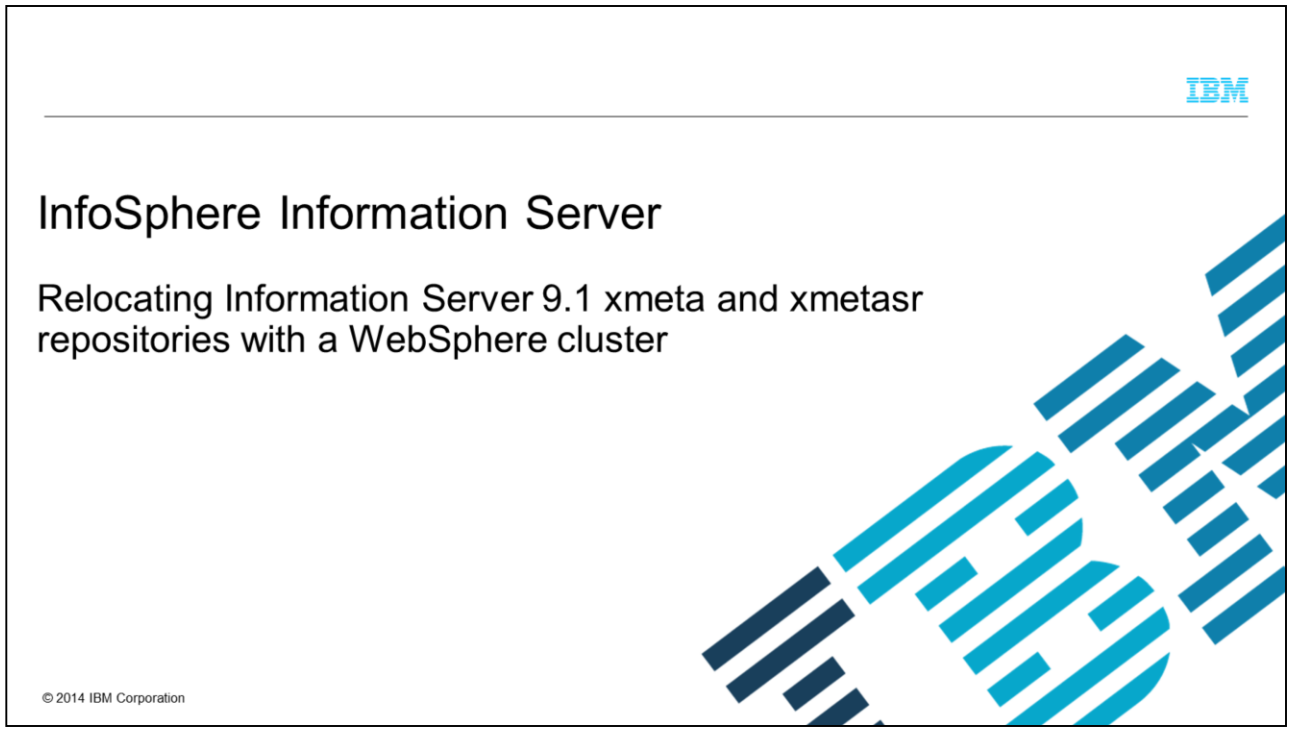

This presentation describes the different databases and repositories included with Information Server version 9.1 and how to relocate xmeta and xmetasr to another database server. This presentation discusses the relocation when using a WebSphere® cluster. If your configuration is using stand-alone WebSphere, see the IBM Education Assistant module on relocating using stand-alone WebSphere.

Objectives

- The databases and repositories in Information Server 9.1
- What repositories and databases need to be relocated
- Backing up XMeta, IS, and key files
- Updating XMETA and XMETASR
  - Update Information Services Framework (ISF) configuration
  - Update WebSphere Application Server configuration

The objectives of this presentation are to describe the databases and repositories that Information Server version 9.1 includes, how to check what you need to relocate, what you need to back up, and how to update xmeta and xmetasr including the Information Services Framework configuration, referred to as the ISF configuration. This presentation also describes what changes you need to make to the WebSphere cluster.

Information Server databases and repositories (1 of 2)

Repository or database	Description	Default database and schema
Active InfoSphere® Information Server metadata repository	Stores the metadata about external data sources that InfoSphere Information Server components govern, manage, and analyze. Normally referred to as the metadata repository.	Database: XMETA Schema: XMETA Database must be the same database that you used for the staging metadata repository.
InfoSphere Information Server staging metadata repository	Stores metadata that you imported from external data sources so that you can examine it before you move it to the active metadata repository.	Database: XMETA Schema: XMETASR Database must be the same database that you used for the active metadata repository.
InfoSphere Information Analyzer analysis database	Stores results of information analysis by InfoSphere Information Analyzer.	Database: IADB Schema: IAUSER Database cannot be the same database that you used for the active or staging metadata repository.
InfoSphere DataStage® and QualityStage operations database Repository.tool=DataStage	Stores monitoring data that the InfoSphere DataStage displays and QualityStage Operations Console.	Database: XMETA Schema: User-defined repository user name, typically DSODB Database can be the same or different as the database that you used for the metadata repository.

