



Migrating WebSphere Process Server



Migrating WebSphere Process Server

Note

Before using this information, be sure to read the general information in the Notices section at the end of this document.

31 March 2008

This edition applies to version 6, release 1, modification 0 of WebSphere Process Server for Multiplatforms (product number 5724-L01) and to all subsequent releases and modifications until otherwise indicated in new editions.

To send us your comments about this document, send an e-mail message to doc-comments@us.ibm.com. We look forward to hearing from you.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2006, 2008. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Migrating from previous versions of WebSphere Process Server and WebSphere Enterprise Service Bus 1

Overview of migrating	1
Premigration considerations	2
Tools for version-to-version migration	6
How data is handled during migration from earlier versions	19
Configuration mapping during product-configuration migration	21
Migrating WebSphere applications.	26
Upgrading databases for migration	27
Upgrading the Common database manually	27
Upgrading the Business Process Choreographer database manually	29
Migrating stand-alone servers	33
Migrating a stand-alone server using the migration wizard	33
Migrating to a stand-alone server using command-line tools.	37
Migrating a network deployment environment	40
Migrating a deployment manager	40
Migrating non-clustered managed nodes	49
Migrating clusters	60
Migrating to a remote system	79
Migrating from an operating system that is no longer supported	83
Verifying migration.	87
Postmigration configuration checking.	88
Rolling back your environment.	88
Rolling back a deployment cell	89
Rolling back a managed node	91
Migrating Cloudscape databases	94
Verifying the Cloudscape v10.1.x automatic migration	96
Upgrading Cloudscape manually	99
Migrating the UDDI registry	103
Migration considerations for Business Process Choreographer	105
Troubleshooting version-to-version migration.	107

Chapter 2. Migrating from previous WebSphere products 117

Migrating from WebSphere InterChange Server	117
Premigration considerations	118
Migrating WebSphere InterChange Server artifacts with the reposMigrate command	124
Postmigration considerations	128
Support for WebSphere Business Integration data handlers	142
Supported WebSphere InterChange Server APIs	143
Limitations when migrating from WebSphere InterChange Server	163
Troubleshooting migration from WebSphere InterChange Server	165
Migrating source artifacts from WebSphere Studio Application Developer Integration Edition.	167
Migrating from WebSphere MQ Workflow.	167

Chapter 3. Deprecated features. 169

Deprecation list.	169
Deprecated features in WebSphere Process server version 6.1	169
Deprecated features in WebSphere Process Server version 6.0.2	173
Deprecated features in WebSphere Process Server version 6.0.1	175
Deprecated features in WebSphere Process Server version 6.0	175
Deprecated features in WebSphere Business Integration Server Foundation version 5.1.1	178
Deprecated features in WebSphere Business Integration Server Foundation version 5.1	178

Chapter 4. Troubleshooting migration 181

Troubleshooting version-to-version migration.	181
Troubleshooting migration from WebSphere InterChange Server	189
Enabling logging and tracing for supported WebSphere InterChange Server APIs	189
Failure trying to serialize an object that is not serializable in a migrated BPEL file	190

Notices 193

Chapter 1. Migrating from previous versions of WebSphere Process Server and WebSphere Enterprise Service Bus

You can migrate your installed applications and configurations from prior versions of WebSphere Process Server and WebSphere® Enterprise Service Bus to WebSphere Process Server version 6.1.

Overview of migrating

Migrate from earlier versions of WebSphere Process Server and WebSphere Enterprise Service Bus.

Moving from one version of WebSphere Process Server to a newer release of WebSphere Process Server, or, in some cases, moving from a version of WebSphere Enterprise Service Bus to a higher release level of WebSphere Process Server is referred to as version-to-version migration. Version-to-version migration occurs when you install a new version of a product, such as WebSphere Process Server, and then copy relevant application and configuration data from the old installation to the new installation. With migration, the new version of the product is installed alongside the older product. Then, data is copied from the old version of the product to the new version of the product. Migrating is different from updating, in which out-of-date files or data of an existing installation are replaced with current information. Refresh packs, interim fixes, and fix packs are examples of updates. For more information on updating, see *Installing fix packs and refresh packs with the Update Installer*.

Migration must take place from an older version of WebSphere Process Server to a newer version that is running on the same operating system. You cannot migrate across different operating systems. For example, if your WebSphere Process Server version 6.0.x is running on Microsoft Windows XP, you can migrate the data from that instance of WebSphere Process Server to WebSphere Process Server version 6.1 running on Windows XP. However, you could not migrate data from WebSphere Process Server version 6.0.x running on Windows XP to WebSphere Process Server version 6.1 running on AIX. In the case of a standalone server only, you can migrate from a prior release of a given operating system to a new supported release of the same operating system. For example, for a standalone server you could migrate WebSphere Process Server version 6.0.x running on AIX 5.2 to WebSphere Process Server version 6.1 running on AIX 5.3. (Refer to “Migrating from an operating system that is no longer supported” on page 83 for instructions on such a migration.)

The following table shows the supported version-to-version migration scenarios for this release of WebSphere Process Server. You can migrate all of the products listed under “Currently installed version” to WebSphere Process Server version 6.1.

Currently installed version	New version
WebSphere Process Server 6.0.1.x	WebSphere Process Server 6.1
WebSphere Process Server 6.0.2.x	WebSphere Process Server 6.1
WebSphere Enterprise Service Bus 6.0.1.x	WebSphere Process Server 6.1
WebSphere Enterprise Service Bus 6.0.2.x	WebSphere Process Server 6.1

Why perform version-to-version migration?

WebSphere Process Server provides user application binary compatibility with prior versions. However, version-to-version migration allows you to preserve WebSphere Process Server configuration data in addition to your applications when moving to a newer version of WebSphere Process Server. Configuration of profiles, cells, clusters, servers and nodes are retained when you perform a version-to-version migration. If you did not perform this migration, and simply installed the new version of WebSphere Process Server, you would need to reconfigure your environment from scratch.

For some releases of WebSphere Process Server, an "in-place update" or maintenance package is available. Such an update will also preserve configuration data. In cases for which no maintenance package is available, such as when moving from WebSphere Process Server version 6.0.x to version 6.1, a version-to-version migration is required to preserve your configuration data.

Related concepts



Development and deployment version levels

Your decision about what version levels of WebSphere Process Server you need in your environment will depend on the version levels with which your applications were developed. Generally applications deployed in a previous version of WebSphere Process Server will run on the next available version of WebSphere Process Server.

Related tasks



Installing fix packs and refresh packs with the Update Installer

You can use the IBM® Update Installer for WebSphere Software to install interim fixes, fix packs, and refresh packs collectively known as maintenance packages. The Update Installer for WebSphere Software is also known as the update installer program, the UpdateInstaller program, and the Update Installation Wizard.

Premigration considerations

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

The following rules and restrictions apply to migration and coexistence if you have WebSphere Process Server version 6.1 installed:

- You can migrate a version 6.0.x deployment manager to a version 6.1.x deployment manager only if they are at the same augmentation level.
- You cannot augment a version 6.1.x stand-alone server that has been migrated from a version 6.0.x stand-alone server.

You can create a new stand-alone profile in version 6.1.x and augment it.

- You cannot augment a version 6.1.x managed server that has been migrated from a 6.0.x managed server.

You can create a new profile in version 6.1.x, augment it, and then add the new node to a version 6.1.x cell that has an augmented deployment manager.

- You can have a mixed cell containing both augmented and unaugmented managed nodes as long as the cell's deployment manager has been augmented to the same augmentation level as the highest augmentation level of any of its managed nodes. For example, if the deployment manager is augmented for WebSphere Process Server, it can successfully manage nodes that have been

augmented for WebSphere Process Server and WebSphere Application Server. However, a deployment manager that has only been augmented for WebSphere Application Server can only manage WebSphere Application Server nodes.

- If you have Business Process Choreographer installed, see “Migration considerations for Business Process Choreographer” on page 105.
- After you have installed WebSphere Process Server version 6.1, you might want to build a complete deployment cell configuration and make sure that it works properly before you attempt to migrate an existing cell or node.

This ensures that your system has all of the necessary prerequisites and supports the new level of WebSphere Process Server.

- High availability manager and core group functionality are included in WebSphere Process Server version 6.0 and later. See Core group migration considerations for core group configuration and topology considerations that might impact your migration from 6.0.x to version 6.1.
- Before you migrate to JDK 5 (introduced in WebSphere Application Server version 6.1 and therefore WebSphere Process Server version 6.1) from JDK 1.4, review your applications for necessary changes based on the Sun Microsystems Java™ specification.

See API and specification migration.

- When migrating a cell with multiple nodes, the applications must remain at the lowest JDK level until all nodes are migrated.

- **Solaris** Java Native Interface (JNI) applications that work with WebSphere Process Server version 6.0.2 on Solaris x64 must be recompiled in a 64-bit environment in order for them to work with WebSphere Process Server version 6.1. This includes all JNI applications that run in a WebSphere Process Server process—code called from an Enterprise JavaBean (EJB) for example.

On Solaris x64, WebSphere Process Server version 6.0.2 runs as a 32-bit application even though the underlying platform is 64-bit. This is because the underlying Java virtual machine is 32-bit. WebSphere Process Server version 6.1 runs as a 64-bit application because the underlying Java virtual machine is 64-bit. JNI applications compiled in a 32-bit environment for version 6.0.2 cannot run in the 64-bit environment of version 6.1.

- WebSphere Process Server version 6.1 can be installed in an environment where it coexists with prior levels of WebSphere Process Server. However, some restrictions exist.

For information about coexistence, including restrictions, see Coexisting with other WebSphere product installations.

- In a cluster, version 6.0.x members and version 6.1 members must never run at the same time. All version 6.0.x cluster members must be stopped before you start the first version 6.1 cluster member. Also, once you start a version 6.1 cluster member, do not start any version 6.0.x cluster members in that cluster.
- During migration, version 6.1 cluster information is distributed throughout the cell. Version 6.0.x nodes that are not at WebSphere Process Server version 6.0.1.3 or later fail to read this information and the cluster function might fail. There are two options for nodes that are not at version 6.0.1.3. One option is to upgrade all version 6.0.x nodes that will be contained in or interoperating with a version 6.1 cell to WebSphere Process Server version 6.0.1.3 or later before migrating your deployment managers to version 6.1. The second option, if you do not wish to move the WebSphere Process Server version to 6.0.1.3 or later, is to update the underlying WebSphere Application Server to version 6.0.2.11 or higher.

- WebSphere Process Server version 6.1 migration converts HTTP transports to channel-framework Web container transport chains.
For more information on version 6.1 transport support, see the following topics:
 - Configuring transport chains
 - HTTP transport channel settings
 - Transport chains
- When migrating a deployment manager, the WebSphere Process Server version 6.1 cell name must match the version 6.0.x cell name.
If you create a profile with a new cell name and use this profile as a migration target, the migration will fail.
- When migrating a federated node, the WebSphere Process Server version 6.1 cell name and node name must match their respective version 6.0.x names.
- If you create a profile that does not meet the migration requirements (such as naming requirements), you can remove the old profile and create a new one rather than uninstalling and reinstalling the WebSphere Process Server version 6.1 product.
- If you migrate a node to WebSphere Process Server version 6.1 then discover that you need to revert back to version 6.0.x, see “Rolling back your environment” on page 88.
- The migration tools create a migration backup directory containing a backup copy of the configuration from the previous version. If you are migrating from version 6.0.x, the space available for this directory should be at least the size of the previous profile’s configuration directory and applications.
- The amount of storage that your system requires during migration to version 6.1 depends on your environment as well as on which migration tool you are using.
 - **WBIPreUpgrade storage requirements**
 - **Location:** Backup directory specified as a parameter of the WBIPreUpgrade command
 - **Amount:** For a rough estimate of your storage requirements when using this command, add the following amounts.
 - Size of the following items for all of the profiles in your old configuration:
 - *profile_root*/installableApps directory
 - *profile_root*/installedApps directory
 - *profile_root*/config directory
 - *profile_root*/properties directory
 - Shared libraries referenced in the libraries.xml configuration files
 - Resource Adapter Archive (RAR) files referenced in the resources.xml configuration files
 - If trace is enabled, which is the default, up to 200 MB (depending on the size and complexity of your configuration)

For more information about this command, see the “WBIPreUpgrade command” on page 12.

 - **WBIPostUpgrade storage requirements**
 - **Location:** New configuration relative to the new *profile_root* directory
 - **Amount:** For a rough estimate of your storage requirements when using this command, add the following amounts.
 - Size of the following items for the old profile that you are migrating:
 - *profile_root*/installableApps directory

- *profile_root*/installedApps directory
- *profile_root*/config directory
- *profile_root*/properties directory
- Shared libraries referenced in the libraries.xml configuration files
- Resource Adapter Archive (RAR) files referenced in the resources.xml configuration files
- If trace is enabled, which is the default, up to 200 MB (depending on the size and complexity of your configuration)

For more information about this command, see the “WBIPostUpgrade command” on page 14.

- Before you migrate a Cloudscape™ database, ensure that any servers hosting applications that are using the Cloudscape database are shut down. Otherwise, the Cloudscape migration will fail.
- After you use the migration tools to migrate to WebSphere Process Server version 6.1, you might need to do some things that are not done automatically by the migration tools.
 - Examine any Lightweight Third Party Authentication (LTPA) security settings that you might have used in version 6.0.x, and make sure that version 6.1 security is set appropriately.
See Lightweight Third Party Authentication.
 - Check the WBIPostUpgrade.log file in the logs directory for details about any JSP objects that the migration tools did not migrate.
If version 6.1 does not support a level for which JSP objects are configured, the migration tools recognize the objects in the output and log them.
 - Review your Java virtual machine settings to verify that you are using the recommended default values as described in
Java virtual machine settings.
 - Verify the results of the automatic Cloudscape database migration, and manually migrate any Cloudscape databases that are not automatically migrated by the tools.
See “Migrating Cloudscape databases” on page 94.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Related tasks

“Rolling back your environment” on page 88

After migrating to a WebSphere Process Server version 6.1 environment, you can roll back to a version 6.0.x environment. This returns the configuration to the state that it was in before migration. After rolling back the environment, you can restart the migration process.

“Migrating Cloudscape databases” on page 94

After you use the migration tools to migrate to WebSphere Process Server version 6.1, you should verify the results of the automatic Cloudscape database migration and manually migrate any Cloudscape database instances that are not automatically migrated by the tools.

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.



WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information



Coexisting with other WebSphere product installations

An installation of WebSphere Process Server, version 6.1 can coexist on the same system with installations of any version of WebSphere Process Server or WebSphere Enterprise Service Bus, and with certain versions of selected WebSphere products.

Configuring transport chains

HTTP transport channel settings

Transport chains

API and specification migration

Creating clusters

Creating application servers

Lightweight Third Party Authentication

Core group migration considerations

Java virtual machine settings

Tools for version-to-version migration

You can perform migration from earlier versions of WebSphere Process Server or WebSphere Enterprise Service Bus by using a migration “wizard” (graphical user interface) or a series of commands or scripts.

To migrate, you can use the migration wizard or a series of commands.

Migration wizard

The migration wizard steps you through the migration process. You will be asked to fill in certain fields and make some choices provided by the wizard or use the defaults. You can invoke the migration wizard from the WebSphere Process Server First Steps panel. You can also invoke it directly from *install_root\bin\wbi_migration.bat* (on Windows® systems) or *install_root/bin/wbi_migration.sh* (on UNIX-based systems).

Migration commands

You can use a series of scripts invoked from the command line to migrate from an earlier version of WebSphere Process Server if you choose not to use the migration wizard. The following commands must be run in the correct sequence:

1. **WBIPreUpgrade** - This command, which you run first, saves the existing WebSphere Process Server configuration and applications into a migration-specific backup directory.
2. **WBIPostUpgrade** - This command, which you run second, processes the contents of the migration-specific backup directory that was created with the **WBIPreUpgrade** command and imports it into the new WebSphere Process Server environment.
3. **WBIProfileUpgrade.ant** - This script needs to be run manually after running the **WBIPreUpgrade** and **WBIPostUpgrade** commands only if you are migrating clustered cells. You might also need to run it in special cases in which migration did not occur successfully. The script updates enterprise applications and configuration settings in a profile. The **WBIPostUpgrade** command invokes this script and for non-clustered scenarios does not have to be manually re-run a second time.

