit to provider=jdbc
3. Save the file.

Connect Pentaho to Your Database

Since you have already copied the JDBC driver to tomcat/lib, you can now connect Pentaho to your database with these steps:

1. Locate the file `<installation path>/pentaho-server/Pentaho-solutions/system/applicationContext-spring-security-jdbc.properties`.
2. Add the correct database information. This example uses PostgreSQL.

```
# The fully qualified Java class name of the JDBC driver to be used
datasource.driver.classname=org.postgresql.Driver

# The connection URL to be passed to our JDBC driver to establish a
connection
datasource.url=jdbc:postgresql://localhost:5432/jdbc_auth

# The connection username to be passed to our JDBC driver to establish a
connection
datasource.username=postgres

# The connection password to be passed to our JDBC driver to establish a
connection
datasource.password=password

# The SQL query that will be used to validate connections from this pool
before returning them to the caller.
# This query must be an SELECT statement that returns at least one row.
# HSQLDB: SELECT 1 FROM INFORMATION_SCHEMA.SYSTEM_USERS
# MySQL, H2, MS-SQL, POSTGRESQL, SQLite: SELECT 1
# ORACLE: SELECT 1 FROM DUAL
datasource.validation.query=SELECT 1
```

```
# the maximum number of milliseconds that the pool will wait (when there
are no available connections)
# for a connection to be returned before throwing an exception, or <= 0 to
wait indefinitely. Default value is -1
datasource.pool.max.wait=-1

# The maximum number of active connections that can be allocated from this
pool at the same time, or negative for no limit. Default value is 8
datasource.pool.max.active=8

# The maximum number of connections that can remain idle in the pool,
without extra ones being destroyed, or negative for no limit. Default value
is 8
datasource.max.idle=4

# The minimum number of active connections that can remain idle in the
pool, without extra ones being created when the evictor runs, or 0 to
create none. Default value is 0
datasource.min.idle=0
```

Map the Administrator Role

Make sure you map the Administrator role correctly using these steps:

1. Locate the file `<installation path>/pentaho-server/Pentaho-solutions/system/applicationContext-pentaho-security.jdbc.xml`.
2. Change the `<entry key>` to the admin role value from the database, from

```
<util:map id="jdbcRoleMap">
<entry key="Admin" value="Administrator"/>
</util:map>
```

   to

```
<util:map id="jdbcRoleMap">
<entry key="DBPentAdmins" value="Administrator"/>
</util:map>
```

3. Save the file.

Understanding Queries Against Your JDBC Security

Now that Pentaho Server has been configured to use your JDBC security, you can find further information about the queries against your database in this section.

Spring Framework Queries

As a user logs in, two queries are fired: one to get information about the user, and one to find out what roles the user belongs to. You can find this in the log with these steps:

1. Locate the file `<installation path>/pentaho-server/tomcat/webapps/pentaho/WEB-INF/classes/log4j.xml`.
2. Add the following categories:

```
<category name="org.springframework.security">
      <priority value="DEBUG"/>
</category>
```

3. In the `Pentaho.log`, you will see:

```
DEBUG
[org.springframework.security.web.authentication.AnonymousAuthenticationFil
ter] SecurityContextHolder not populated with anonymous token, as it
already contained:
'org.springframework.security.authentication.UsernamePasswordAuthentication
Token@f81a4943: Principal:
org.springframework.security.core.userdetails.User@fc211e2b: Username:
skemparaju; Password: [PROTECTED]; Enabled: true; AccountNonExpired: true;
credentialsNonExpired: true; AccountNonLocked: true; Granted Authorities:
Authenticated,DBPentSales,DBPentUsers; Credentials: [PROTECTED];
Authenticated: true; Details:
org.springframework.security.web.authentication.WebAuthenticationDetails@25
5f8: RemoteIpAddress: 127.0.0.1; SessionId:
4713A775E8A737E8ED4A3E2E768B5653; Granted Authorities: Authenticated,
DBPentSales, DBPentUsers'
```