You might need to change the configuration of databases and repositories that you used to run InfoSphere Information Server or its components. You might need to make this change for various reasons: a change in the physical computer that hosts the repository, a change in host name or port number of the computer that hosts the repository, a change to the high-availability configuration, or a password change.

You might deploy some repositories as separate schemas within the same database. This deployment is always the case with the active metadata repository and the staging metadata repository. You can also create other repositories as schemas in this same database or in a separate database. Such a separate database can also hold multiple repositories. This slide and the following slide list and describe all of the databases and repositories that IBM ships with Information Server.

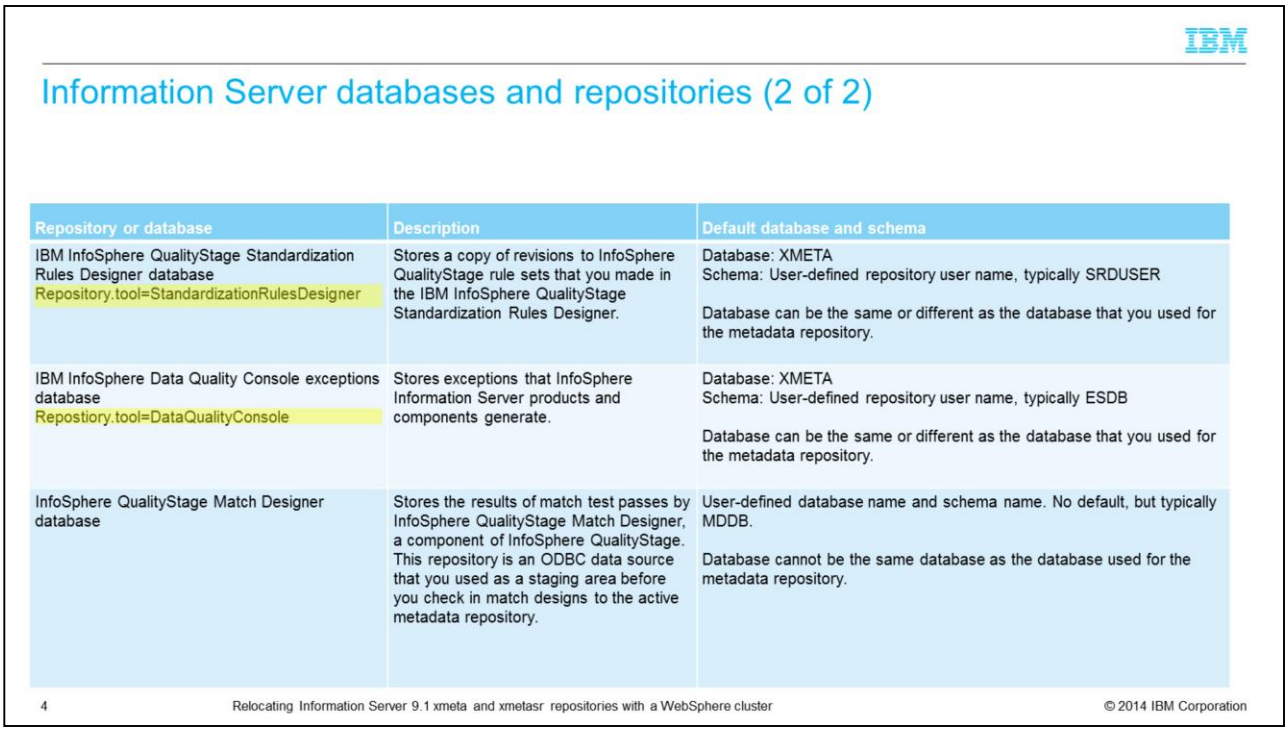

Use the table displayed on this slide to decide which databases or repositories that you need to relocate. When moving xmeta, it is important to check which of these databases you registered with the xmeta repository. You need to relocate any database that you registered with xmeta as well. The InfoSphere Information Analyzer Analysis Database, or IADB, and the QualityStage Match Designer Database, referred to as MDDB cannot be in the same database, so this presentation does not discuss them. Refer to separate IBM Education Assistant modules on relocating these databases.

Listing repositories and databases

- List repositories and databases
  - UNIX® or Linux®
    cd <is_installPath>/ASBServer/bin
    ./RepositoryAdmin.sh -listRepositories
  - Windows®
    cd <is_installPath>\ASBServer\bin
    .\RepositoryAdmin –listRepositories
- Example:
  $ ./RepositoryAdmin.sh –listRepositories
  dsodb
  QSSRDDB

The next step is to run the RepositoryAdmin tool to list out all of the installed repositories. Use the appropriate command that is displayed on this slide for your operating system.