Database upgrade scripts

Normally, if updates are required to any databases supporting WebSphere Process Server components, the update is completed automatically when the migrated server process is started. However, if the server process does not have sufficient permissions (that is, if it has not been configured with the correct user ID with sufficient permissions for the applicable database), then you must update the databases manually using scripts provided.

Related tasks

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

Migration wizard

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

The version-to-version migration wizard is the graphical interface to the primary WebSphere Process Server migration tools: the **WBIPreUpgrade** command and the **WBIPostUpgrade** command.

Note: The migration wizard cannot run in a non-graphical environment. Examples of non-graphical environments include telnet sessions. If you want to run migration in a non-graphical environment, use the **WBIPreUpgrade** and **WBIPostUpgrade** commands.

What the migration wizard does

The migration wizard uses the **WBIPreUpgrade** command and **WBIPostUpgrade** command to migrate the data and applications from the older version to the newer version of WebSphere Process Server.

This step transfers applications and configuration information for the older version’s server resources, security, variables, and virtual hosts to the newer version’s server. All stored information is in XML files in the *profile_dir/config/cells* directory of each product.



The WBIPreUpgrade tool saves selected files from the *install_root* and *profile_root* directory to a backup directory you specify on a wizard panel. Migration saves files to the following subdirectories in the backup directory:

- *websphere_backup*
 - bin
 - cloudscape
 - lib
 - ProcessChoreographer/Staff
 - properties
- *profile_root*
 - bin
 - config
 - etc
 - event
 - installedApps
 - installedConnectors
 - properties

Later, the migration wizard uses the WBIPostUpgrade tool to restore the environment in the backup directory into the newer WebSphere Process Server installation.

Accessing the migration wizard

Invoke the migration wizard in one of the following ways:

- From the WebSphere Process Server First Steps console, select **Migration wizard**.
- Run one of the following scripts (depending upon your operating system) stored in the *install_dir/bin* directory:
 1.  On UNIX[®] and Linux[®] systems: `wbi_migration.sh`
 2.  On Windows systems: `wbi_migration.bat`

Before you begin

The migration wizard prompts you for information as you proceed. Before you invoke the wizard, collect the following information:

Cell name

(Required for migration of deployment managers only.) Name of the cell managed by the deployment manager that you are migrating. The new version cell name must match the name in the old version's configuration.

Installation root directory

See "WBIPreUpgrade command" on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See "WBIPreUpgrade command" on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See "WBIPostUpgrade command" on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

“WBIPostUpgrade command” on page 14

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Running the migration wizard: Procedure

1. Read the Welcome panel to learn about the migration process, and then click **Next**.
2. Select or specify a previous version of WebSphere Process Server from which to migrate, and then click **Next**.
Select the check box and enter the location of the previous installation if it does not appear in the selection list.
3. Select the previous version’s profile that you want to migrate, and then click **Next**.
4. Select the profile from the list of valid profiles for the version 6.1 installation or select **Create new profile**, and then click **Next**.
Select the check box to create a backup copy of the new profile’s configuration before migrating the old profile. If you select the check box, the backup copy of the target profile will be written to *profile_root/temp/MigrationBackup.time_stamp.zip*.
5. If you selected **Create new profile** on the last panel, enter a name for the profile and a host name (name of the system on which the profile will reside) and then click **Next**. If possible, use the same host name that was used for the older profile. Otherwise, after migration, you will need to use the administrative console to update the host name manually in other places where it is configured.
6. Specify a migration backup directory in which to place a backup copy of the configuration from the previous version, and then click **Next**.
The directory is created if it does not already exist. If the directory exists, it should be empty because the backup operation might overwrite existing backup files.
7. Specify where the migrated applications should be located and then click **Next**.
You can choose any one of the following options:
 - Keep the applications in the same directories in which they are currently located.

Restrictions: If you choose this option, the location is shared by the existing installation and the new installation. If you keep the migrated applications in the same locations as those of the previous version, the following restrictions apply:

- Mixed-node support limitations must be followed. This means that the following support cannot be used when evoking the wsadmin command:
 - Precompile JSP
 - Use Binary Configuration
 - Deploy EJB
 - You risk losing the migrated applications unintentionally if you later delete applications from these locations when administering (uninstalling for example) your version 6.0.x installation.
 - Choose to install the applications in the default directory of the target version.
 - Specify the directory in which to install the migrated applications.
8. Select one of the options for assigning port values, and then click **Next**. You can choose to do any one of the following with the port values:
- Use the port values assigned to the previous installation.
 - Use the port values assigned to the target profile.
 - Define the port values as a block.
- If you select this option, specify the first value of the block of consecutive port numbers to assign.
9. Select the check box if you want to migrate to support script compatibility, and then click **Next**.

If you select this option, the migration wizard creates the following configuration definitions of the version 6.0.x of WebSphere Process Server:

- Transport
- ProcessDef
- version 6.0.x SSL
- version 6.0.x ORB service threadpool

instead of the following version 6.1 configuration definitions:

- Channels
- ProcessDefs
- version 6.1 SSL
- version 6.1 ORB service threadpool

Select this option in order to minimize impacts to existing administration scripts. If you have existing **wsadmin** scripts or programs that use third-party configuration APIs to create or modify your existing configuration definitions, for example, you might want to select this option during migration.

Note: This is meant to provide a temporary transition until all of the nodes in the environment are at the newer level (for example, WebSphere Process Server version 6.1.) When they are all at the new level, you should perform the following actions:

- a. Modify your administration scripts to use all of the settings for the new version (in this case version 6.1).
- b. Use the `convertScriptCompatibility` command to convert your configurations to match all of the settings corresponding to the new version.

See the `convertScriptCompatibility` command.

Note: When following the directions at this link to use the `convertScriptCompatibility` command, use the `WBIPostUpgrade` command rather than the `WASPostUpgrade` command.

10. Check the information in the summary panel and make sure that it is correct, and then click **Next** to start the migration.
If you selected the option to create a new target profile, panels show the beginning and results of that creation.
Panels show the progress of the migration process.
If the migration is not successful, the wizard displays a failure panel. If the migration is partially successful, the wizard displays a warning panel. Correct any problems and retry the migration.
If the postmigration is successful, the wizard displays an indication of success.
11. Click **Finish** to exit the migration wizard.

Results

You can now start the migrated server in the WebSphere Process Server environment at the new release level.

Related concepts

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

 Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

Related reference

“WBIPreUpgrade command” on page 12

Use the `WBIPreUpgrade` command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

 `WBIPostUpgrade` command

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the `backupDirectory` that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the `profile_root/installedApps` directory for the new installation.

Related information

Java virtual machine settings

`convertScriptCompatibility` command

WBIPreUpgrade command

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

Details

The WBIPreUpgrade command saves selected files from the *install_root* and *profile_root* directories to a backup directory you specify. Migration saves files to the following subdirectories in the backup directory:

- *websphere_backup*
 - bin
 - cloudscape
 - lib
 - ProcessChoreographer/Staff
 - properties
- *profile_root*
 - bin
 - config
 - etc
 - event
 - installedApps
 - installedConnectors
 - properties

Location

The command file is located in, and should be run from, the *install_dir/bin* directory.

Syntax

The syntax is as follows: UNIX Linux

```
WBIPreUpgrade.sh backupDirectory  
                  currentWebSphereDirectory  
                  [-traceString trace_spec [-traceFile file_name ]]
```

Windows

```
WBIPreUpgrade.bat backupDirectory  
                  currentWebSphereDirectory  
                  [-traceString trace_spec [-traceFile file_name ]]
```

Parameters

The parameters are as follows:

backupDirectory

This is a required parameter and must be the first parameter that you specify. The value *backupDirectory* specifies the name of the directory where the command script stores the saved configuration.

This is also the directory from which the WBIPostUpgrade command reads the configuration.

If the directory does not exist, the `WBIPreUpgrade` command script creates it.

currentWebSphereDirectory

This is a required parameter and must be the second parameter that you specify. This can be any edition of WebSphere Process Server for which migration is supported.

The value *currentWebSphereDirectory* specifies the name of the installation root directory for the existing WebSphere Process Server installation.

-traceString

This is an optional parameter. The value *trace_spec* specifies the trace information that you want to collect.

To gather all trace information, specify `"*=all=enabled"` (with quotation marks).

If you do not specify the `-traceString` or `-traceFile` parameter, the command creates a trace file by default and places it in the *backupDirectory/logs* directory.

If you specify this parameter, you must also specify the `-traceFile` parameter.

-traceFile

This is an optional parameter. The value *file_name* specifies the name of the output file for trace information.

If you do not specify the `-traceString` or `-traceFile` parameter, the command creates a trace file by default and places it in the *backupDirectory/logs* directory.

If you specify the `-traceString` parameter but do not specify the `-traceFile` parameter, the script does not generate a trace file.

Logging

The `WBIPreUpgrade` tool displays status to the screen while it runs. The tool also saves a more extensive set of logging information in the `WBIPreUpgrade.time_stamp.log` file written to the *backupDirectory/logs* directory, where *backupDirectory* is the value specified for the `backupDirectory` parameter. You can view the `WBIPreUpgrade.time_stamp.log` file with a text editor.

Related concepts

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related reference

“WBIPostUpgrade command” on page 14

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the *backupDirectory* that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

“WBIProfileUpgrade script” on page 18

Use the `WBIProfileUpgrade` script to update application and configuration settings in a WebSphere Process Server profile when you are migrating clusters and in some other special situations.

Related information

`clientUpgrade` command

WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Location

The command file is located and should be run in the *install_root/bin* directory.

Syntax

The syntax is as follows:

UNIX

Linux

```
WBIPostUpgrade.sh backupDirectory
    [-oldProfile profile_name]
    [-profileName profile_name]
    [-scriptCompatibility true | false]
    [-portBlock port_starting_number]
    [-backupConfig true | false]
    [-replacePorts true | false]
    [-keepAppDirectory true | false]
    [-keepDmgrEnabled true | false]
    [-appInstallDirectory user_specified_directory]
    [-traceString trace_spec [-traceFile file_name]]
```

Windows

```
WBIPostUpgrade.bat backupDirectory
    [-oldProfile profile_name]
    [-profileName profile_name]
    [-scriptCompatibility true | false]
    [-portBlock port_starting_number]
    [-backupConfig true | false]
    [-replacePorts true | false]
    [-keepAppDirectory true | false]
    [-keepDmgrEnabled true | false]
    [-appInstallDirectory user_specified_directory]
    [-traceString trace_spec [-traceFile file_name]]
```

Parameters

The parameters are as follows:

backupDirectory

This is a required parameter. The value *backupDirectory* specifies the name of the directory in which the WBIPreUpgrade tool stored the saved configuration and files and from which the WBIPostUpgrade tool reads the configuration and files.

-oldProfile

This is an optional parameter for migrating instances or profiles from previous versions. The instance or profile must already exist in the migration backup directory before running this command.

If the *-oldProfile* parameter is not specified, the default profile will be used. If no default profile is found, the system will report an error.

-profileName

This is an optional parameter for migrating to specific profiles. The value

profile_name specifies the name of the profile, already created in the more recent version of WebSphere Process Server, to which the script migrates your configuration. You must have already created this profile before calling the `WBIPostUpgrade` command.

If the `-profileName` parameter is not specified, the default profile will be used. If no default profile is found, the system will report an error.

-scriptCompatibility

This is an optional parameter used to specify whether or not migration should create the following version 6.0.x configuration definitions:

- Transport
- ProcessDef
- SSL for version 6.0.x

instead of the following version 6.1 configuration definitions:

- Channels
- ProcessDefs
- SSL for version 6.1

The default is true.

Specify true for this parameter in order to minimize impacts to existing administration scripts. If you have existing `wsadmin` scripts or programs that use third-party configuration APIs to create or modify the version 6.0.x configuration definitions, for example, you might want to specify true for this option during migration.

Note: This is meant to provide a temporary transition until all of the nodes in the environment are at the version 6.1 level. When they are all at the new level, you should perform the following actions:

1. Modify your administration scripts to use all of the version 6.1 settings.
2. Use the `convertScriptCompatibility` command to convert your configurations to match all of the version 6.1 settings.

For more information see the `convertScriptCompatibility` command .

-portBlock

This is an optional parameter. The *port_starting_number* value specifies the first of a block of consecutive port numbers to assign when the command script runs.

-backupConfig

This is an optional parameter used to specify whether the existing configuration is saved before any changes are made by the `WBIPostUpgrade` tool. The default is true—that is, `WBIPostUpgrade` saves a copy of the current configuration into the *profile_name/temp* directory.

Use the `restoreConfig` command to restore that configuration as required. See the `restoreConfig` command.

-replacePorts

This optional parameter is used to specify how to map port values for virtual hosts and web-container transport ports.

- False

Do not replace the version 6.0.x port definitions during migration.

- The previous version’s configuration is left alone – no channels are deleted.

- The following four named channels are set to values that are equivalent to the values set for the previous release:
 - WC_adminhost
 - WC_defaulthost
 - WC_adminhost_secure
 - WC_defaulthost_secure
- The migration process creates transports or channels, based on the `-scriptCompatibility` setting, for any ports in the previous release.
- The migration process sets all non-web-container ports to the values set for the previous release.
- Port conflicts might arise if the migration process creates a transport or channel that is the same as one defined in the web container.

This is the default.

- True

Replace all virtual host alias port settings during migration with version 6.1 port definitions.

By default, the migration process adds configuration data from the previous environment to the data in the new WebSphere Process Server environment. In some cases, however, this might not be the desired behavior for these port values. For example, existing port definitions from the earlier release might have been carefully set to avoid port conflicts with other products; in such cases, it is likely that you would want to migrate these settings into the new version's configuration. Specify true for this parameter to cause any ports of matching virtual hosts to be removed from the newer version's configuration before the new values are added.

- All transport channels associated with the web container are deleted except for the following four named channels, which are set to values that are equivalent to the values set for the previous release:
 - WC_adminhost
 - WC_defaulthost
 - WC_adminhost_secure
 - WC_defaulthost_secure
- The migration process creates transports or channels, based on the `-scriptCompatibility` setting, for any ports in the previous release.
- The migration process sets all non-web-container ports to the values set for the previous release.

-keepAppDirectory

This is an optional parameter used to specify whether to install all applications to the same directories in which they are currently located. The default is false.

If this parameter is specified as true, each individual application retains its location.

Restrictions: If this parameter is specified as true, the location is shared by the existing WebSphere Process Server and the new installation. If you keep the migrated applications in the same locations as those of the previous version, the following restrictions apply:

- The version 6.1 mixed-node support limitations must be followed. This means that the following support cannot be used when evoking the `wsadmin` command:
 - Precompile JSP

- Use Binary Configuration
- Deploy EJB
- You risk losing the migrated applications unintentionally if you later delete applications from these locations when administering (uninstalling for example) your previously existing installation.

-keepDmgrEnabled

This is an optional parameter used to specify whether to disable the existing WebSphere Process Server deployment manager. The default is false.

If this parameter is specified as true, you can use the existing deployment manager while the migration is completed. It is only valid when you are migrating a deployment manager; it is ignored in all other migrations.

Caution: Use this parameter with care.

- The reason that existing WebSphere Process Server deployment manager configurations normally are stopped and disabled is to prevent multiple deployment managers from managing the same nodes. You must stop the existing deployment manager before you start using the new version's deployment manager. The most likely error conditions that will occur if this is not done are port conflicts when the second instance of the deployment manager is started.
- Specifying true for this parameter means that any configuration changes made in the old configuration during migration might not be migrated.

-appInstallDirectory

This is an optional parameter used to pass the directory name to use when installing all applications during migration. The default of *profile_name*\installedApps is used if this parameter is not specified.

Quotes must be used around the directory name if one or more blanks are in the name.

-traceString

This is an optional parameter. The value *trace_spec* specifies the trace information that you want to collect.

To gather all trace information, specify "**=all=enabled*" (with quotation marks).