   The queries that result in this information are found in `<installation path>/pentaho-server/Pentaho-solutions/system/applicationContext-spring-security-jdbc.xml`.
4. The `usersByUsernameQuery` loads the username and password:

```
SELECT username, password, enabled FROM USERS WHERE username = ? ORDER BY
username
```

5. The `authoritiesByUsernameQuery` loads the roles the user belongs to:

```
SELECT username, authority FROM GRANTED_AUTHORITIES WHERE username = ?
ORDER BY authority
```

Pentaho Queries

When you start Pentaho, it will connect to the database to gather information about the user and authorities for Pentaho authorization. The Administrator role and user are used to retrieve this information.

To see this information in the Pentaho.log, follow these steps:

1. Locate the file `<installation path>/pentaho-server/tomcat/webapps/pentaho/WEB-INF/classes/log4j.xml`.
2. Add the following categories:

```
<category name="org.pentaho.platform.engine.security">
      <priority value="DEBUG"/>
</category>
```

3. In the `Pentaho.log`, you will see:

```
DEBUG
[org.pentaho.platform.plugin.services.security.userrole.jdbc.JdbcUserRoleLi
stService$AllAuthoritiesMapping] RdbmsOperation with SQL [SELECT
distinct(authority) as authority FROM AUTHORITIES ORDER BY authority]
compiled

DEBUG
[org.pentaho.platform.plugin.services.security.userrole.jdbc.JdbcUserRoleLi
stService$AllUserNamesInRoleMapping] RdbmsOperation with SQL [SELECT
distinct(username) as username FROM GRANTED_AUTHORITIES where authority = ?
ORDER BY username] compiled

DEBUG
[org.pentaho.platform.plugin.services.security.userrole.jdbc.JdbcUserRoleLi
stService$AllUserNamesMapping] RdbmsOperation with SQL [SELECT
distinct(username) as username FROM USERS ORDER BY username] compiled
```

   Note that these are the same queries shown in `<installation path>/pentaho-server/Pentaho-solutions/system/applicationContext-pentaho-security-jdbc.xml`.
4. The `allAuthoritiesQuery`, used to show all the roles on the Authorities table, is written as:

```
SELECT distinct(authority) as authority FROM AUTHORITIES ORDER BY authority
```

5. These roles are later displayed under **Users & Roles** on the Administration Perspective in the Pentaho User Console (PUC):