Check repository location

- Check repository location

```
./RepositoryAdmin.sh -displayRepository -rn dsodb
DatabasePlatform.databaseType=DB2
DatabasePlatform.version=10.1
DatabaseServer.host=ipsvm00079.swg.usma.ibm.com
DatabaseServer.port=50000
Database.name=xmeta
Database.alias=null
Database.location=/opt/IS91/IBM/InformationServer/Repos/dsodb
Repository.name=dsodb
Repository.description=DSODB
Repository.tool=DataStage
Repository.context=
Repository.schema=dsodb
RepositoryConnection.name=dsodb
RepositoryConnection.userName=dsodb
RepositoryConnection.password={iisenc}N2RHakj6gLz7fCJ2yknhIg==
RepositoryConnection.connectionURL=jdbc:db2://ipsvm00079.swg.usma.ibm.com:50000/dsodb
RepositoryConnection.managedDataSourceName=Tablespace.name=DSODBSPACE
```


Next, check each of the repositories that are listed in the last step to see which database they are in. You need to relocate all databases in xmeta, so it is important to make note of these databases to ensure that you complete all necessary repository relocations. Use the RepositoryAdmin command that is shown in this slide to obtain the value of Database.name. The –rn argument takes a repository name and is case-sensitive. Be sure to enter the name exactly as it looks in the previous step.

If the Database.name value is equal to xmeta, then you need to move the repository. It is possible that there might not be any additional databases in xmeta. If you changed the repository names from the defaults at installation time and it is unclear which repository you are working with based on the repository name, check the field that is called Repository.tool and match the name to the proper repository seen on the chart on slides 3 and 4.

Back up and restore

- Backup
  - Backup databases, repositories, Information Server, and WebSphere
  - Back up all files being changed
    - InformationServer/ASBServer/bin/sql/database.properties
    - InformationServer/ASBServer/apps/lib/ojb-conf.jar
  - **Do not leave copy of ojb-conf.jar in Information Server or WebSphere folder hierarchy**
- Restore
  - Restore databases and/or repositories to new target system

Before you make any changes to the databases or repositories, Information Server, or WebSphere, it is good practice to take a complete backup of all installations. It is safest to do a cold backup of the WebSphere Application Server by stopping WebSphere before you do the backup. It is also good practice to make a copy of all the files that you changed during this process to make it easier to revert if necessary. This slide displays the files that are key to make copies of. Ensure that you do not leave the backup of ojb-conf.jar in the IBM Information Server or WebSphere folder hierarchy. Back up the databases and repositories on the source system and restore to the target system using the backup and restore tools that are provided with the database. Back up the affected files before changing them in this procedure.

Updating xmeta and xmetasr

- Xmeta- Active InfoSphere Information Server metadata repository
- Xmetasr - InfoSphere Information Server staging metadata repository
- Same database but different schema names
- Must move them together

The first topics that this presentation discusses are the metadata repository database, commonly known as xmeta, and the metadata staging repository, commonly known as xmetasr. These two repositories are in the same database but in their own schemas; you must move them together.

Updates to ISF and WebSphere configuration (1 of 4)

- Stop WebSphere Application cluster members
- Create temporary empty directory on your Domain Server and make it current working directory
  - **Windows:**
    mkdir c:\tmp\isftmp
    cd \tmp\isftmp
  - **Linux or UNIX:**
    mkdir /tmp/isftmp
    cd /tmp/isftmp

The next step is to update the ISF configuration. To do this update, first you need to stop the WebSphere Application cluster members.

After stopping the WebSphere Application cluster members, create a temporary directory on your domain server and set it as your current working directory.

Update ISF and WebSphere configuration (2 of 4)

- Extract ojb-conf.jar into temp directory
  - **Linux or UNIX:**
    /opt/IBM/WebSphere/AppServer/java/bin/jar xf /opt/IBM/InformationServer/ASBServer/apps/lib/ojb-conf.jar
  - **Windows:**
    C:\IBM\WebSphere\AppServer\java\bin\jar xf c:\IBM\InformationServer\ASBServer\apps\lib\ojb-conf.jar
- Edit repository_database.xml
  Linux or UNIX: `vi repository_database.xml`
  Windows: `write repository_database.xml`
- File contains multiple dbalias entires
  - Update all dbalias host name and port number attributes

While in your newly created temp directory, extract the ojb-conf.jar file using the jar utility of a JDK, for example, the JDK in WebSphere.

There are example commands displayed on this slide. This command extracts ojb-conf.jar and places the contents in your temp directory.

Next, edit the repository_database.xml file that is in your temp directory. Use the vi command for Linux and UNIX or open the file in WordPad by using the write command if on Windows. Search for all of the dbalias attributes. Use the table that is displayed on the next slide to correctly update this field. Edit every dbalias attribute in the file with the new host name and port number, and save the file.