If you do not specify the `-traceString` or `-traceFile` parameter, the command creates a trace file by default and places it in the *backupDirectory*/logs directory.

If you specify this parameter, you must also specify the `-traceFile` parameter.

-traceFile

This is an optional parameter. The value *file_name* specifies the name of the output file for trace information.

If you do not specify the `-traceString` or `-traceFile` parameter, the command creates a trace file by default and places it in the *backupDirectory*/logs directory.

If you specify the `-traceString` parameter but do not specify the `-traceFile` parameter, the script does not generate a trace file.

Logging

The `WBIPostUpgrade` tool displays status to the screen while running. This tool also saves a more extensive set of logging information in the `WBIPostUpgrade.time_stamp.log` file located in the *backupDirectory*/logs directory.

You can view the `WBIPostUpgrade.time_stamp.log` file with a text editor.

Related concepts

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related reference

“WBIPreUpgrade command” on page 12

Use the `WBIPreUpgrade` command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

“WBIProfileUpgrade script”

Use the `WBIProfileUpgrade` script to update application and configuration settings in a WebSphere Process Server profile when you are migrating clusters and in some other special situations.

Related information

`convertScriptCompatibility` command

`restoreConfig` command

WBIProfileUpgrade script

Use the `WBIProfileUpgrade` script to update application and configuration settings in a WebSphere Process Server profile when you are migrating clusters and in some other special situations.

Purpose

`WBIProfileUpgrade` is an Apache™ Ant script that updates support applications supplied with WebSphere Process Server and configuration settings in a WebSphere Process Server profile. It is normally invoked automatically by the migration wizard or the `WBIPostUpgrade` command. However, when you migrate a cluster you will need to invoke `WBIProfileUpgrade` manually once the migration wizard or the `WBIPostUpgrade` command has been run on the cluster members. In addition, you might need to invoke `WBIProfileUpgrade` when manually migrating internal support applications or in some situations in which an error has occurred during migration.

Syntax

The syntax is as follows:

UNIX

Linux

```
profile_root/bin/ws_ant.sh -f install_root/util/WBIProfileUpgrade.ant  
-DmigrationDir=backupDirectory [-Dcluster=clustername ]
```

Windows

```
profile_root\bin ws_ant.bat -f install_root/util/WBIProfileUpgrade.ant  
-DmigrationDir=backupDirectory [-Dcluster=clusterName ]
```

Note: On Windows systems, specify paths for the parameters using either a single forward slash, such as `c:/migration/backup` or a double backslash, such as `c:\\migration\\backup`.

Parameters

profile_root

The directory in which the profile being migrated resides. The `ws_ant` command resides in the `bin` subdirectory of this profile directory.

install_root

The directory in which WebSphere Process Server is installed.

backupDirectory

This is a required parameter. The *backupDirectory* specifies the name of the directory in which the migration-specific backup directory created by the `WBIPreUpgrade` command resides.

clusterName

This specifies the name of the cluster that needs to be migrated.

Logging

The `WBIPreUpgrade` script creates a log when it runs that is written to the *backupDirectory/logs/WBIPreUpgrade.ant-time_stamp.log* directory, where *backupDirectory* is the value specified for the `backupDirectory` parameter. You can view the `WBIPreUpgrade.ant-time_stamp.log` file with a text editor.

Examples

Note: The following examples are each single command lines but are shown on multiple lines for readability purposes only.

UNIX

Linux

```
/opt/IBM/WPS61/profiles/DMgr/bin>ws_ant.sh -f
../util/WBIPreUpgrade.ant -DmigrationDir=/tmp/migrationBackup
-Dcluster=clusterA
```

Windows

```
C:\IBM\WPS61\profiles\DMgr\bin>ws_ant.bat -f
..\util\WBIPreUpgrade.ant -DmigrationDir=c:\\temp\\migrationBackup
-Dcluster=clusterA
```

Related reference

“`WBIPreUpgrade` command” on page 12

Use the `WBIPreUpgrade` command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

“`WBIPostUpgrade` command” on page 14

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the *backupDirectory* that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

How data is handled during migration from earlier versions

The WebSphere Process Server version-to-version migration tools will handle different sets of data (enterprise application data, configuration data, and system application data) in different ways.

Configuration data migration

The version-to-version migration tools (wizard or scripts) will automatically apply the configuration settings from the previous profile to the new profile created during the migration process. In cases in which the new profile has already been configured and values in the old profile and new profile do not match, the values will be handled as follows:

- The installation directory name that has already been configured in the new profile will be retained in the new profile.
- Any values from the old profile (other than the installation directory name) will replace non-matching values in the new profile.

Application migration

Your user applications (any applications not provided with the WebSphere Process Server product) are binary-compatible for the migration scenarios supported. (Refer to “Overview of migrating” on page 1 for supported migration scenarios.) All user applications will be automatically migrated to the new server. You should not have to modify any part of the application to have it run on the newer version of WebSphere Process Server.

Note: For version 6.0.1 WebSphere Adapters, some additional steps might be required for compatibility. For more information on this or any other exception, see the WebSphere Process Server technotes at the WebSphere Process Server technotes web site.

Except for sample applications, applications that are provided as part of the WebSphere Process Server product are migrated to the latest version of those applications. These are handled as follows:

- For all system applications—applications that reside in the *install_root* /systemApps directory, the newer version is installed.
- For all support applications—applications provided with WebSphere Process Server, such as the Business Rules Manager and Business Process Choreographer applications, older versions are updated to the latest version.

Sample applications are handled differently. For stand-alone profiles, the migration process will not install any sample applications. To make sample applications available for a stand-alone profile, you can install them using the installation wizard for the later version of WebSphere Process Server. For network deployment profiles, any samples installed with the previous version of WebSphere Process Server will be installed during migration to the new version.

Database migration

If you have a Cloudscape database, the migration tools will migrate the database configuration automatically, with certain exceptions. See “Migrating Cloudscape databases” on page 94 for more information. In addition, the Cloudscape database will be converted to a Derby database, which is the successor to Cloudscape and is supported by WebSphere Process Server version 6.1.

If you have a database other than Cloudscape, the migration tools will automatically migrate the provider and datasource definitions for each existing data source and provider. However, database schema upgrades may also be required, which could require special attention. If the server process has the

required database permissions and, in the case of some databases, meets other requirements, the schema upgrades will happen automatically when the server is first started.

If the server process does not have the required permissions or meet other requirements, or if you would like to manually upgrade your database schemas, then you will need to use the scripts provided. See “Upgrading databases for migration” on page 27 for more information.

Related concepts

“Overview of migrating” on page 1

Migrate from earlier versions of WebSphere Process Server and WebSphere Enterprise Service Bus.

Related tasks

“Migrating Cloudscape databases” on page 94

After you use the migration tools to migrate to WebSphere Process Server version 6.1, you should verify the results of the automatic Cloudscape database migration and manually migrate any Cloudscape database instances that are not automatically migrated by the tools.

Configuration mapping during product-configuration migration

Various configurations are mapped during product-configuration migration.

Migration always involves migrating a single profile to another single profile on the same machine or a separate machine. Examples include a WebSphere Process Server version 6.0.x deployment manager migrating to a version 6.1 deployment manager profile and a version 6.0.x stand-alone server migrating to a version 6.1 stand-alone server profile.

Many migration scenarios are possible. The migration tools map objects and attributes existing in the version from which you are migrating to the corresponding objects and attributes in the newer version’s environment.

Bootstrap port

The migration tools map a non-default value directly into the version 6.1 environment. If the `-portBlock` parameter is specified during the call to `WBIPostUpgrade`, however, a new port value is given to each server that is migrated to version 6.1.

Command-line parameters

The migration tools convert appropriate command-line parameters to Java virtual machine (JVM) settings in the server process definition. Most settings are mapped directly. Some settings are not migrated because their roles in the WebSphere Process Server version 6.1 configuration do not exist, have different meanings, or have different scopes.

For information on how to change the process-definition settings, see Process definition settings in the WebSphere Application Server Network Deployment, version 6.1 information center. For information on how to change the Java virtual machine settings, see Java virtual machine settings in the WebSphere Application Server Network Deployment, version 6.1 information center.

Java heap size for migrating EAR files

When migrating all WebSphere Process Server EAR files to version 6.1 using the wsadmin tool, the WBIPostUpgrade tool uses the default maximum Java heap size value of 64 MB to install the EAR files.

If a version 6.1 EAR file fails to install during migration because the Java heap size is not large enough, you will see a message similar to the following message:

```
java.lang.OutOfMemoryError JVMXE006:OutOfMemoryError
```

Increase the maximum Java heap size and follow the example below to install the application.

Example of installing the application on WebSphere Process Server version 6.1

Assume that:

Installation root

C:\WebSphere\DeploymentManager

Number signs (###)

Maximum heap size value

EAR_file_name

Name of the EAR file

app_name

Name of the application

cluster_name

Name of the cluster on which the EAR file should be installed

The command is displayed on more than one line for clarity.

```
wsadmin -conntype NONE
        -javaoption
        -Xmx###m
        -c "$AdminApp install
           C:\\WebSphere\\DeploymentManager\\installableApps\\
           EAR_file_name>
        {-nodeployejb
         -appname app_name
         -cluster cluster_name}"
```

Migration of a version 6.0.x node to a version 6.1 node

You can migrate a WebSphere ESB version 6.0.x node that belongs to a cell to WebSphere ESB version 6.1 without removing the node from the cell.

Migrate the deployment manager first, before migrating any base nodes in the cell.

Use the same cell name when migrating from version 6.0.x to version 6.1. If you use a different cell name, federated nodes cannot successfully migrate to the WebSphere ESB version 6.1 cell.

Migrating a base WebSphere ESB node that is within a cell to version 6.1 also migrates the node agent to version 6.1. A cell can have some version 6.1 nodes and other nodes that are at version 6.0.x levels.

Policy file

WebSphere Process Server version 6.1 migrates all the policy files that are installed with version 6.0.x policy files with the following characteristics:

- Any comments located in the version 6.1 policy file will be preserved. Any comments contained in the version 6.0.x policy file will not be included in the version 6.1.
- Migration will not attempt to merge permissions or grants; it is strictly an add-type migration. If the permission or grant is not located in the version 6.1 file, the migration will bring it over.
- Security is a critical component; thus, the migration makes any additions at the end of the original .policy file right after the comment MIGR0372I: Migrated grant permissions follow. This is done to help administrators verify any policy file changes that the migration has made.

Properties and lib/app directories

Migration copies files from prior version directories into the WebSphere Process Server version 6.1 configuration.

Property files

WebSphere migrates all the WebSphere Process Server version 6.1 property files that are installed with version 6.0.x by merging settings into the version 6.1 property files.

Migration does not overlay property files.

Resource adapter archives (RARs) referenced by J2C resources

RARs that are referenced by J2C resources are migrated if those RARs are in the old WebSphere Process Server installation. In this case, the RARs are copied over to the corresponding location in the new WebSphere Process Server installation. Relational Resource Adapter RARs will not be migrated.

Migrating cluster-level resources:

Cluster-level resources are configured in resourcexxx.xml files under the cluster directories. For example:

```
<resources.j2c:J2CResourceAdapter xmi:id="J2CResourceAdapter_1112808424172"
  name="ims" archivePath="{WAS_INSTALL_ROOT}\installedConnectors\x2.rar">
  ...
</resources.j2c:J2CResourceAdapter>
```

If you have a cluster-level resource, this resource must be in the same location on each cluster member (node). Using the above example, therefore, each cluster member must have the RAR file installed at location {WAS_INSTALL_ROOT}\installedConnectors\x2.rar. {WAS_INSTALL_ROOT} is resolved on each cluster member to get the exact location.

In the migration of a deployment manager, the tools migrate the cluster files on the deployment manager, including the resourcexxx.xml files.

In the migration of a managed node, the tools process each J2C adapter. Files such as RAR files are migrated as follows from version 6.0.x to version 6.1.

Migration from version 6.0.x to version 6.1 copies files such as RAR files from WAS_INSTALL_ROOT to WAS_INSTALL_ROOT and from USER_INSTALL_ROOT to USER_INSTALL_ROOT.

If you have a RAR file in the WAS_INSTALL_ROOT for version 6.0.x, for example, the migration tools do not automatically copy the file from

WAS_INSTALL_ROOT to USER_INSTALL_ROOT. This maintains the integrity of the cluster-level J2C resources. If you hardcoded a path to a RAR file (archivePath="C:/WAS/installedConnectors/x2.rar" for example) in version 6.0.x, however, the version 6.1 migration tools cannot change the archivePath attribute to reflect this because that would break all of the other cluster members that have not been migrated.

Samples

During the migration of the deployment manager, no WebSphere Process Server samples for federated nodes are migrated. Equivalent version 6.1 samples are available for all version 6.1 samples.

Security

Java 2 security is enabled by default when you enable security in WebSphere Process Server version 6.1. Java 2 security requires you to grant security permissions explicitly.

There are several techniques that you can use to define different levels of Java 2 security in version 6.1. One is to create a was.policy file as part of the application to enable all security permissions. The migration tools call the wsadmin command to add an existing was.policy file in the version 6.1 properties directory to enterprise applications as they are being migrated.

When migrating to WebSphere Process Server version 6.1, your choice of whether or not to migrate to support script compatibility results in one of two different outcomes.

- If you choose to migrate to support script compatibility, your security configuration is brought over to version 6.1 without any changes. This is the default.
- If you choose not to migrate to support script compatibility, the security configuration is converted to the default configuration for WebSphere Process Server version 6.1. The default version 6.1 security configuration acts almost the same as in the previous versions, but there are some changes.

For example, existing keyfiles and trustfiles are moved out of the SSLConfig repertoire and new keystore and truststore objects are created.

In order to retain the same security settings, you need to migrate the WebSphere Application Server security settings that might have been set for version 6.0.x. For more information on migrating your security configurations to version 6.1, see *Migrating, coexisting, and interoperating - Security considerations in the WebSphere Application Server Network Deployment, version 6.1 information center*.

Stdin, stdout, stderr, passivation, and working directories

The location for these directories is typically within the installation directory of a previous version. The default location for stdin, stdout, and stderr is the logs directory of the WebSphere Process Server version 6.1 installation root.

The migration tools attempt to migrate existing passivation and working directories. Otherwise, appropriate version 6.1 defaults are used.

For more information on passivation directories, see EJB container settings. For more information on working directories, see Process definition settings.

In a coexistence scenario, using common directories between versions can create problems.

Transport ports

The migration tools migrate all ports. The tools log a port-conflict warning if a port is already defined in the configuration. You must resolve any port conflicts before you can run servers at the same time.

If the `-portBlock` parameter is specified in the `WBIPostUpgrade` command, a new value is assigned to each transport that is migrated.

For more information on the `WBIPostUpgrade` command, see “`WBIPostUpgrade` command” on page 14.

For further information on transport chains and channels, see `Transport chains`.

You must manually add virtual host alias entries for each port. For more information, see `Configuring virtual hosts`.

Web modules

The specification level of the Java 2 Platform, Enterprise Edition (J2EE) implemented in WebSphere Process Server version 6.0.x required behavior changes in the Web container for setting the content type. If a default servlet writer does not set the content type, not only does the Web container no longer default to it but the Web container returns the call as “null.” This situation might cause some browsers to display resulting Web container tags incorrectly. To prevent this problem from occurring, migration sets the `autoResponseEncoding` IBM extension to “true” for Web modules as it migrates enterprise applications.

Related concepts

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Migrating WebSphere applications” on page 26

You should not have to modify any existing WebSphere Process Server applications to migrate them. However, more information about migrating different types of WebSphere applications is available in the WebSphere Application Server Network Deployment Information Center.

“Migrating clusters” on page 60

Migrate clusters by migrating, in turn, each profile that contains cluster members following special procedures. Take additional steps if you want to minimize down time of cluster services.

Related tasks

“Migrating non-clustered managed nodes” on page 49

Migrate a WebSphere Process Server managed node, choosing from several methods depending on your needs.