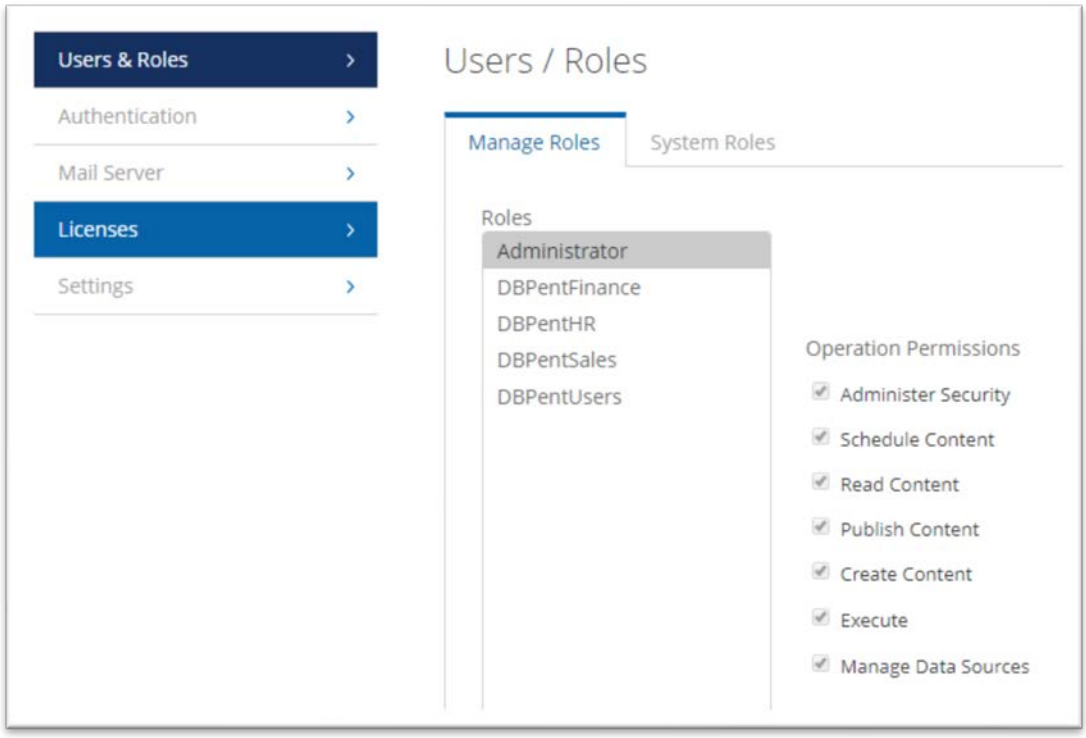

Figure 2: Users & Roles in the PUC

6. The `allUsernamesInRoleQuery`, which gets all users who belong to a specific role, is written as:

```
SELECT distinct(username) as username FROM GRANTED_AUTHORITIES where
authority = ? ORDER BY username
```

7. The `allUsernamesQuery`, which shows the users on the Share or Permissions tab, is written as:

```
SELECT distinct(username) as username FROM USERS ORDER BY username
```