Update ISF and WebSphere configuration (3 of 4)

DB2®	dbalias="//host:port/dbname" Example: dbalias="//db2host:50000/xmeta"
DB2 cluster or HADR database system	dbalias="//*host:port/dbname*;clientRerouteAlternateServerName=*alternate_host*;clientRerouteAlternatePortNumber=*alternate_host_port*;queryCloseImplicit=2;"
Oracle	dbalias="oracle://host:port;SID=dbname" Example: dbalias="oracle://oracleHost:1521;SID=xmeta"
Oracle RAC	dbalias="oracle://*host:port*;ServiceName=*SID*;AlternateServers=(*host:port*, *host:port*,*host:port*, …) " Example: dbalias="oracle://rac1:1521;serviceName=orcl;alternateServers=(rac1:1521,rac2:1521,rac3:1521) "
MS SQLServer	dbalias="sqlserver://host:port;DatabaseName=dbname" Example: dbalias="sqlserver://sqlHost:1433;DatabaseName=xmeta"
MS SQLServer using Named Instance	dbalias="sqlserver://host\named_instance:port;DatabaseName=dbname" Example: dbalias="sqlserver://sqlHost\my_instance:1433;DatabaseName=xmeta"

This slide displays the dbalias values and examples for DB2, Oracle, and SQL Server. If you are using SQL Server named instances, note that the format is sqlserverHostname\named_instance.

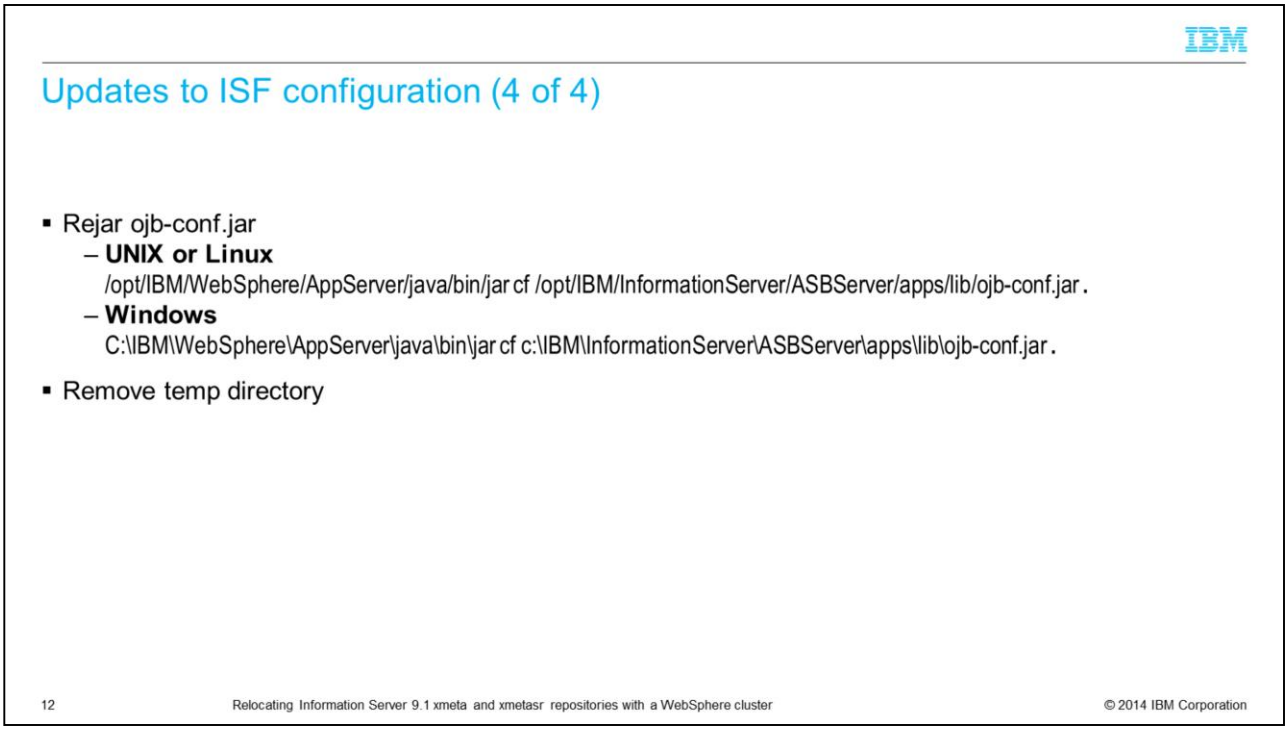

After updating the dbalias attribute, rejar ojb-conf.jar with the updated repository_database.xml file using the jar utility of a JDK. For example, the JDK in WebSphere. Be sure that you are still in your temp directory. This slide includes example commands. Be sure that your paths are correct for your installation. You must remember to put the “space dot” at the end of the jar command.

After this step completes, delete the temp directory.

Test changes

- Run PropertyAdmin from ASBServer
  - **UNIX or Linux:**
    bin/PropertyAdmin.sh –d
  - **Windows:**
    bin\PropertyAdmin.bat -d

Next, check that the new ojb-conf.jar is correct. To do this step, run the PropertyAdmin command in ASBServer. You need to be sure that this command returns successfully. If it does not, go back and check the changes that you made to ojb-conf.jar before continuing.