Related reference

 `WBIPostUpgrade` command

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the `backupDirectory` that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to

migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

Process definition settings

Java virtual machine settings

Migrating, coexisting, and interoperating - Security considerations

EJB container settings

Transport chains

Configuring virtual hosts

Task overview: Using enterprise beans in applications

Migrating WebSphere applications

You should not have to modify any existing WebSphere Process Server applications to migrate them. However, more information about migrating different types of WebSphere applications is available in the WebSphere Application Server Network Deployment Information Center.

Your applications—that is, any applications not provided with the WebSphere Process Server product—are binary-compatible for the migration scenarios supported. (Refer to “Overview of migrating” on page 1 for supported migration scenarios.) You should not have to modify any part of the application to have it run on the newer version of WebSphere Process Server.

Note: For version 6.0.1 WebSphere Adapters, some additional steps might be required for compatibility. For more information on this or any other exception, see the WebSphere Process Server technotes at the WebSphere Process Server technotes web site.

For more information about migrating certain types of WebSphere applications, see Migrating WebSphere Applications and its subsidiary topics in the WebSphere Application Server Network Deployment, version 6.1 Information Center. Because WebSphere Process Server is based upon , the same information applies.

For information about migrating WebSphere Adapters, refer to the documentation for your adapter in the WebSphere Integration Developer documentation in the IBM WebSphere Business Process Management Version 6.1 information center .

Note that applications developed on a more recent version of WebSphere Process Server will not run on older versions. For details about run-time compatibility, see Development and deployment version levels.

Related concepts

“Overview of migrating” on page 1

Migrate from earlier versions of WebSphere Process Server and WebSphere Enterprise Service Bus.



Development and deployment version levels

Your decision about what version levels of WebSphere Process Server you need in your environment will depend on the version levels with which your applications were developed. Generally applications deployed in a previous version of WebSphere Process Server will run on the next available version of WebSphere Process Server.

Upgrading databases for migration

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

About this task

Normally, database changes required by new versions of WebSphere Process Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, if the server has insufficient permissions to access the database schema, or if other database-specific requirements are not met, you must update the database manually.

For WebSphere Process Server version 6.1, the databases requiring schema upgrades are the following:

- Common database (default name WPRCSDB)
- Business Process Choreographer (default name BPEDB)

Perform the manual schema upgrade after migration, but before starting any server that uses the database.

For information about additional conditions required for automatic upgrade of Business Process Choreographer see “Upgrading the Business Process Choreographer database manually” on page 29.

For instructions on manually upgrading these databases, see the following subtopics.

Related tasks

“Migrating Cloudscape databases” on page 94

After you use the migration tools to migrate to WebSphere Process Server version 6.1, you should verify the results of the automatic Cloudscape database migration and manually migrate any Cloudscape database instances that are not automatically migrated by the tools.

Upgrading the Common database manually

After migrating the server from a previous version, you must upgrade to a new database schema for the “Common” database before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

Before you begin

- You should have already run the migration wizard or migration scripts to migrate the server, or in the case of a cluster, the servers in the cluster.
- Make sure the server or, if applicable, servers in the cluster remain stopped (do not start them after the migration wizard or scripts have run before completing the database upgrade).

About this task

Any database that is accessed by a migrated server needs to have its schema updated before you start the server. In the case of a cluster, any database that is accessed by any of the migrated cluster members needs to have its schema updated before you start any of the cluster members. You must upgrade manually

if the database user that is defined for the data source does not have sufficient authorization to modify the database schema. For the Common database, the database user that is configured for the data source must be authorized to perform all of the following operations: create and alter tables and create and drop indexes and views.

Procedure

1. Make sure that you are using a user ID with sufficient authority to update the database schema.
2. Back up the database.
3. Locate the directory where the database scripts are located:
 - **Windows** On Windows platforms: *install_root\dbscripts\component_name\database_type*
 - **Linux** **UNIX** On Linux and UNIX platforms: *install_root/dbscripts/component_name/database_type*

Where:

install_root

is the root directory into which WebSphere Process Server version 6.1 has been installed.

component_name

is CommonDB.

database_type

is a name corresponding to the database product you are using.

Applicable database types and their directory names are as follows:

Database type	Directory name
DB2 Universal Database™ (for all operating systems except z/OS® and i5/OS®)	DB2
DB2® for z/OS and OS/390® v. 7.x	DB2zOSV7
DB2 for z/OS and OS/390 v. 8.x and Version 9.x	DB2zOSV8
Derby	Derby Note: If you have an existing Cloudscape database for a previous (6.0.x) installation of WebSphere Process Server, the database must first be upgraded from a Cloudscape database to a Derby database. See “Migrating Cloudscape databases” on page 94. Then, once you have a Derby database that corresponds to that same (6.0.x) WebSphere Process Server version, you must then upgrade the database to correspond to WebSphere Process Server version 6.1 with the scripts in this directory.
Informix®	Informix
Oracle	Oracle
Microsoft® SQL Server	SQLServer

For specific database product names and versions see Database specifications.

4. Locate the migration scripts for the database and current schema version, where *schema_version* has the value 601 for version 6.0.1, or 602 for version

6.0.2. For example, to upgrade the Common database schema from WebSphere Process Server 6.0.1 to version 6.1, you will need `install_root/dbscripts/CommonDB/upgradeSchema601.sql`. To upgrade the Common database schema from WebSphere Process Server 6.0.2 to version 6.1 `install_root/dbscripts/CommonDB/upgradeSchema602.sql`.

Note: Other scripts exist in the same CommonDB directory, including some used for portions of the database upgrade process. Use only those named `upgradeSchema601.sql` or `upgradeSchema602.sql`, which run all of the required scripts at one time.

5. Copy the appropriate script from the directory in which you found it to the system on which the database is running.
6. Check the SQL scripts, and modify them, if necessary, to meet your requirements. This includes any scripts in the directory that are named according to this syntax: `upgradeSchema nnn` where nnn corresponds to a product version number. For example, you might need to change a user name, password, or file path.
7. Using your database client, connect to the database. This is to ensure that it can connect.
8. Run your version of the upgrade SQL scripts. For information on how to execute a `.sql` script with your database, refer to the documentation for your database product.
9. If there are any errors, or failure is indicated in your database client output, fix the reported errors and retry step 8.

Results

The database schema has been updated. When the server is started for the first time after the upgrade, the data is migrated according to the new schema. After the data has been migrated, version 6.0.x servers cannot run against the database.

Related tasks

“Upgrading the Business Process Choreographer database manually”

After migrating a server that has Business Process Choreographer configured, the schema for the associated Business Process Choreographer database must be upgraded before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

Upgrading the Business Process Choreographer database manually

After migrating a server that has Business Process Choreographer configured, the schema for the associated Business Process Choreographer database must be upgraded before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

Before you begin

- You should have already run the migration wizard or migration scripts to migrate the server, or in the case of a cluster, the servers in the cluster.
- Make sure the server or, if applicable, servers in the cluster remain stopped (do not start them after the migration wizard or scripts have run before completing the database upgrade).

About this task

The database associated with Business Process Choreographer that is accessed by a migrated server needs to have its schema updated before you start the server. In the case of a cluster, any such database that is accessed by any of the migrated cluster members needs to have its schema updated before you start any of the cluster members. You must upgrade manually if either of the following conditions are true:

- You did not use the default table spaces for the Business Process Choreographer database. If you used the sample Business Process Choreographer configuration, or have created all of the database objects in the default table spaces specified in the sample SQL scripts, your database uses the default table spaces. This is typically the case for a test environment.
- The database user that is configured for the BPEDB data source is not authorized to perform all of the following operations: create and alter tables, create and drop indexes and views, and for the table SCHEMA_VERSION: query, update, delete, and insert.

Procedure

1. Make sure that you are using a user ID with sufficient authority to update the database schema.
2. Back up the database.
3. Locate the directory where the database scripts are located:
 - **Windows** On Windows platforms: *install_root\dbscripts\component_name\database_type*
 - **Linux** **UNIX** On Linux and UNIX platforms: *install_root/dbscripts/component_name/database_type*

Where:

install_root

is the root directory into which WebSphere Process Server version 6.1 has been installed.

component_name

is ProcessChoreographer.

database_type

is a name corresponding to the database product you are using.

Applicable database types and their directory names are as follows:

Database type	Directory name
DB2 Universal Database (for all operating systems except z/OS and i5/OS)	DB2
DB2 for z/OS and OS/390 v. 7.x	DB2zOSV7
DB2 for z/OS and OS/390 v. 8.x and Version 9.x	DB2zOSV8

Database type	Directory name
Derby	Derby Note: If you have an existing Cloudscape database for a previous (6.0.x) installation of WebSphere Process Server, the database must first be upgraded from a Cloudscape database to a Derby database. See “Migrating Cloudscape databases” on page 94. Then, once you have a Derby database that corresponds to that same (6.0.x) WebSphere Process Server version, you must then upgrade the database to correspond to WebSphere Process Server version 6.1 with the scripts in this directory.
Informix	Informix
Oracle	Oracle
Microsoft SQL Server	SQLServer

For specific database product names and versions see Database specifications.

4. Locate the migration scripts for your database and current schema version, where *schema_version* has the value 601 for version 6.0.1, or 602 for version 6.0.2.

For DB2 on Linux, UNIX, and Windows:

Use one of the following upgrade scripts:

- `upgradeSchemaschema_version.sql` to create new database objects in the table spaces that were created using the `createTablespace.sql` script at schema creation time.
- `upgradeSchemaschema_versionnonp.sql` to create new objects in the default table space.

For DB2 on z/OS and OS/390:

Run these two scripts, in the following order:

- a. `upgradeTablespacesschema_version.sql` to upgrade the table spaces prior to upgrading the database objects.
- b. `upgradeSchemaschema_version.sql` to upgrade the database objects after you upgraded the table spaces.

For Derby:

Use one of the following upgrade scripts:

- `upgradeSchemaschema_version.sql` to upgrade the schema using a schema qualifier.
- `upgradeSchemaschema_versionnonp.sql` to upgrade the schema without using a custom schema qualifier.

For Informix Dynamic Server:

Use one of the following upgrade scripts:

- `upgradeSchemaschema_version.sql` to create new database objects in the database spaces that were created using the `createDbospace.sh` or `createDbospace.bat` shell script at schema creation time.
- `upgradeSchemaschema_versionnonp.sql` to create new objects in the default database space.

For Oracle

Use the script `upgradeSchemaschema_version.sql`.

For Microsoft SQL Server

Use one of the following upgrade scripts:

- `upgradeSchemaschema_version.sql` to upgrade the schema using a custom schema qualifier.
- `upgradeSchemaschema_versionnonp.sql` to upgrade the schema without using a custom schema qualifier.
- `upgradeSchema602Unicode.sql` if you upgrade from WebSphere Process Server version 6.0.2 and you have created a 6.0.2. schema with UNICODE support using the `createSchemaUnicode.sql` script or `createDatabaseUnicode.sql` script and if you want to use a custom schema qualifier.
- `upgradeSchema602UnicodeNonp.sql` if you upgrade from WebSphere Process Server version 6.0.2 and you have created a 6.0.2. schema with UNICODE support using the `createSchemaUnicode.sql` script or `createDatabaseUnicode.sql` script and if you do not want to use a custom schema qualifier.

Note: The `upgradeSchema...Nonp.sql` versions of the Microsoft SQL Server upgrade scripts do not make use of a schema qualifier. They upgrade the database objects in the user schema.

5. Copy the appropriate script or scripts from the directory in which you found it to the system on which the database is running.
6. Check the SQL script or scripts you just copied, and modify them, if necessary, to meet your requirements. For example, you might need to change a user name, password, schema qualifier, or file path.
7. Using your database client, connect to the database. This is to ensure that it can connect.
8. Run your version of the upgrade SQL scripts. For information on how to execute a `.sql` script with your database, refer to the documentation for your database product.
9. If there are any errors, or failure is indicated in your database client output, fix the reported errors and retry step 8.

Results

The database schema has been updated. When the server is started for the first time after the upgrade, the data is migrated according to the new schema. After the data has been migrated, version 6.0.x servers cannot run against the database.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Related tasks

“Upgrading the Common database manually” on page 27

After migrating the server from a previous version, you must upgrade to a new database schema for the “Common” database before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

“Migrating non-clustered managed nodes using the migration wizard” on page 49

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server using the migration wizard.



Migrating a managed node using command-line tools

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server with the command-line tools.

“Migrating a cluster” on page 60

To migrate a cluster, migrate each profile containing a member of that cluster one at a time. The migration requires extra steps not required for a non-clustered environment.

“Migrating a cluster with minimal down time” on page 69

To migrate a cluster while minimizing down time, first migrate approximately half of the profiles contributing to the cluster, then migrate the other half.

Perform the extra steps required for cluster migration after you have migrated the first set of profiles.

Migrating stand-alone servers

Migrate a stand-alone WebSphere Process Server server, choosing from several methods depending on your needs.

Before you begin

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

Select the appropriate migration scenario for information on how to migrate a WebSphere Process Server stand-alone server from an older version to a newer version of WebSphere Process Server. (For example, a WebSphere Process Server Version 6.0.2 stand-alone server to a WebSphere Process Server Version 6.1 stand-alone server.)

- “Migrating a stand-alone server using the migration wizard”
This topic contains instructions for migrating an older version WebSphere Process Server stand-alone server to a newer version stand-alone server using the migration wizard (a graphical user interface).
- “Migrating to a stand-alone server using command-line tools” on page 37
This topic contains instructions for migrating an older version WebSphere Process Server stand-alone server to a newer version stand-alone server using the migration command line tools rather than the migration wizard.
- “Migrating to a remote system” on page 79
This topic contains instructions for migrating an older version WebSphere Process Server server to a newer version server on a remote system.
- “Migrating from an operating system that is no longer supported” on page 83
This topic contains instructions for migrating a WebSphere Process Server stand-alone server that is running on an operating system that WebSphere Process Server no longer supports.

Tip: For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Migrating a stand-alone server using the migration wizard

Migrate a stand-alone server from an older version to a newer version of WebSphere Process Server using the migration wizard.

Before you begin

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.

Note: If you plan to migrate to a new physical computer system then use the alternate procedure described in “Migrating to a remote system” on page 79. Or, if you plan to upgrade the version of the operating system on the computer that is running WebSphere Process Server, then use the alternate procedure described in “Migrating from an operating system that is no longer supported” on page 83.

- A non-federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up your existing configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the stopServer command from the *profile_dir*/bin directory for the profile of the affected server, or stop the server from the profile’s First steps console.

For more information about the stopServer command see the stopServer command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the -username and -password parameters on the stopServer command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:

```
stopServer.bat server1
```

On the Windows operating system, even if security is enabled, the -username and -password parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Optional: Create a new version 6.1 profile with tools outside the migration wizard.

The migration wizard allows you to create a new profile based on an existing version 6.0.x profile or to use an existing version 6.1 profile for your target profile. To create such an existing version 6.1 profile you can use the Profile Management Tool or the `manageprofiles` command to create a new (version 6.1) profile before you run the migration wizard. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured. The migration wizard will detect the new profile and display its name as a possible choice for your target profile. See *Creating profiles* for information about creating profiles.

4. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name




See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

5. Invoke the migration wizard.

Invoke the migration wizard in one of the following ways:

- From the WebSphere Process Server First Steps console, select **Migration wizard**.
- Run one of the following scripts (depending upon your operating system) stored in the `install_dir/bin` directory:
 - a.   On UNIX and Linux systems: `wbi_migration.sh`
 - b.  On Windows systems: `wbi_migration.bat`

For information about what processing the migration wizard actually performs, see “What the migration wizard does” on page 7.

6. Follow the prompts for the migration wizard as described in “Running the migration wizard” on page 9.
7. If required, manually update the databases used by WebSphere Process Server. Normally, database changes required by new versions of WebSphere Process

Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, in cases in which the server has insufficient permissions to access the database schema, or other database-specific requirements are not met, you must update the database manually. For more information see “Upgrading databases for migration” on page 27.

Results

You have now migrated your stand-alone server.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

“Running the migration wizard” on page 9

 Migrating standalone servers using command-line tools

Migrate a stand-alone server from an older version to a newer version of WebSphere Process Server using the command-line tools.

 Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

“Migrating to a remote system” on page 79

Use the migration tools to migrate from an older version on one system to a newer version of WebSphere Process Server on a different, remote system. (Stand-alone servers only.)

“Migrating from an operating system that is no longer supported” on page 83

Use the migration tools to migrate an earlier WebSphere Process Server release that is running on an operating system that the newer version does not support. (Stand-alone servers only.)

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

backupConfig command

stopServer command

Migrating to a stand-alone server using command-line tools

Migrate a stand-alone server from an older version to a newer version of WebSphere Process Server using the command-line tools.

Before you begin

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.

Note: If you plan to migrate to a new physical computer system then use the alternate procedure described in “Migrating to a remote system” on page 79. Or, if you plan to upgrade the version of the operating system on the computer that is running WebSphere Process Server, then use the alternate procedure described in “Migrating from an operating system that is no longer supported” on page 83.

- A non-federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up your existing configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the stopServer command from the *profile_dir/bin* directory for the profile of the affected server, or stop the server from the profile's First steps console.

For more information about the stopServer command see the stopServer command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the -username and -password parameters on the stopServer command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:

```
stopServer.bat server1
```

On the Windows operating system, even if security is enabled, the -username and -password parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Use the Profile Management Tool or the **manageprofiles** command to create a new profile with the newer version of WebSphere Process Server. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured. See Creating profiles for information about creating profiles.
4. Run the WBIPreUpgrade command, specifying the migration backup directory name and the existing WebSphere Process Server directory name. The WBIPreUpgrade tool saves selected files from the *install_root* and *profile_root* directories to a backup directory you specify. See the "WBIPreUpgrade command" on page 12 for details.
5. Run the WBIPostUpgrade command, specifying the migration backup directory. The WBIPostUpgrade tool restores the environment in the backup directory into the new WebSphere Process Server stand-alone server installation. See the "WBIPostUpgrade command" on page 14 for details.
6. If required, manually update the databases used by WebSphere Process Server. Normally, database changes required by new versions of WebSphere Process

Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, in cases in which the server has insufficient permissions to access the database schema, or other database-specific requirements are not met, you must update the database manually. For more information see “Upgrading databases for migration” on page 27.

Results

You have now migrated your stand-alone server.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

“Migrating a stand-alone server using the migration wizard” on page 33
Migrate a stand-alone server from an older version to a newer version of WebSphere Process Server using the migration wizard.

Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

 WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

backupConfig command

stopServer command

Migrating a network deployment environment

Migrate a WebSphere Process Server network deployment environment.

Before you begin

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

About this task

To migrate a network deployment environment, you need to first migrate the deployment manager, and then migrate its managed nodes.

Select the appropriate migration scenario for information on how to migrate to a WebSphere Process Server version 6.1 deployment cell.

Note: For managed nodes that are not in a clustered environment, see “Migrating non-clustered managed nodes” on page 49. For managed nodes in a clustered environment, see “Migrating clusters” on page 60.

Tip: For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Migrating a deployment manager

Migrate a WebSphere Process Server deployment manager, choosing from several methods depending on your needs.

Before you begin

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

Select the appropriate migration scenario for information on how to migrate a WebSphere Process Server deployment manager from an older version to a newer

version of WebSphere Process Server. (For example, a WebSphere Process Server Version 6.0.2 deployment manager to a WebSphere Process Server Version 6.1 deployment manager.)

- “Migrating a deployment manager using the migration wizard”
This topic contains instructions for migrating an older version WebSphere Process Server deployment manager to a newer version deployment manager using the migration wizard (a graphical user interface).
- “Migrating a deployment manager using command-line tools” on page 45
This topic contains instructions for migrating an older version WebSphere Process Server deployment manager to a newer version deployment manager using the migration command line tools rather than the migration wizard.

Tip: For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Related concepts

“Migrating clusters” on page 60

Migrate clusters by migrating, in turn, each profile that contains cluster members following special procedures. Take additional steps if you want to minimize down time of cluster services.

Migrating a deployment manager using the migration wizard

Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the migration wizard.

Before you begin

Note: Migrate the WebSphere Process Server version 6.0.x deployment manager to version 6.1 before migrating the managed nodes that comprise the cell. The deployment manager must always be at the highest release and fix level within a cell in order for it to manage all nodes in the cell. A version 6.1 deployment manager can manage version 6.0.1, version 6.0.2, and version 6.1 managed nodes. This allows a cell to be upgraded to a new release one node at a time, with minimal impact to the applications that are running within the cell. See *Coexisting with other WebSphere product installations* for information on restrictions on using mixed-release cells.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A deployment manager profile, created with the older WebSphere Process Server version, resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for details about disk space requirements.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.

- Use the backupConfig command or your own preferred backup utility to back up your existing configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Cell name

Name of the cell managed by the deployment manager that you are migrating. The new version cell name must match the name in the old version's configuration.

Node name

Name of the node that you are migrating. The new version node name must match the name in the old version's configuration.

Installation root directory

See "WBIPreUpgrade command" on page 12 for a description of the -currentWebSphereDirectory parameter.

Migration backup directory name

See "WBIPreUpgrade command" on page 12 for a description of the -backupDirectory parameter.

Source profile name

See "WBIPostUpgrade command" on page 14 for a description of the -oldProfile parameter.

Target profile name

See "WBIPostUpgrade command" on page 14 for a description of the -profileName parameter.

Port value assignments (optional)

See "WBIPostUpgrade command" on page 14 for a description of the -replacePorts and -portBlock parameters.

3. Optional: Create a new version 6.1 profile with tools outside the migration wizard.

The migration wizard allows you to create a new profile based on an existing version 6.0.x profile or to use an existing version 6.1 profile for your target profile. To create such an existing version 6.1 profile you can use the Profile Management Tool or the manageprofiles command to create a new (version 6.1) profile before you run the migration wizard. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured. The migration wizard will detect the new profile and display its name as a possible choice for your target profile. See Creating profiles for information about creating profiles.

4. Stop the deployment manager that you are about to migrate. Use the **stopManager** command from the deployment manager's *profile_dir/bin* directory or from the deployment manager's First steps console.

See stopManager for more information about the stopManager command.

Linux

For example, use the following command on a Linux platform:





```
./stopManager.sh
```

If you have security enabled, specify the `-username` and `-password` parameters of the command.

You can migrate a deployment manager whether it is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the old version deployment manager before you can start the new version deployment manager that you are installing, however, so it makes sense to stop it now.

5. Invoke the migration wizard from the WebSphere Process Server to which you are migrating.

Invoke the migration wizard in one of the following ways:

- From the WebSphere Process Server First Steps console, select **Migration wizard**.
- Run one of the following scripts (depending upon your operating system) stored in the `install_dir/bin` directory:
 - a.   On UNIX and Linux systems: `wbi_migration.sh`
 - b.  On Windows systems: `wbi_migration.bat`

For information about what processing the migration wizard actually performs, see “What the migration wizard does” on page 7.

6. Follow the prompts for the migration wizard as described in “Running the migration wizard” on page 9.
7. If you need to manually update the Common database, do so now. See “Upgrading the Common database manually” on page 27 for instructions. Normally, database changes required by new versions of WebSphere Process Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, in cases in which the server has insufficient permissions to access the database schema, or other database-specific requirements are not met, you must update the database manually.

Note: If a server or cluster has Business Process Choreographer configured, upgrade the database after you migrate the server or cluster, not when you migrate the deployment manager.

8. Start the version 6.1 deployment manager. Use the `startManager` command from the `profile_dir/bin` directory or the First steps console. See `startManager` command for more information on the `startManager` command.
9. Optional: Uninstall the version 6.0.x deployment manager.

Perform this step only after you are certain that you have successfully migrated the configuration of the deployment manager that you intend to delete. For more information about uninstalling, see Uninstalling the software.

Results

Your deployment manager is now migrated.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

“Running the migration wizard” on page 9

 Migrating a deployment manager using command-line tools

Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the command-line tools.

“Migrating non-clustered managed nodes” on page 49

Migrate a WebSphere Process Server managed node, choosing from several methods depending on your needs.

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

 Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Rolling back a deployment cell” on page 89

You can use the `restoreConfig` and `wsadmin` commands to roll back a migrated WebSphere Process Server version 6.1 deployment cell to version 6.0.x. This returns the configuration to the state that it was in before migration. After rolling back the deployment cell, you can restart the migration process.

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

 WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command

from the *backupDirectory* that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

`backupConfig` command

`stopServer` command

 Uninstalling the software

Learn about the different ways of uninstalling IBM WebSphere Process Server.

 Coexisting with other WebSphere product installations

An installation of WebSphere Process Server, version 6.1 can coexist on the same system with installations of any version of WebSphere Process Server or WebSphere Enterprise Service Bus, and with certain versions of selected WebSphere products.

Migrating a deployment manager using command-line tools

Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the command-line tools.

Before you begin

Note: Migrate the WebSphere Process Server version 6.0.x deployment manager to version 6.1 before migrating the managed nodes that comprise the cell. The deployment manager must always be at the highest release and fix level within a cell in order for it to manage all nodes in the cell. A version 6.1 deployment manager can manage version 6.0.1, version 6.0.2, and version 6.1 managed nodes. This allows a cell to be upgraded to a new release one node at a time, with minimal impact to the applications that are running within the cell. See *Coexisting with other WebSphere product installations* for information on restrictions on using mixed-release cells.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A deployment manager profile, created with the older WebSphere Process Server version, resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for details about disk space requirements.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the `backupConfig` command or your own preferred backup utility to back up your existing configuration if you want to be able to restore it to its previous

state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Cell name

Name of the cell managed by the deployment manager that you are migrating. The new version cell name must match the name in the old version's configuration.

Node name

Name of the node that you are migrating. The new version node name must match the name in the old version's configuration.

Installation root directory

See "WBIPreUpgrade command" on page 12 for a description of the -currentWebSphereDirectory parameter.

Migration backup directory name

See "WBIPreUpgrade command" on page 12 for a description of the -backupDirectory parameter.

Source profile name

See "WBIPostUpgrade command" on page 14 for a description of the -oldProfile parameter.

Target profile name

See "WBIPostUpgrade command" on page 14 for a description of the -profileName parameter.

Port value assignments (optional)

See "WBIPostUpgrade command" on page 14 for a description of the -replacePorts and -portBlock parameters.

3. Use the Profile Management Tool or the **manageprofiles** command to create a new profile with the newer version of WebSphere Process Server. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured. See Creating profiles for information about creating profiles.
4. Stop the deployment manager that you are about to migrate. Use the **stopManager** command from the deployment manager's *profile_dir/bin* directory or from the deployment manager's First steps console. See stopManager for more information about the stopManager command.

Linux

For example, use the following command on a Linux platform:

```
./stopManager.sh
```

If you have security enabled, specify the -username and -password parameters of the command.

You can migrate a deployment manager whether it is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the old version deployment manager before you can start the new version deployment manager that you are installing, however, so it makes sense to stop it now.

5. Run the `WBIPreUpgrade` command, specifying the migration backup directory name and the existing WebSphere Process Server directory name. The `WBIPreUpgrade` tool saves selected files from the `install_root` and `profile_root` directories to a backup directory you specify. See the “`WBIPreUpgrade` command” on page 12 for details.
6. Run the `WBIPostUpgrade` command, specifying the migration backup directory. The `WBIPostUpgrade` tool restores the environment in the backup directory into the new WebSphere Process Server stand-alone server installation. See the “`WBIPostUpgrade` command” on page 14 for details.
7. If you need to manually update the Common database, do so now. See “`Upgrading the Common database manually`” on page 27 for instructions. Normally, database changes required by new versions of WebSphere Process Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, in cases in which the server has insufficient permissions to access the database schema, or other database-specific requirements are not met, you must update the database manually.

Note: If a server or cluster has Business Process Choreographer configured, upgrade the database after you migrate the server or cluster, not when you migrate the deployment manager.

8. Start the version 6.1 deployment manager. Use the `startManager` command from the `profile_dir/bin` directory or the First steps console. See `startManager` command for more information on the `startManager` command.
9. Optional: Uninstall the version 6.0.x deployment manager.
Perform this step only after you are certain that you have successfully migrated the configuration of the deployment manager that you intend to delete. For more information about uninstalling, see `Uninstalling the software`.

Results

Your deployment manager is now migrated.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “`Migration considerations for Business Process Choreographer`” on page 105. Finally, perform the checks described in “`Postmigration configuration checking`” on page 88.

Related concepts

“`Migration considerations for Business Process Choreographer`” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

“`Premigration considerations`” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“`Migration wizard`” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

“Migrating a deployment manager using the migration wizard” on page 41
Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the migration wizard.

“Migrating non-clustered managed nodes” on page 49
Migrate a WebSphere Process Server managed node, choosing from several methods depending on your needs.

“Verifying migration” on page 87
Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107
Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88
After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

“Upgrading databases for migration” on page 27
In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Rolling back a deployment cell” on page 89
You can use the `restoreConfig` and `wsadmin` commands to roll back a migrated WebSphere Process Server version 6.1 deployment cell to version 6.0.x. This returns the configuration to the state that it was in before migration. After rolling back the deployment cell, you can restart the migration process.

Related reference

“WBIPreUpgrade command” on page 12
Use the `WBIPreUpgrade` command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

WBIPostUpgrade command

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the *backupDirectory* that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

`backupConfig` command

`stopServer` command

Uninstalling the software

Learn about the different ways of uninstalling IBM WebSphere Process Server.

Coexisting with other WebSphere product installations

An installation of WebSphere Process Server, version 6.1 can coexist on the same system with installations of any version of WebSphere Process Server or WebSphere Enterprise Service Bus, and with certain versions of selected WebSphere products.

Migrating non-clustered managed nodes

Migrate a WebSphere Process Server managed node, choosing from several methods depending on your needs.

Before you begin

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

Select the appropriate subtopic for information on how to migrate a WebSphere Process Server non-clustered managed node from an older version to a newer version of WebSphere Process Server.

Note: If your managed nodes are part of a cluster, follow the instructions in “Migrating clusters” on page 60.

Tip: For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Related tasks

“Migrating a cluster” on page 60

To migrate a cluster, migrate each profile containing a member of that cluster one at a time. The migration requires extra steps not required for a non-clustered environment.

“Migrating a cluster with minimal down time” on page 69

To migrate a cluster while minimizing down time, first migrate approximately half of the profiles contributing to the cluster, then migrate the other half. Perform the extra steps required for cluster migration after you have migrated the first set of profiles.

“Migrating a deployment manager using the migration wizard” on page 41
Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the migration wizard.

 Migrating a deployment manager using command-line tools

Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the command-line tools.

Migrating non-clustered managed nodes using the migration wizard

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server using the migration wizard.

Before you begin

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.

- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.
- The deployment manager that manages the managed node that you intend to migrate has already been migrated to the newer version of WebSphere Process Server, and is running.

Note: Migrating a WebSphere Process Server version 6.0.x managed node to a version 6.1 managed node requires that you first migrate the version 6.0.x deployment manager to a version 6.1 deployment manager. See “Migrating a deployment manager” on page 40 for instructions. Complete the migration of the deployment manager before you proceed with the instructions in this topic.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up the managed node configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

About this task

After you migrate an older version deployment manager to a newer version of WebSphere Process Server, the newer version deployment manager runs in compatibility mode by default, where it can manage both older and newer versions of WebSphere Process Server. For example, after migration, a version 6.1 deployment manager can manage both version 6.0.x and version 6.1 release nodes. The managed nodes of the version 6.0.x deployment manager are now running as version 6.0.x managed nodes in the version 6.1 deployment manager. See Coexisting with other WebSphere product installations for information on restrictions on using mixed-release cells.

Over time, migrate each version 6.0.x WebSphere Process Server managed node (server managed by a version 6.1 deployment manager) to a version 6.1 managed node. After migrating all the version 6.0.x managed nodes, use the **convertScriptCompatibility** script to change the deployment manager from a node that supports backward compatibility of version 6.0.x administration scripts to a node that supports only version 6.1. See the convertScriptCompatibility command.