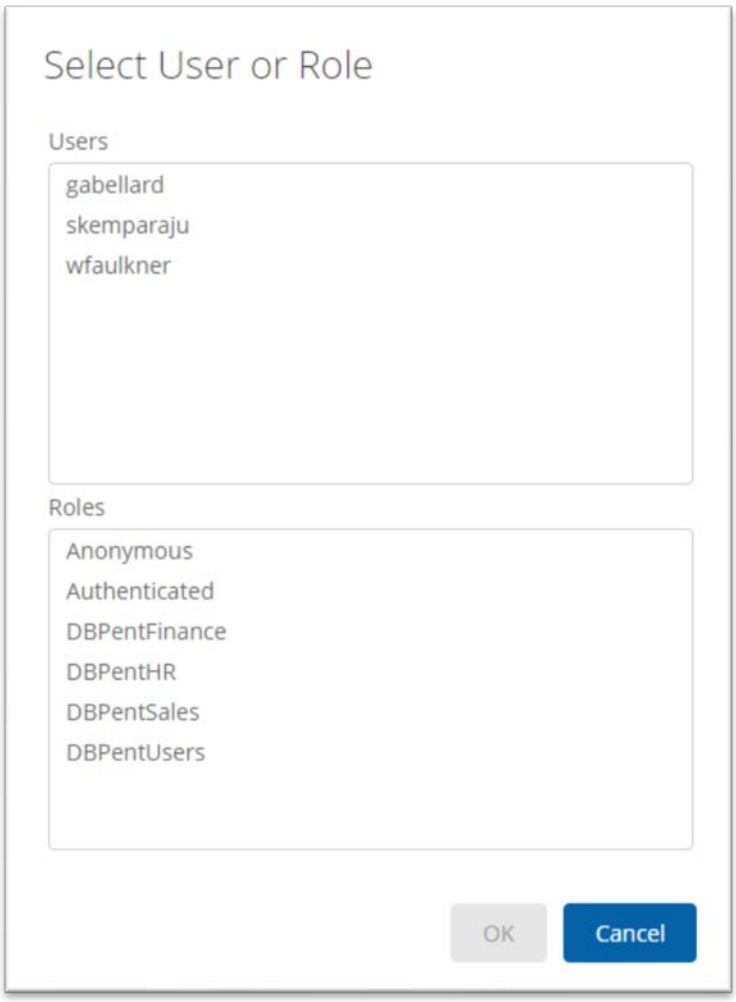

Figure 3: Select User or Role

8. Check the `Pentaho.log` for the Administrator role and user:

```
DEBUG
[org.springframework.security.core.userdetails.cache.EhCacheBasedUserCache]
Cache hit: true; username: gabellard

DEBUG [org.pentaho.platform.engine.security.SecurityHelper]
rolesForUser:[Authenticated, Administrator, DBPentUsers, Anonymous]
```

These must be active in your database, or you will not be able to connect to your database to extract the roles and users necessary for the Pentaho Server to function properly.

Known Issues

In this section are a few known issues to be aware of, as well as solutions for each.

Database and Table Structure are Different

If you are already using a different table structure for your JDBC authentication, make sure you use an alias for the different field names, as illustrated in this example:

```
SELECT userid as username, 'password' as password, 'enabled' as enabled
FROM USERS_ROLES WHERE userid= ? ORDER BY userid
```

Browse File Keeps Spinning with No Results

After configuring JDBC security, you may run into an issue where your browse file constantly spins but does not show anything. The `Catalina` log shows an example:

```
SEVERE: The RuntimeException could not be mapped to a response, re-throwing
to the HTTP container

org.pentaho.platform.api.repository2.unified.UnifiedRepositoryException:
exception while getting tree rooted at path "/"

Reference number: 2f863f91-f38f-4176-91a2-a0fb43a73af2

at
org.pentaho.platform.repository2.unified.ExceptionLoggingDecorator.callLogT
hrow(ExceptionLoggingDecorator.java:512)
```

The `Pentaho.log` will show comparable results. This can happen for any of the following reasons:

- One of the queries in either configuration file is returning a null value.
- Passwords, roles, or granted roles are null.
- Users are not properly disabled.
- A role is not properly mapped to a user in the `granted_authorities` table.

To fix this, follow these steps:

1. Locate the file `<installation path>/pentaho-server/tomcat/webapps/Pentaho/WEB-INF/classes/log4j.xml`.
2. Add the following category to the file:

```
<category
name="org.pentaho.platform.repository2.unified.ExceptionLoggingDecorator">

<priority value="DEBUG"/>

</category>
```

3. Restart the Pentaho Server
4. Log back in again and choose **Browse Files**.

Passwords Stored in Cleartext

By default, Pentaho connects to your database using a cleartext password stored in the file `<installation-path>/pentaho-server/Pentaho-solutions/system/applicationContext-spring-security-jdbc.properties`.

A workaround to this is to use an account or database login that is only for this database. You will need `READ ONLY` permissions for this. Do not use something like a system administrator account or similar.

Related Information

Here are some links to information that you may find helpful while using this best practices document:

- Installation
- Pentaho Components Reference
- Security
- Setting Up User Security
- Spring Framework
- Starting and Stopping the Pentaho Server

Finalization Checklist

This checklist is designed to be added to any implemented project that uses this collection of best practices, to verify that all items have been considered and reviews have been performed. (Compose specific questions about the topics in the document and put them in the table.)

Name of the Project: _____

Date of the Review: _____

Name of the Reviewer: _____

Item	Response	Comments
Did you install Pentaho before beginning?	YES_____ NO_____	
Did you familiarize yourself with how Pentaho uses authentication and authorization?	YES_____ NO_____	
Are you familiar with Spring Framework?	YES_____ NO_____	
Did you follow the table declaration settings for JDBC security?	YES_____ NO_____	
Did you stop the Pentaho server before starting your configuration?	YES_____ NO_____	
Did you use aliases for the different field names, if you have a different table structure for your JDBC authentication?	YES_____ NO_____	