Propagate changes to WebSphere

- Be sure to have 1.5 GB+ free space
  - AIX/Linux - /tmp
  - Solaris/HPUX - /var/tmp
  - Windows - %TEMP%
- Run FilePropagator command
  - UNIX or Linux
    <IS_HOME>/ASBServer/bin/FilePropagator.sh -user *wasadmin_user* -password *wasadmin_password* -apps
  - Windows
    <IS_HOME>\ASBServer\bin\FilePropagator.bat -user *wasadmin_user* -password *wasadmin_password* -apps

The next step is to propagate the changes to WebSphere using the FilePropagator tool. This tool requires the system temporary directory have at least 1.5 gigabytes of free space. This location varies depending on your operating system. Run FilePropagator as displayed on this slide. Be sure to set wasadmin_user to the WebSphere primary administrative user. FilePropagator propagates many files to the cluster and takes some time to complete. Do not break out of this command.

Updates to ISF configuration

- Edit database.properties
  - **Linux or UNIX:**
    vi <IS_HOME>/ASBServer/bin/sql/database.properties
  - **Windows:**
    write C:\<IS_HOME>\ASBServer\bin\sql\database.properties
- Find and update URL parameter

DB2	url=jdbc\:db2\://NewServer.com\:port/dbname
Oracle	url=jdbc\:ibm\:oracle\://*host*\:*port*;SID=*SID*
Oracle RAC	url=jdbc\:ibm\:oracle\://*host*\:*port*;serviceName=*service;alternateServer=(host*\:*port, host*\:*port, host*\:*port, …)*
SQL Server	url=jdbc\:ibm\:sqlserver\://*host:port*\;DatabaseName=*dbname*
SQL Server - Named Instance	url=jdbc\:ibm\:sqlserver\://*host**instance_name*\:*port*;DatabaseName=*dbname*

15 Relocating Information Server 9.1 xmeta and xmetasr repositories with a WebSphere cluster

The next step is to edit the database.properties file in the InformationServer/ASBServer/bin/sql directory. Find and update the URL parameter to reflect the new repository server name and port. After updating, save the file.

Test changes

- Test changes to database.properties
  - UNIX or Linux
    <IS_HOME>/ASBServer/bin/AppServerAdmin.sh –w –time 0
  - Windows
    <IS_HOME>\ASBServer\bin\AppServerAdmin.bat –w –time 0

Next, test the changes to the database.properties file. Run the AppServerAdmin command that is displayed on this slide that is appropriate for your operating system. The command just needs to return successfully. If it errors, go back and verify the changes that are made to the database.properties file.

Synchronize nodes (1 of 2)

- Manually run WebSphere node synchronization
  - Administrative console
    - System Administration => Nodes => Synchronize
  - If unable to log in to the administrative console
    - Restart WebSphere Deployment Manager

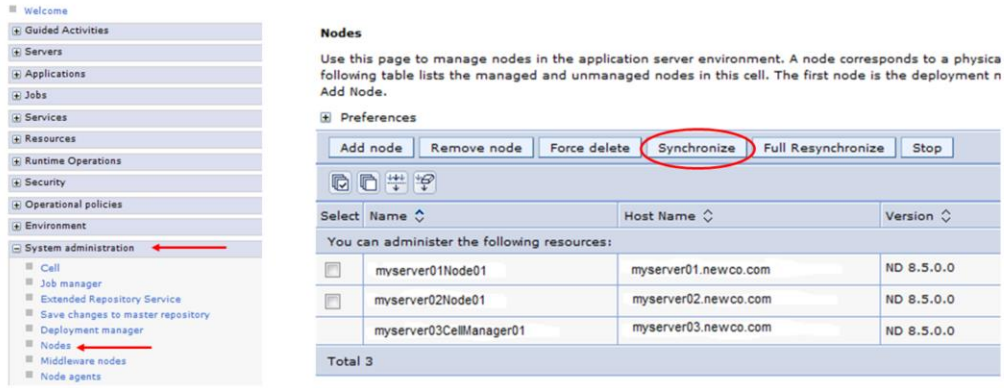

Next, it is a good idea to manually run WebSphere node synchronization to sync everything properly. You can do this synchronization through the administrative console. From the administrative console, click System Administration, Nodes, and click the Synchronize button.

In some cases, you might not be able to log in to the administrative console after making the changes in the previous slides. If so, restart the WebSphere Deployment manager to get back into the administrative console.

Synchronize nodes (2 of 2)

- If node agent is not running on a particular node
  - Manually run WebSphere node synchronization
    - Administrative console
    - If node agents are not running

      UNIX or Linux:
      <was_profile_dir>/bin/syncNode.sh <dmgr_host> <dmgr_port>

      Windows:
      <was_profile_dir>\bin\syncNode.bat <dmgr_host> <dmgr_port>
- Restart the Node Agents
- Restart the Deployment Manager

If there is a node or nodes in the cluster where the node agent is not running, you cannot do the synchronization by way of the administrative console. In this case, you can do the synchronization by running the syncNode command that is displayed on this slide on the node profile that you need to synchronize. dmgr_host is the name of the host running the Deployment Manager and dmgr_port is the port the Deployment Manager is running on. The default value for dmgr_port is 8879.