Note: When following the directions at this link to use the convertScriptCompatibility command, use the WBIPostUpgrade command rather than the WASPostUpgrade command.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the stopServer command from the *profile_dir/bin* directory for the profile of the affected server, or stop the server from the profile's First steps console. For more information about the stopServer command see the stopServer command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the -username and -password parameters on the stopServer command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:
stopServer.bat server1

On the Windows operating system, even if security is enabled, the -username and -password parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Stop the node agent of the node to be migrated. Issue one of the following commands to stop the nodeagent process, depending on platform (where *profile_root* represents the installation directory of the federated node):
 - **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat*

If security is enabled, use one of the following commands instead:

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh -username user_ID -password password*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat -username user_ID -password password*

Note: A node can either be stopped or running during migration. You can migrate a managed node without stopping it (stopping its node agent); however, you must stop the old node before you can start the new one (the node running the newer version of WebSphere Process Server. At the same time, it is not necessary to have the node running to migrate its configuration. The migration tools can retrieve all the configuration data while the node is stopped.

4. Optional: Create a new version 6.1 profile with tools outside the migration wizard.

You can create either a version 6.1 stand-alone server profile or a custom profile as a target. If you create a custom profile, do not federate the node before migration. The migration tools federate the node during migration.

Note: For migration to be successful, you must use the same node name and cell name for each node that you migrate from version 6.0.x to version 6.1.

Tip: If you make any cell-level changes to the new version 6.1 node before migration, such as changes to virtual-host information, these changes will be lost during migration. Therefore, you should wait until after the node has been migrated before making any such changes. Otherwise, you will have to manually remake all of the changes, such as any changes to the virtual-host and host-alias information, to the new cell after migration using the administrative console running on the deployment manager.

5. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)



See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

6. Ensure that the version 6.1 deployment manager is up and running.

Note: You can migrate a version 6.0.x node whether the node is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the version 6.0.x node before you can start the version 6.1 node that you are installing, however, so it makes sense to stop it now.

7. Invoke the migration wizard.

Invoke the migration wizard in one of the following ways:

- From the WebSphere Process Server First Steps console, select **Migration wizard**.
- Run one of the following scripts (depending upon your operating system) stored in the `install_dir/bin` directory:
 - a.  On UNIX and Linux systems: `wbi_migration.sh`
 - b.  On Windows systems: `wbi_migration.bat`

The migration wizard copies the configuration and applications from the version 6.0.x managed node to the version 6.1 managed node. After migrating all of the data, the wizard federates the version 6.1 managed node into the deployment manager cell.

8. Stop (if they are not stopped already) the server and node agent. If the server has not already been stopped, stop the server as described in step 2. If the node agent has not already been stopped, stop the node agent as described in step 3.

9. If you are migrating a server that has Business Process Choreographer configured and you need to upgrade its associated database manually, do this step now. See “Upgrading the Business Process Choreographer database manually” on page 29.
10. Restart the node agent. To start a node agent, run the command `profile_root\bin\startNode` (where `profile_root` represents the installation directory of the managed node).
 - **Linux** **UNIX** **On Linux and UNIX platforms:** `profile_root/bin/startNode.sh`
 - **Windows** **On Windows platforms:** `profile_root\bin\startNode.bat`
11. Start the server or servers running on this node. Start each server using the `startServer` command, the administrative console, or the profile’s First steps console. For more information see Starting an application server.
12. Repeat steps 1-11 for each additional managed node you wish to migrate.
13. If you chose the compatibility option (which is the default), and if all of your nodes are completely migrated to WebSphere Process Server version 6.1, run the `convertScriptCompatibility` script to remove backward compatibility from the version 6.1 deployment manager.

Issue the `convertScriptCompatibility` command from the `bin` directory.

 - **UNIX** **Linux** `install_root/bin/convertScriptCompatibility.sh`
 - **Windows** `install_root\bin\convertScriptCompatibility.bat`

See the `convertScriptCompatibility` command.

Results

Your non-clustered managed nodes are now migrated.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

“Running the migration wizard” on page 9

 Migrating a managed node using command-line tools

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server with the command-line tools.

“Upgrading the Business Process Choreographer database manually” on page 29

After migrating a server that has Business Process Choreographer configured, the schema for the associated Business Process Choreographer database must be

upgraded before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Rolling back a managed node” on page 91

You can use the `restoreConfig` and `wsadmin` commands to roll back a migrated WebSphere Process Server version 6.1 managed node to the state that it was in before migration. For each managed node that you want to roll back, you must roll back the managed node itself and the corresponding changes made to the master repository located on the deployment manager.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

Related reference

“WBIPreUpgrade command” on page 12

Use the `WBIPreUpgrade` command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

WBIPostUpgrade command

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the `backupDirectory` that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the `profile_root/installedApps` directory for the new installation.

Related information

Starting an application server

`convertScriptCompatibility` command

`startManager` command

`backupConfig` command

`stopServer` command

Migrating non-clustered managed nodes using the command-line tools

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server with the command-line tools.

Before you begin

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.
- The deployment manager that manages the managed node that you intend to migrate has already been migrated to the newer version of WebSphere Process Server, and is running.

Note: Migrating a WebSphere Process Server version 6.0.x managed node to a version 6.1 managed node requires that you first migrate the version 6.0.x deployment manager to a version 6.1 deployment manager. See “Migrating a deployment manager” on page 40 for instructions. Complete the migration of the deployment manager before you proceed with the instructions in this topic.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up the managed node configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

About this task

After you migrate an older version deployment manager to a newer version of WebSphere Process Server, the newer version deployment manager runs in compatibility mode by default, where it can manage both older and newer versions of WebSphere Process Server. For example, after migration, a version 6.1 deployment manager can manage both version 6.0.x and version 6.1 release nodes. The managed nodes of the version 6.0.x deployment manager are now running as version 6.0.x managed nodes in the version 6.1 deployment manager. See Coexisting with other WebSphere product installations for information on restrictions on using mixed-release cells.

Over time, migrate each version 6.0.x WebSphere Process Server managed node (server managed by a version 6.1 deployment manager) to a version 6.1 managed node. After migrating all the version 6.0.x managed nodes, use the **convertScriptCompatibility** script to change the deployment manager from a node that supports backward compatibility of version 6.0.x administration scripts to a node that supports only version 6.1. See the convertScriptCompatibility command.

Note: When following the directions at this link to use the convertScriptCompatibility command, use the WBIPostUpgrade command rather than the WASPostUpgrade command.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the stopServer command from the *profile_dir/bin* directory for the profile of the affected server, or stop the server from the profile's First steps console.

For more information about the stopServer command see the stopServer command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the -username and -password parameters on the stopServer command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:
stopServer.bat server1

On the Windows operating system, even if security is enabled, the -username and -password parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Stop the node agent of the node to be migrated. Issue one of the following commands to stop the nodeagent process, depending on platform (where *profile_root* represents the installation directory of the federated node):

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat*

If security is enabled, use one of the following commands instead:

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh -username user_ID -password password*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat -username user_ID -password password*

Note: A node can either be stopped or running during migration. You can migrate a managed node without stopping it (stopping its node agent); however, you must stop the old node before you can start the new one (the node running the newer version of WebSphere Process Server. At the same time, it is not necessary to have the node running to migrate its configuration. The migration tools can retrieve all the configuration data while the node is stopped.

4. Create a new profile. Use the Profile Management tool or the manageprofiles command to create a new profile with the newer version of WebSphere Process Server.

You can create either a version 6.1 stand-alone server profile or a custom profile as a target. If you create a custom profile, do not federate the node before migration. The migration tools federate the node during migration. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured.

Note: For migration to be successful, you must use the same node name and cell name for each node that you migrate from version 6.0.x to version 6.1.

Tip: If you make any cell-level changes to the new version 6.1 node before migration, such as changes to virtual-host information, these changes will be lost during migration. Therefore, you should wait until after the node has been migrated before making any such changes. Otherwise, you will have to manually remake all of the changes, such as any changes to the virtual-host and host-alias information, to the new cell after migration using the administrative console running on the deployment manager.

5. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

6. Ensure that the version 6.1 deployment manager is up and running.

Note: You can migrate a version 6.0.x node whether the node is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the version 6.0.x node before you can start the version 6.1 node that you are installing, however, so it makes sense to stop it now.

7. Run the WBIPreUpgrade command, specifying the migration backup directory name and the existing WebSphere Process Server directory name. The WBIPreUpgrade tool saves selected files from the `install_root` and `profile_root` directories to a backup directory you specify. See the “WBIPreUpgrade command” on page 12 for details.
8. Run the WBIPostUpgrade command, specifying the migration backup directory. The WBIPostUpgrade tool restores the environment in the backup directory into the new WebSphere Process Server stand-alone server installation. See the “WBIPostUpgrade command” on page 14 for details.

9. Stop (if they are not stopped already) the server and node agent. If the server has not already been stopped, stop the server as described in step 2. If the node agent has not already been stopped, stop the node agent as described in step 3.
10. If you are migrating a server that has Business Process Choreographer configured and you need to upgrade its associated database manually, do this step now. See “Upgrading the Business Process Choreographer database manually” on page 29.
11. Restart the node agent. To start a node agent, run the command `profile_root\bin\startNode` (where `profile_root` represents the installation directory of the managed node).
 - **Linux** **UNIX** On Linux and UNIX platforms: `profile_root/bin/startNode.sh`
 - **Windows** On Windows platforms: `profile_root\bin\startNode.bat`
12. Start the server or servers running on this node. Start each server using the `startServer` command, the administrative console, or the profile’s First steps console. For more information see Starting an application server.
13. Repeat steps 1- 12 for each additional managed node you wish to migrate.

Note: You need to perform step 7 (running `WBIPreUpgrade`) again only if the version 6.0.x system has been re-configured since the first time you ran `WBIPreUpgrade`.

14. If you chose the compatibility option (which is the default), and if all of your nodes are completely migrated to WebSphere Process Server version 6.1, run the `convertScriptCompatibility` script to remove backward compatibility from the version 6.1 deployment manager.

Issue the `convertScriptCompatibility` command from the `bin` directory.

- **UNIX** **Linux** `install_root/bin/convertScriptCompatibility.sh`
- **Windows** `install_root\bin\convertScriptCompatibility.bat`

See the `convertScriptCompatibility` command.

Results

Your non-clustered managed nodes are now migrated.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Migration wizard” on page 7

The version-to-version migration wizard is a graphical interface that guides you through migrating from an older version to a newer version of WebSphere Process Server.

Related tasks

“Migrating non-clustered managed nodes using the migration wizard” on page 49

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server using the migration wizard.

“Upgrading the Business Process Choreographer database manually” on page 29

After migrating a server that has Business Process Choreographer configured, the schema for the associated Business Process Choreographer database must be upgraded before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Rolling back a managed node” on page 91

You can use the `restoreConfig` and `wsadmin` commands to roll back a migrated WebSphere Process Server version 6.1 managed node to the state that it was in before migration. For each managed node that you want to roll back, you must roll back the managed node itself and the corresponding changes made to the master repository located on the deployment manager.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

Related reference

“WBIPreUpgrade command” on page 12

Use the `WBIPreUpgrade` command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

WBIPostUpgrade command

Use the `WBIPostUpgrade` command for WebSphere Process Server to retrieve the saved configuration that was created by the `WBIPreUpgrade` command from the *backupDirectory* that you specified. The `WBIPostUpgrade` command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

Starting an application server

`convertScriptCompatibility` command

`startManager` command

`backupConfig` command

`stopServer` command

Migrating clusters

Migrate clusters by migrating, in turn, each profile that contains cluster members following special procedures. Take additional steps if you want to minimize down time of cluster services.

Migrating a cluster requires you to stop all node agents and servers that contribute to a cluster before migrating each profile. In addition, you must run the `WBIPProfileUpgrade` script for each cluster. Choose from the following subtopics, depending on your needs:

Related tasks

“Migrating a deployment manager” on page 40

Migrate a WebSphere Process Server deployment manager, choosing from several methods depending on your needs.

Migrating a cluster

To migrate a cluster, migrate each profile containing a member of that cluster one at a time. The migration requires extra steps not required for a non-clustered environment.

Before you begin

Note: In a cluster, version 6.0.x members and version 6.1 members must never run at the same time. All version 6.0.x cluster members must be stopped before you start the first version 6.1 cluster member. Also, once you start a version 6.1 cluster member, do not start any version 6.0.x cluster members in that cluster.

You must have an existing cell containing at least one cluster running on an older version of WebSphere Process Server (for example, version 6.0.x) that you wish to migrate to a newer version (for example, version 6.1). In addition, you must have installed the new version of WebSphere Process Server.

About this task

Following these steps will ensure that you retain cluster functionality on the new version of WebSphere Process Server.

Procedure

1. Migrate the deployment manager. Follow one of the instruction sets listed in “Migrating a deployment manager” on page 40 to complete this task.
2. Make sure that the new deployment manager is running.
3. Identify the profiles involved.
 - a. Identify an older version profile that contains cluster members.
 - b. Identify which other clusters this profile contributes to; that is, if the profile defines servers that are members of any other clusters, identify those clusters.
 - c. Identify all other profiles in the same cell that contribute cluster members to any of the clusters identified in step 3b.
 - d. Identify all the node agents and process servers defined by any of the profiles identified in step 3c.

All of the profiles identified in step 3c, and all of the corresponding node agents and servers identified in step 3d will be involved in the migration.

4. Stop all the node agents and servers identified in step 3d.

5. Migrate each profile identified in step 3c on page 60, one at a time, but **do not start** any new agents or servers. To do so, follow one of the instruction sets listed in either “Migrating cluster members using the migration wizard” on page 62 or “Migrating cluster members using the command-line tools” on page 65.
6. On the system that hosts the WebSphere Process Server version 6.1 deployment manager profile, navigate to the *install_dir/util* directory. This directory contains the `WBIPprofileUpgrade` script, `WBIPprofileUpgrade.ant`.
7. Run the `WBIPprofileUpgrade` script for each cluster defined in the migrated profiles. Run `WBIPprofileUpgrade` for each cluster defined in step 3b on page 60. For instructions on running `WBIPprofileUpgrade`, see “`WBIPprofileUpgrade` script” on page 18.
8. If you are migrating a server that has Business Process Choreographer configured and you need to upgrade its associated database manually, do this step now. See “Upgrading the Business Process Choreographer database manually” on page 29.
9. Start all the new (migrated) node agents and servers that are members of the cluster.

Results

The cluster is now migrated to the new version of WebSphere Process Server.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
 If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Related tasks

“Migrating a deployment manager” on page 40

Migrate a WebSphere Process Server deployment manager, choosing from several methods depending on your needs.

“Migrating non-clustered managed nodes” on page 49

Migrate a WebSphere Process Server managed node, choosing from several methods depending on your needs.

“Migrating a cluster with minimal down time” on page 69

To migrate a cluster while minimizing down time, first migrate approximately half of the profiles contributing to the cluster, then migrate the other half. Perform the extra steps required for cluster migration after you have migrated the first set of profiles.

“Upgrading the Business Process Choreographer database manually” on page 29

After migrating a server that has Business Process Choreographer configured, the schema for the associated Business Process Choreographer database must be upgraded before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

Related reference

“WBIPProfileUpgrade script” on page 18

Use the WBIPProfileUpgrade script to update application and configuration settings in a WebSphere Process Server profile when you are migrating clusters and in some other special situations.

Migrating cluster members using the migration wizard:

Migrate cluster members from an older version to a newer version of WebSphere Process Server using the migration wizard.

Before you begin

Note: These instructions are part of the larger procedure to migrate all the servers in your cluster. Follow the instructions in “Migrating a cluster” on page 60 or “Migrating a cluster with minimal down time” on page 69 before performing the steps described here.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.
- The deployment manager that manages the managed node that you intend to migrate has already been migrated to the newer version of WebSphere Process Server, and is running.

Note: Migrating a WebSphere Process Server version 6.0.x managed node to a version 6.1 managed node requires that you first migrate the version 6.0.x deployment manager to a version 6.1 deployment manager. See “Migrating a deployment manager” on page 40 for instructions. Complete the migration of the deployment manager before you proceed with the instructions in this topic.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.

- Use the `backupConfig` command or your own preferred backup utility to back up the managed node configuration if you want to be able to restore it to its previous state after migration. See the `backupConfig` command. Make sure that you note the exact name and location of this backed-up configuration.