Once the synchronization is complete, restart the node agents and the deployment manager if not already restarted from the previous slide.

Update Version.xml

- Update <IS_HOME>/Version.xml on Information Server server
- Open Version.xml in text editor
  - Locate following XML element
    <PersistedVariable encrypted="false" name="**xmeta.db.hostname**" persistent="true" readonly="false" value="myserver"/>
    Modify "*value*" attribute with new xmeta server name
  - SQLServer with named instances required format:
    value="myserver\named_instance"
  - Locate next XML element
    <PersistedVariable encrypted="false" name="**xmeta.db.port**" persistent="true" readonly="false" value="50000"/>
    Modify "*value*" attribute with new port number

Next, update the values in the Version.xml file. This file contains the installation records that the installation program uses. Keeping the file current avoids problems with future installations. The Version.xml file is in the IBM InfoSphere Information Server installation directory on the same server as the ojb-conf.jar file.

Open the Version.xml file in a text editor and locate the PersistedVariable XML element that has the name attribute equal to xmeta.db.hostname. You need to modify the value attribute to contain the correct xmeta server name. Note that if you are using SQL Server with named instances, you need to use the format of servername\named_instance for the xmeta server name.

Next, you need to locate the PersistedVariable XML element that has the name attribute equal to xmeta.db.port. Modify the value for port if your port number has changed. Save your changes.

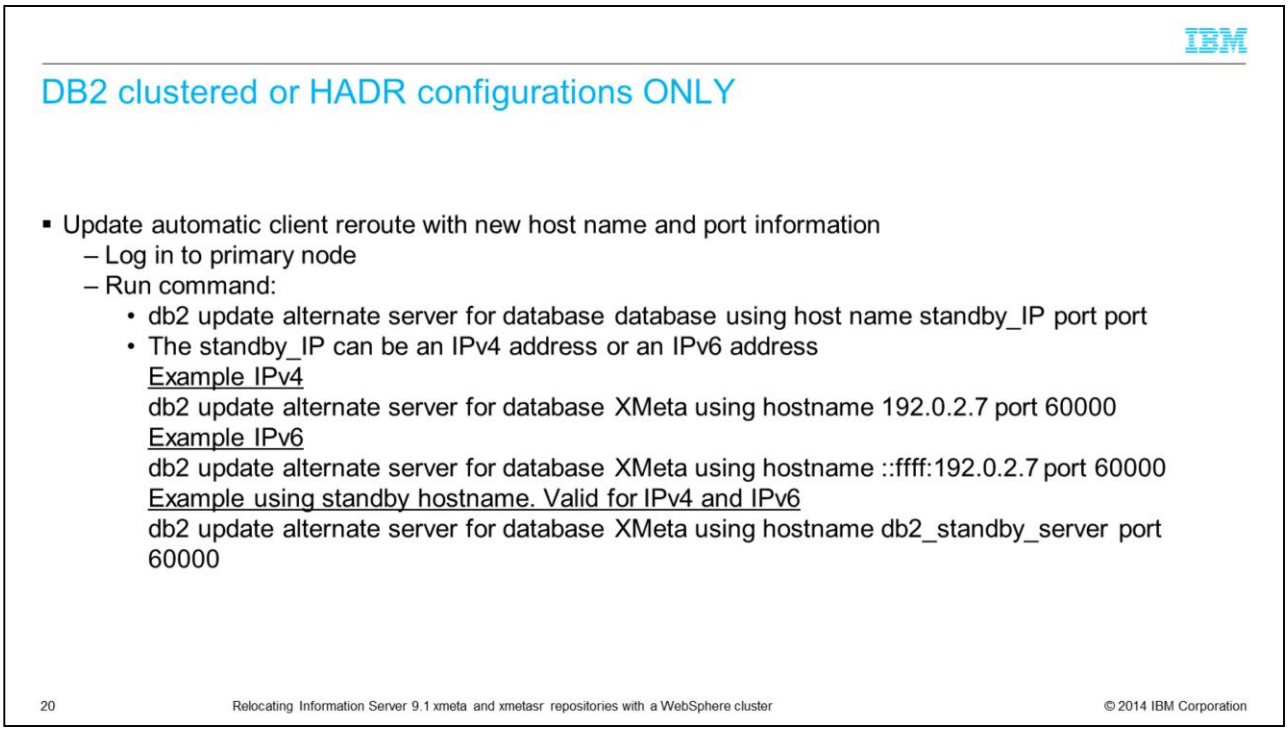

In an IBM InfoSphere Information Server installation with a clustered DB2 database system setup, you must update the automatic client reroute with the new host name and port. This slide displays the format of the update alternate server command along with examples of the command. If you are not using Information Server with a clustered DB2 database, skip this step.

Update WebSphere Application Server configuration (1 of 7)

- Log in to WebSphere Application Server administrative console
- Update all data sources that are highlighted in yellow

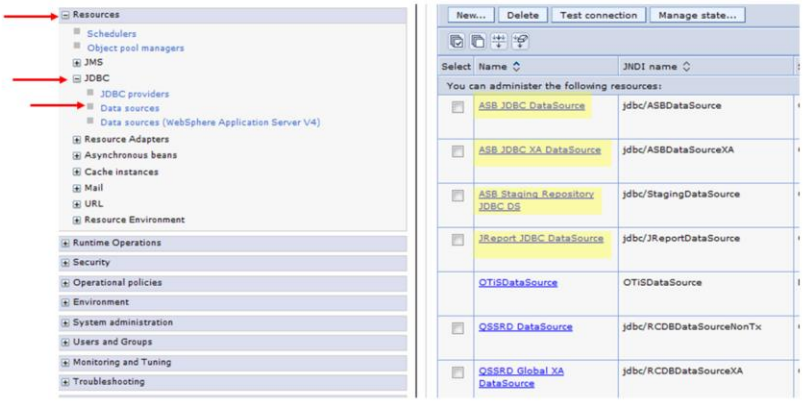

Log in to the WebSphere Application Server administrative console. Click the Resources tab on the left side and then click JDBC and then Data sources underneath. You need to complete the changes that are described in the next few slides on all of the data sources that are highlighted in yellow on this slide. To start, click the first data source, ASB JDBC DataSource. If you do not see the data sources as displayed on this slide, go to the Scope section, click the drop-down, and select All scopes. All of the data sources should now appear in the box.

Update WebSphere Application Server configuration (2 of 7)

- Modify connection properties – DB2

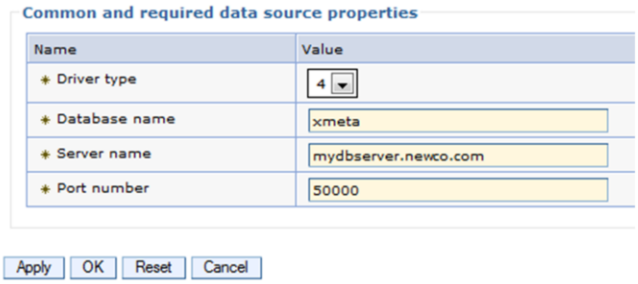

For xmeta on DB2, scroll to the bottom of the screen for the data source and modify the connection properties as required. Click Apply.

Update WebSphere Application Server configuration (3 of 7)

- Specifying the alternate address and port for a DB2 clustered database and DB2 HADR
  - From the data source properties page that you are editing
    - Click Connection Pool Properties under Additional Properties
    - Click Connection pool custom properties under additional properties
    - Update
      - clientRerouteAlternateServerName
      - clientRerouteAlternatePort

If you are using a DB2 cluster or HADR for xmeta, you need to update the alternate address and port. On the data source page that is seen in the previous slide, click the Connection pool properties under Additional Properties on the upper right side of the screen. Next, click Connection poll custom properties under Additional Properties. Update the clientRerouteAlternateServerName and clientRerouteAlternatePort properties and click Apply.

Update WebSphere Application Server configuration (4 of 7)

- Modify connection properties – Oracle

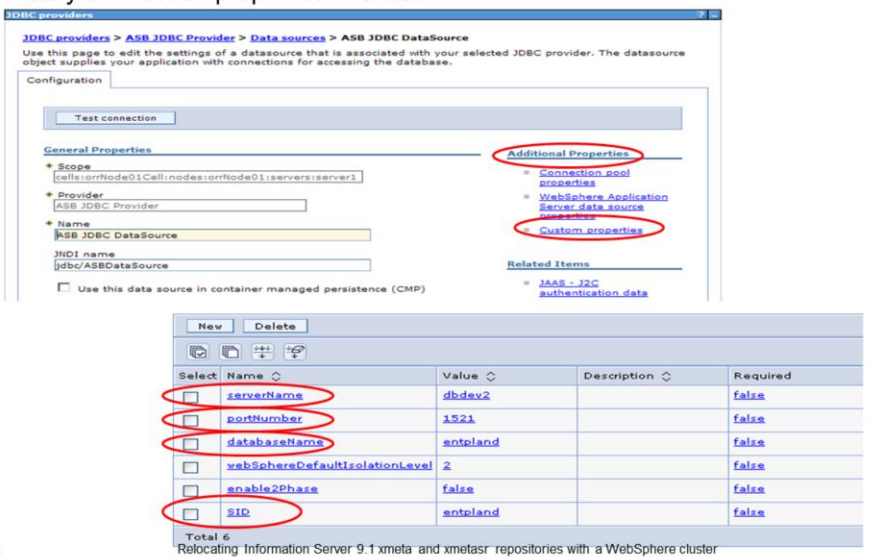

For xmeta on Oracle, on the Configuration tab, click Custom Properties on the right side under the Additional Properties heading. Once in Custom Properties, click the settings that have changed and set them to the appropriate values. Click Apply.