About this task

After you migrate an older version deployment manager to a newer version of WebSphere Process Server, the newer version deployment manager runs in compatibility mode by default, where it can manage both older and newer versions of WebSphere Process Server. For example, after migration, a version 6.1 deployment manager can manage both version 6.0.x and version 6.1 release nodes. The managed nodes of the version 6.0.x deployment manager are now running as version 6.0.x managed nodes in the version 6.1 deployment manager. See *Coexisting with other WebSphere product installations* for information on restrictions on using mixed-release cells.

Over time, migrate each version 6.0.x WebSphere Process Server managed node (server managed by a version 6.1 deployment manager) to a version 6.1 managed node. After migrating all the version 6.0.x managed nodes, use the `convertScriptCompatibility` script to change the deployment manager from a node that supports backward compatibility of version 6.0.x administration scripts to a node that supports only version 6.1. See the `convertScriptCompatibility` command.

Note: When following the directions at this link to use the `convertScriptCompatibility` command, use the `WBIPostUpgrade` command rather than the `WASPostUpgrade` command.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the `stopServer` command from the `profile_dir/bin` directory for the profile of the affected server, or stop the server from the profile’s First steps console.

For more information about the `stopServer` command see the `stopServer` command.

Linux For example, issue the following command to stop `server1` on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the `-username` and `-password` parameters on the `stopServer` command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop `server1` on a Windows system:

```
stopServer.bat server1
```

On the Windows operating system, even if security is enabled, the `-username` and `-password` parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Stop the node agent of the node to be migrated. Issue one of the following commands to stop the nodeagent process, depending on platform (where *profile_root* represents the installation directory of the federated node):

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat*

If security is enabled, use one of the following commands instead:

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh -username user_ID -password password*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat -username user_ID -password password*

Note: A node can either be stopped or running during migration. You can migrate a managed node without stopping it (stopping its node agent); however, you must stop the old node before you can start the new one (the node running the newer version of WebSphere Process Server. At the same time, it is not necessary to have the node running to migrate its configuration. The migration tools can retrieve all the configuration data while the node is stopped.

4. Optional: Create a new version 6.1 profile with tools outside the migration wizard.

You can create either a version 6.1 stand-alone server profile or a custom profile as a target. If you create a custom profile, do not federate the node before migration. The migration tools federate the node during migration.

Note: For migration to be successful, you must use the same node name and cell name for each node that you migrate from version 6.0.x to version 6.1.

Tip: If you make any cell-level changes to the new version 6.1 node before migration, such as changes to virtual-host information, these changes will be lost during migration. Therefore, you should wait until after the node has been migrated before making any such changes. Otherwise, you will have to manually remake all of the changes, such as any changes to the virtual-host and host-alias information, to the new cell after migration using the administrative console running on the deployment manager.

5. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the *-currentWebSphereDirectory* parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the *-backupDirectory* parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

6. Ensure that the version 6.1 deployment manager is up and running.

Note: You can migrate a version 6.0.x node whether the node is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the version 6.0.x node before you can start the version 6.1 node that you are installing, however, so it makes sense to stop it now.

7. Invoke the migration wizard.

Invoke the migration wizard in one of the following ways:

- From the WebSphere Process Server First Steps console, select **Migration wizard**.
- Run one of the following scripts (depending upon your operating system) stored in the `install_dir/bin` directory:
 - a. `UNIX` `Linux` On UNIX and Linux systems: `wbi_migration.sh`
 - b. `Windows` On Windows systems: `wbi_migration.bat`

The migration wizard copies the configuration and applications from the version 6.0.x managed node to the version 6.1 managed node. After migrating all of the data, the wizard federates the version 6.1 managed node into the deployment manager cell.

8. Repeat steps 1-7 on page 52 for each cluster member you wish to migrate.
9. If you chose the compatibility option (which is the default), and if all of your nodes are completely migrated to WebSphere Process Server version 6.1, run the `convertScriptCompatibility` script to remove backward compatibility from the version 6.1 deployment manager.

Issue the `convertScriptCompatibility` command from the `bin` directory.

- `UNIX` `Linux` `install_root/bin/convertScriptCompatibility.sh`
- `Windows` `install_root\bin\convertScriptCompatibility.bat`

See the `convertScriptCompatibility` command.

Results

The cluster member profiles are now migrated.

Complete the cluster migration by completing steps 6-9 in “Migrating a cluster” on page 60 or steps 7-12 in “Migrating a cluster with minimal down time” on page 69.

Related tasks

“Migrating cluster members using the command-line tools”

Migrate cluster members from an older version to a newer version of WebSphere Process Server with the command-line tools.

Migrating cluster members using the command-line tools:

Migrate cluster members from an older version to a newer version of WebSphere Process Server with the command-line tools.

Before you begin

Note: These instructions are part of the larger procedure to migrate all the servers in your cluster. Follow the instructions in “Migrating a cluster” on page 60 or “Migrating a cluster with minimal down time” on page 69 before performing the steps described here.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.
- The deployment manager that manages the managed node that you intend to migrate has already been migrated to the newer version of WebSphere Process Server, and is running.

Note: Migrating a WebSphere Process Server version 6.0.x managed node to a version 6.1 managed node requires that you first migrate the version 6.0.x deployment manager to a version 6.1 deployment manager. See “Migrating a deployment manager” on page 40 for instructions. Complete the migration of the deployment manager before you proceed with the instructions in this topic.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up the managed node configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

About this task

After you migrate an older version deployment manager to a newer version of WebSphere Process Server, the newer version deployment manager runs in compatibility mode by default, where it can manage both older and newer versions of WebSphere Process Server. For example, after migration, a version 6.1 deployment manager can manage both version 6.0.x and version 6.1 release nodes. The managed nodes of the version 6.0.x deployment manager are now running as version 6.0.x managed nodes in the version 6.1 deployment manager. See Coexisting with other WebSphere product installations for information on restrictions on using mixed-release cells.

Over time, migrate each version 6.0.x WebSphere Process Server managed node (server managed by a version 6.1 deployment manager) to a version 6.1 managed

node. After migrating all the version 6.0.x managed nodes, use the **convertScriptCompatibility** script to change the deployment manager from a node that supports backward compatibility of version 6.0.x administration scripts to a node that supports only version 6.1. See the **convertScriptCompatibility** command.

Note: When following the directions at this link to use the **convertScriptCompatibility** command, use the **WBIPostUpgrade** command rather than the **WASPostUpgrade** command.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the **stopServer** command from the *profile_dir/bin* directory for the profile of the affected server, or stop the server from the profile’s First steps console. For more information about the **stopServer** command see the **stopServer** command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the **-username** and **-password** parameters on the **stopServer** command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:
stopServer.bat server1

On the Windows operating system, even if security is enabled, the **-username** and **-password** parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Stop the node agent of the node to be migrated. Issue one of the following commands to stop the nodeagent process, depending on platform (where *profile_root* represents the installation directory of the federated node):

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat*

If security is enabled, use one of the following commands instead:

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh -username user_ID -password password*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat -username user_ID -password password*

Note: A node can either be stopped or running during migration. You can migrate a managed node without stopping it (stopping its node agent); however, you must stop the old node before you can start the new one (the node running the newer version of WebSphere Process Server. At the same time, it is not necessary to have the node running to migrate its configuration. The migration tools can retrieve all the configuration data while the node is stopped.

4. Create a new profile. Use the Profile Management tool or the `manageprofiles` command to create a new profile with the newer version of WebSphere Process Server.

You can create either a version 6.1 stand-alone server profile or a custom profile as a target. If you create a custom profile, do not federate the node before migration. The migration tools federate the node during migration.

You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured.

Note: For migration to be successful, you must use the same node name and cell name for each node that you migrate from version 6.0.x to version 6.1.

Tip: If you make any cell-level changes to the new version 6.1 node before migration, such as changes to virtual-host information, these changes will be lost during migration. Therefore, you should wait until after the node has been migrated before making any such changes. Otherwise, you will have to manually remake all of the changes, such as any changes to the virtual-host and host-alias information, to the new cell after migration using the administrative console running on the deployment manager.

5. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

6. Ensure that the version 6.1 deployment manager is up and running.

Note: You can migrate a version 6.0.x node whether the node is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the version 6.0.x node before you can start the version 6.1 node that you are installing, however, so it makes sense to stop it now.

7. Run the `WBIPreUpgrade` command, specifying the migration backup directory name and the existing WebSphere Process Server directory name. The `WBIPreUpgrade` tool saves selected files from the `install_root` and `profile_root` directories to a backup directory you specify. See the “`WBIPreUpgrade` command” on page 12 for details.
8. Run the `WBIPostUpgrade` command, specifying the migration backup directory. The `WBIPostUpgrade` tool restores the environment in the backup directory into the new WebSphere Process Server stand-alone server installation. See the “`WBIPostUpgrade` command” on page 14 for details.
9. Repeat steps 1-8 (with the possible exception of step 7).

Note: You need to perform step 7 (running `WBIPreUpgrade`) again only if the version 6.0.x system has been re-configured since the first time you ran `WBIPreUpgrade`.

10. If you chose the compatibility option (which is the default), and if all of your nodes are completely migrated to WebSphere Process Server version 6.1, run the `convertScriptCompatibility` script to remove backward compatibility from the version 6.1 deployment manager.

Issue the `convertScriptCompatibility` command from the `bin` directory.

- `UNIX` `Linux` `install_root/bin/convertScriptCompatibility.sh`
- `Windows` `install_root\bin\convertScriptCompatibility.bat`

See the `convertScriptCompatibility` command.

Results

The cluster member profiles are now migrated.

Complete the cluster migration by completing steps 6-9 in “Migrating a cluster” on page 60 or steps 7-12 in “Migrating a cluster with minimal down time.”

Related tasks

“Migrating cluster members using the migration wizard” on page 62
Migrate cluster members from an older version to a newer version of WebSphere Process Server using the migration wizard.

Migrating a cluster with minimal down time

To migrate a cluster while minimizing down time, first migrate approximately half of the profiles contributing to the cluster, then migrate the other half. Perform the extra steps required for cluster migration after you have migrated the first set of profiles.

Before you begin

Note: In a cluster, version 6.0.x members and version 6.1 members must never run at the same time. All version 6.0.x cluster members must be stopped before you start the first version 6.1 cluster member. Also, once you start a version 6.1 cluster member, do not start any version 6.0.x cluster members in that cluster.

You must have an existing cell containing at least one cluster running on an older version of WebSphere Process Server (for example, version 6.0.x) that you wish to migrate to a newer version (for example, version 6.1). In addition, you must have installed the new version of WebSphere Process Server.

About this task

Following these steps will ensure that you retain cluster functionality on the new version of WebSphere Process Server.

Procedure

1. Migrate the deployment manager. Follow one of the instruction sets listed in “Migrating a deployment manager” on page 40 to complete this task.
2. Make sure that the new deployment manager is running.
3. Identify the profiles involved.
 - a. Identify an older version profile that contains cluster members.
 - b. Identify which other clusters this profile contributes to; that is, if the profile defines servers that are members of any other clusters, identify those clusters.
 - c. Identify all other profiles in the same cell that contribute cluster members to any of the clusters identified in step 3b.
 - d. Identify all the node agents and process servers defined by any of the profiles identified in step 3c.

All of the profiles identified in step 3c, and all of the corresponding node agents and servers identified in step 3d will be involved in the migration.

4. Define two groups of profiles out of the full set of profiles identified in step 3 on page 60. Divide the profiles into roughly half (If the total number of profiles is an odd number, one group will consist of one more in number than the other group). You will migrate one set of servers while the other set is still running, thus reducing the amount of time when all servers in the cluster are stopped.
5. Stop all the node agents and servers that are defined by the first set of profiles that you will migrate.
6. Migrate each profile in the first set, one at a time, but **do not start** any new agents or servers. Follow one of the instruction sets listed in either “Migrating cluster members using the migration wizard” on page 62 or “Migrating cluster members using the command-line tools” on page 65.
7. Stop the remaining node agents and servers; that is, those defined by the second set of profiles. This action starts the period of time when cluster services will be unavailable.
8. On the system that hosts the WebSphere Process Server version 6.1 deployment manager profile, navigate to the *install_dir/util* directory. This directory contains the *WBIPProfileUpgrade* script, *WBIPProfileUpgrade.ant*.
9. Run the *WBIPProfileUpgrade* script for each cluster defined in the profiles that have been migrated so far. (That is, run *WBIPProfileUpgrade* for each cluster defined in step 3 on page 60.) For instructions on running *WBIPProfileUpgrade*, see “*WBIPProfileUpgrade* script” on page 18.
10. If you are migrating a server that has Business Process Choreographer configured and you need to upgrade its associated database manually, do this step now. See “Upgrading the Business Process Choreographer database manually” on page 29.
11. Start all the new (migrated) node agents and servers; that is, the node agents and servers corresponding to the profiles that have been migrated so far.
12. Migrate each profile in the second set of profiles. As with the first set, to migrate, follow one of the instruction sets listed in either “Migrating cluster members using the migration wizard” on page 62 or “Migrating cluster

members using the command-line tools” on page 65. This time, you can start the migrated node agents and servers as you proceed to migrate each managed node.

Results

The cluster is now migrated to the new version of WebSphere Process Server.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Related tasks

“Migrating a deployment manager” on page 40

Migrate a WebSphere Process Server deployment manager, choosing from several methods depending on your needs.

“Migrating non-clustered managed nodes” on page 49

Migrate a WebSphere Process Server managed node, choosing from several methods depending on your needs.

“Migrating a cluster” on page 60

To migrate a cluster, migrate each profile containing a member of that cluster one at a time. The migration requires extra steps not required for a non-clustered environment.

“Upgrading the Business Process Choreographer database manually” on page 29

After migrating a server that has Business Process Choreographer configured, the schema for the associated Business Process Choreographer database must be upgraded before you start the server. You must upgrade manually if the database user that is defined for the data source does not have sufficient authorization to modify the database schema.

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

Related reference

“WBIPProfileUpgrade script” on page 18

Use the WBIPProfileUpgrade script to update application and configuration settings in a WebSphere Process Server profile when you are migrating clusters and in some other special situations.

Migrating cluster members using the migration wizard:

Migrate cluster members from an older version to a newer version of WebSphere Process Server using the migration wizard.

Before you begin

Note: These instructions are part of the larger procedure to migrate all the servers in your cluster. Follow the instructions in “Migrating a cluster” on page 60 or “Migrating a cluster with minimal down time” on page 69 before performing the steps described here.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A federated profile, created with the older WebSphere Process Server version resides on the same system.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.
- The deployment manager that manages the managed node that you intend to migrate has already been migrated to the newer version of WebSphere Process Server, and is running.

Note: Migrating a WebSphere Process Server version 6.0.x managed node to a version 6.1 managed node requires that you first migrate the version 6.0.x deployment manager to a version 6.1 deployment manager. See “Migrating a deployment manager” on page 40 for instructions. Complete the migration of the deployment manager before you proceed with the instructions in this topic.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up the managed node configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

About this task

After you migrate an older version deployment manager to a newer version of WebSphere Process Server, the newer version deployment manager runs in compatibility mode by default, where it can manage both older and newer versions of WebSphere Process Server. For example, after migration, a version 6.1 deployment manager can manage both version 6.0.x and version 6.1 release nodes. The managed nodes of the version 6.0.x deployment manager are now running as version 6.0.x managed nodes in the version 6.1 deployment manager. See Coexisting with other WebSphere product installations for information on restrictions on using mixed-release cells.

Over time, migrate each version 6.0.x WebSphere Process Server managed node (server managed by a version 6.1 deployment manager) to a version 6.1 managed

node. After migrating all the version 6.0.x managed nodes, use the **convertScriptCompatibility** script to change the deployment manager from a node that supports backward compatibility of version 6.0.x administration scripts to a node that supports only version 6.1. See the **convertScriptCompatibility** command.

Note: When following the directions at this link to use the **convertScriptCompatibility** command, use the **WBIPostUpgrade** command rather than the **WASPostUpgrade** command.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the **stopServer** command from the *profile_dir/bin* directory for the profile of the affected server, or stop the server from the profile’s First steps console.