Update WebSphere Application Server configuration (5 of 7)

- Modify connection properties – Oracle RAC

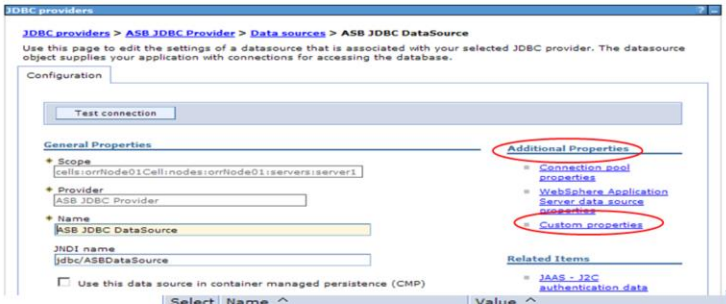

Select	Name	Value	Description	Required
You can administer the following resources:				
	serverName	rac1		false
	portNumber	1521		false
	webSphereDefaultIsolationLevel	2		false
	enable2Phase	false		false
	serviceName	orcl		false
	alternateServers	(rac1:1521,rac2:1521,rac3:1521)		false

This slide displays an example of an Oracle RAC configuration. Click the values that have changed and set the values appropriately.

Update WebSphere Application Server configuration (6 of 7)

- Modify connection properties – SQL Server

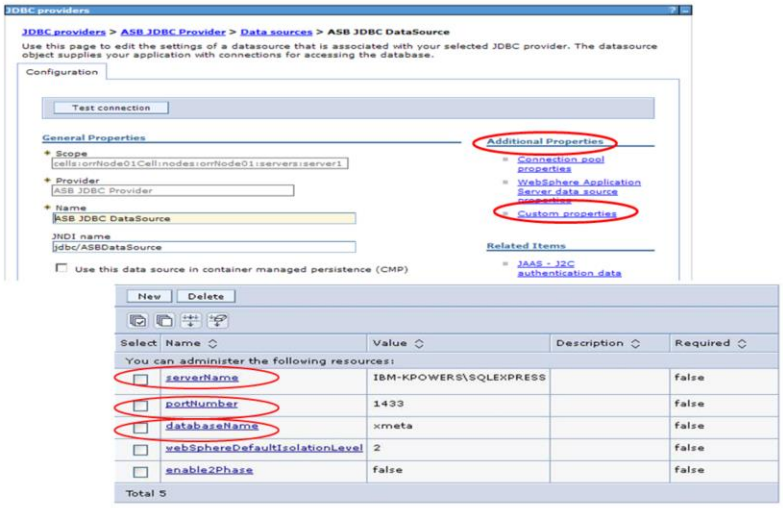

This slide displays an example on SQL Server. Click the settings that have changed and set them to the appropriate values. If you are using SQL Server named instances, be sure that your server name is in the format of servername\instancename as displayed in the example on this slide. Click Apply.

Update WebSphere Application Server configuration (7 of 7)

In the Messages box at the top of the page, click Save to save to the master configuration.

Test connection

- Test connection

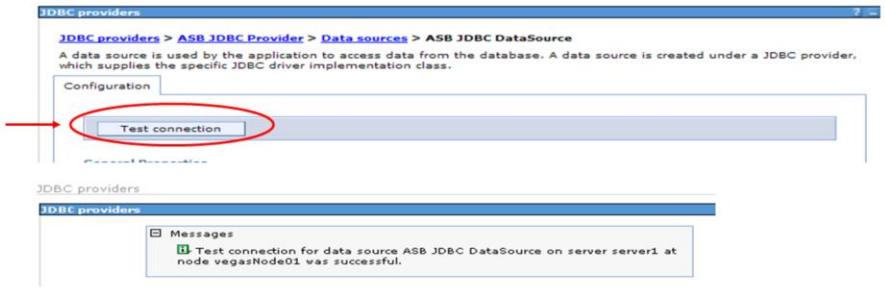

- Restart
  - WebSphere cluster members

28 Relocating Information Server 9.1 xmeta and xmetasr repositories with a WebSphere cluster © 2014 IBM Corporation

After you save the changes, test your new connection by clicking the Test connection button at the top of the Data Sources page where you made the server changes. If the connection is successful, you see a message at the top of the screen indicating success. If it is unsuccessful, go back and check the modified data source settings.

Once the test completes successfully, go back and modify the remaining data sources in the same manor. After you make, save, and successfully test all the changes, restart the WebSphere cluster members.

Relocating more repositories

- Follow IBM Education Assistant modules for more repository and database relocation

If you need to relocate more repositories or databases, see the IBM Education Assistant module for each repository or database.

Trademarks, disclaimer, and copyright information