For more information about the **stopServer** command see the **stopServer** command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the **-username** and **-password** parameters on the **stopServer** command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:

```
stopServer.bat server1
```

On the Windows operating system, even if security is enabled, the **-username** and **-password** parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Stop the node agent of the node to be migrated. Issue one of the following commands to stop the nodeagent process, depending on platform (where *profile_root* represents the installation directory of the federated node):

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat*

If security is enabled, use one of the following commands instead:

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh -username user_ID -password password*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat -username user_ID -password password*

Note: A node can either be stopped or running during migration. You can migrate a managed node without stopping it (stopping its node agent); however, you must stop the old node before you can start the new one (the node running the newer version of WebSphere Process Server. At the same time, it is not necessary to have the node running to migrate its configuration. The migration tools can retrieve all the configuration data while the node is stopped.

4. Optional: Create a new version 6.1 profile with tools outside the migration wizard.

You can create either a version 6.1 stand-alone server profile or a custom profile as a target. If you create a custom profile, do not federate the node before migration. The migration tools federate the node during migration.

Note: For migration to be successful, you must use the same node name and cell name for each node that you migrate from version 6.0.x to version 6.1.

Tip: If you make any cell-level changes to the new version 6.1 node before migration, such as changes to virtual-host information, these changes will be lost during migration. Therefore, you should wait until after the node has been migrated before making any such changes. Otherwise, you will have to manually remake all of the changes, such as any changes to the virtual-host and host-alias information, to the new cell after migration using the administrative console running on the deployment manager.

5. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

6. Ensure that the version 6.1 deployment manager is up and running.

Note: You can migrate a version 6.0.x node whether the node is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the version 6.0.x node before you can start the version 6.1 node that you are installing, however, so it makes sense to stop it now.

7. Invoke the migration wizard.

Invoke the migration wizard in one of the following ways:

- From the WebSphere Process Server First Steps console, select **Migration wizard**.

- Run one of the following scripts (depending upon your operating system) stored in the *install_dir/bin* directory:

- UNIX** **Linux** On UNIX and Linux systems: `wbi_migration.sh`
- Windows** On Windows systems: `wbi_migration.bat`

The migration wizard copies the configuration and applications from the version 6.0.x managed node to the version 6.1 managed node. After migrating all of the data, the wizard federates the version 6.1 managed node into the deployment manager cell.

8. Repeat steps 1-7 on page 52 for each cluster member you wish to migrate.
9. If you chose the compatibility option (which is the default), and if all of your nodes are completely migrated to WebSphere Process Server version 6.1, run the `convertScriptCompatibility` script to remove backward compatibility from the version 6.1 deployment manager.

Issue the `convertScriptCompatibility` command from the bin directory.

- **UNIX** **Linux** `install_root/bin/convertScriptCompatibility.sh`
- **Windows** `install_root\bin\convertScriptCompatibility.bat`

See the `convertScriptCompatibility` command.

Results

The cluster member profiles are now migrated.

Complete the cluster migration by completing steps 6-9 in “Migrating a cluster” on page 60 or steps 7-12 in “Migrating a cluster with minimal down time” on page 69.

Related tasks

“Migrating cluster members using the command-line tools” on page 65
Migrate cluster members from an older version to a newer version of WebSphere Process Server with the command-line tools.

Migrating cluster members using the command-line tools:

Migrate cluster members from an older version to a newer version of WebSphere Process Server with the command-line tools.

Before you begin

Note: These instructions are part of the larger procedure to migrate all the servers in your cluster. Follow the instructions in “Migrating a cluster” on page 60 or “Migrating a cluster with minimal down time” on page 69 before performing the steps described here.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- If you are migrating on the same physical computer system on which the older version of WebSphere Process Server resides, you have installed the new version of WebSphere Process Server on the same system.
- A federated profile, created with the older WebSphere Process Server version resides on the same system.

- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.
- The deployment manager that manages the managed node that you intend to migrate has already been migrated to the newer version of WebSphere Process Server, and is running.

Note: Migrating a WebSphere Process Server version 6.0.x managed node to a version 6.1 managed node requires that you first migrate the version 6.0.x deployment manager to a version 6.1 deployment manager. See “Migrating a deployment manager” on page 40 for instructions. Complete the migration of the deployment manager before you proceed with the instructions in this topic.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up the managed node configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

About this task

After you migrate an older version deployment manager to a newer version of WebSphere Process Server, the newer version deployment manager runs in compatibility mode by default, where it can manage both older and newer versions of WebSphere Process Server. For example, after migration, a version 6.1 deployment manager can manage both version 6.0.x and version 6.1 release nodes. The managed nodes of the version 6.0.x deployment manager are now running as version 6.0.x managed nodes in the version 6.1 deployment manager. See Coexisting with other WebSphere product installations for information on restrictions on using mixed-release cells.

Over time, migrate each version 6.0.x WebSphere Process Server managed node (server managed by a version 6.1 deployment manager) to a version 6.1 managed node. After migrating all the version 6.0.x managed nodes, use the **convertScriptCompatibility** script to change the deployment manager from a node that supports backward compatibility of version 6.0.x administration scripts to a node that supports only version 6.1. See the convertScriptCompatibility command.

Note: When following the directions at this link to use the convertScriptCompatibility command, use the WBIPostUpgrade command rather than the WASPostUpgrade command.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Log on as the root user on a Linux or UNIX system, or as a member of the Administrator group on a Windows system.
2. Stop the version 6.0.x server if it is running on the node to be migrated. Use the stopServer command from the *profile_dir*/bin directory for the profile of the affected server, or stop the server from the profile’s First steps console.

For more information about the stopServer command see the stopServer command.

Linux For example, issue the following command to stop server1 on a Linux system:

```
./stopServer.sh server1
```

If security is enabled, specify the -username and -password parameters on the stopServer command. The user name provided must be a member of the operator or administrator role.

Windows Issue the following command to stop server1 on a Windows system:
stopServer.bat server1

On the Windows operating system, even if security is enabled, the -username and -password parameters do not have to be specified if the server is running as a Windows service. In this case, the parameters are automatically passed into the script that the Windows service uses to shut down the system.

Note: A server can either be stopped or running during migration. You can migrate a server without stopping it; however, you must stop the old server before you can start the new one (the server running the newer version of WebSphere Process Server). At the same time, it is not necessary to have the server running to migrate its configuration. The migration tools can retrieve all the configuration data while the server is stopped.

3. Stop the node agent of the node to be migrated. Issue one of the following commands to stop the nodeagent process, depending on platform (where *profile_root* represents the installation directory of the federated node):

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat*

If security is enabled, use one of the following commands instead:

- **Linux** **UNIX** **On Linux and UNIX platforms:** *profile_root/bin/stopNode.sh -username user_ID -password password*

- **Windows** **On Windows platforms:** *profile_root\bin\stopNode.bat -username user_ID -password password*

Note: A node can either be stopped or running during migration. You can migrate a managed node without stopping it (stopping its node agent); however, you must stop the old node before you can start the new one (the node running the newer version of WebSphere Process Server. At the same time, it is not necessary to have the node running to migrate its configuration. The migration tools can retrieve all the configuration data while the node is stopped.

4. Create a new profile. Use the Profile Management tool or the manageprofiles command to create a new profile with the newer version of WebSphere Process Server.

You can create either a version 6.1 stand-alone server profile or a custom profile as a target. If you create a custom profile, do not federate the node before migration. The migration tools federate the node during migration.

You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured.

Note: For migration to be successful, you must use the same node name and cell name for each node that you migrate from version 6.0.x to version 6.1.

Tip: If you make any cell-level changes to the new version 6.1 node before migration, such as changes to virtual-host information, these changes will be lost during migration. Therefore, you should wait until after the node has been migrated before making any such changes. Otherwise, you will have to manually remake all of the changes, such as any changes to the virtual-host and host-alias information, to the new cell after migration using the administrative console running on the deployment manager.

5. Identify, in advance, the pre-existing information required by the migration wizard, as listed below:

Installation root directory

See “WBIPreUpgrade command” on page 12 for a description of the `-currentWebSphereDirectory` parameter.

Migration backup directory name

See “WBIPreUpgrade command” on page 12 for a description of the `-backupDirectory` parameter.

Source profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-oldProfile` parameter.

Target profile name

See “WBIPostUpgrade command” on page 14 for a description of the `-profileName` parameter.

Port value assignments (optional)

See “WBIPostUpgrade command” on page 14 for a description of the `-replacePorts` and `-portBlock` parameters.

6. Ensure that the version 6.1 deployment manager is up and running.

Note: You can migrate a version 6.0.x node whether the node is running or stopped. The migration tools can retrieve all the configuration data either way. You must stop the version 6.0.x node before you can start the version 6.1 node that you are installing, however, so it makes sense to stop it now.

7. Run the WBIPreUpgrade command, specifying the migration backup directory name and the existing WebSphere Process Server directory name. The WBIPreUpgrade tool saves selected files from the `install_root` and `profile_root` directories to a backup directory you specify. See the “WBIPreUpgrade command” on page 12 for details.
8. Run the WBIPostUpgrade command, specifying the migration backup directory. The WBIPostUpgrade tool restores the environment in the backup directory into the new WebSphere Process Server stand-alone server installation. See the “WBIPostUpgrade command” on page 14 for details.
9. Repeat steps 1-8 (with the possible exception of step 7).

Note: You need to perform step 7 (running WBIPreUpgrade) again only if the version 6.0.x system has been re-configured since the first time you ran WBIPreUpgrade.

10. If you chose the compatibility option (which is the default), and if all of your nodes are completely migrated to WebSphere Process Server version 6.1, run the `convertScriptCompatibility` script to remove backward compatibility from the version 6.1 deployment manager.

Issue the `convertScriptCompatibility` command from the `bin` directory.

- **UNIX** **Linux** `install_root/bin/convertScriptCompatibility.sh`
 - **Windows** `install_root\bin\convertScriptCompatibility.bat`
- See the `convertScriptCompatibility` command.

Results

The cluster member profiles are now migrated.

Complete the cluster migration by completing steps 6-9 in “Migrating a cluster” on page 60 or steps 7-12 in “Migrating a cluster with minimal down time” on page 69.

Related tasks

“Migrating cluster members using the migration wizard” on page 62
 Migrate cluster members from an older version to a newer version of WebSphere Process Server using the migration wizard.

Migrating to a remote system

Use the migration tools to migrate from an older version on one system to a newer version of WebSphere Process Server on a different, remote system. (Stand-alone servers only.)

Before you begin

Note: This procedure is supported for stand-alone servers only.

Make sure that the following conditions are met before you start the migration process:

- Your new (remote) system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the `backupConfig` command or your own preferred backup utility to back up your existing configuration if you want to be able to restore it to its previous state after migration. See the `backupConfig` command. Make sure that you note the exact name and location of this backed-up configuration.

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

Typically, you can use the migration tools to upgrade from an older version to a newer version on the same system. However, some scenarios require that you migrate the old version’s configuration on one machine to the new version of WebSphere Process Server on a different system. One of these scenarios is when you install new systems for your environment based on the newer version but need to migrate your existing older configuration from other systems.

If the node that you are migrating is part of a deployment cell, migrate the older deployment manager to the newer WebSphere Process Server version as described in “Migrating a deployment manager using the migration wizard” on page 41 or

“Migrating a deployment manager using command-line tools” on page 45 first before continuing this procedure. The release level of the deployment manager must always be at the same or higher release level as any nodes within the cell that it manages.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

About this task

The **WBIPreUpgrade** command saves the existing older version configuration into a migration-specific backup directory. The **WBIPostUpgrade** command uses this directory to add the old configuration settings to the newer version environment.

Procedure

1. Obtain the WebSphere Process Server CD-ROMs and (if you have DVD media) the WebSphere Process Server DVD for the version to which you are migrating (for example version 6.1).

For WebSphere Process Server 6.1 you will need the following:

- WebSphere Process Server version 6.1 , Disk 1 of 2, *or* WebSphere Process Server version 6.1 DVD
- WebSphere Process Server version 6.1, WebSphere Application Server Network Deployment Supplements

You will need to copy files from the product installation media to a directory on the system hosting your older version (version 6.0.x) of WebSphere Process Server. The files in these directories constitute a special environment that you can use to run the **WBIPreUpgrade** command without installing the newer version of WebSphere Process Server.

2. Create a new directory on your local system hosting your older version (version 6.0.x) of WebSphere Process Server called `migration_copy`. (This directory can be any name you choose. We are using “`migration_copy`” here for explanatory purposes.)
3. Copy the files from the product media into the directory you just made, `migration_copy`, on your local version 6.0.x system. Copy the entire directories listed below:
 - a. On either the WebSphere Process Server version 6.1 , Disk 1 of 2 CD-ROM, *or* WebSphere Process Server version 6.1 DVD, the directory called `migration`.
 - b. On the WebSphere Process Server version 6.1, WebSphere Application Server Network Deployment Supplements CD, the directory called `migration`.

Note: The files in this directory and the files in the `migration` directory from the WebSphere Process Server version 6.1 media should combine into a single migration directory on your system.

- c. On the WebSphere Process Server version 6.1, WebSphere Application Server Network Deployment Supplements CD, the directory called `JDK`.

You should now have the following directory structure on your system:

```
migration_copy/  
  migration/  
    JDK/
```

4. Navigate to the `migration_copy/migration/bin` directory.

5. Save the current configuration using the **WBIPreUpgrade** script from the `migration_copy/bin` directory. Save the configuration in the migration-specific backup directory on the system hosting the older version. See the “WBIPreUpgrade command” on page 12 for details.

Linux

UNIX

On Linux or UNIX:

```
./WBIPreUpgrade.sh /filepath/migration_specific_backup  
currentWebSphereDirectory
```

Windows

On Windows:

```
WBIPreUpgrade C:\filepath\migration_specific_backup  
currentWebSphereDirectory
```

See the “WBIPreUpgrade command” on page 12 for parameter descriptions. The **WBIPreUpgrade** command provides status to the screen and to log files in the `migration_specific_backup` directory. ASCII log file names start with the text `WBIPreUpgrade` and include a date and timestamp.

6. Copy the `migration_specific_backup` directory from the older version WebSphere Process Server system to the newer version system.
Use the **ftp** command, shared storage, or some other mechanism to copy the directory to the new machine.
7. Install the newer version of WebSphere Process Server .
See *Installing and configuring WebSphere Process Server* for installation information.
8. Use the Profile Management Tool or the **manageprofiles** command to create a new profile with the newer version of WebSphere Process Server. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured. See *Creating profiles* for information about creating profiles.
9. Add the configuration from the older version of WebSphere Process Server configuration to the newer version’s configuration using the **WBIPostUpgrade** command. Use the **WBIPostUpgrade** command in the `install_root/bin` directory of the new installation to add the older version’s configuration (which you copied into a directory on the new system in step 6) to the newer version’s configuration.

See the “WBIPostUpgrade command” on page 14 for details.

Linux

UNIX

Linux or UNIX:

```
./WBIPostUpgrade.sh /filepath/migration_specific_backup
```

Windows

Windows:

```
WBIPostUpgrade C:\filepath\migration_specific_backup
```

The **WBIPostUpgrade** tool records detailed information specific to each enterprise bean it deploys in the `migration_specific_backup/WBIPostUpgrade.log` file.

10. Modify the configuration using the newer WebSphere Process Server version’s administrative console.
 - a. Change user IDs and passwords to match security requirements.
You might have to change user IDs and passwords if they are not identical to those in use on the system hosting the older version.
 - b. Change other system-specific information.

The configuration might refer to other software products or configurations that do not exist on the new system. For example, the old system might have a database. Modify the data source to point to the database on the old system.

11. If required, manually update the databases used by WebSphere Process Server. Normally, database changes required by new versions of WebSphere Process Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, in cases in which the server has insufficient permissions to access the database schema, or other database-specific requirements are not met, you must update the database manually. For more information see “Upgrading databases for migration” on page 27.

Results

You have migrated WebSphere Process Server from the older version to a remote system hosting the newer version.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Overview of migrating” on page 1

Migrate from earlier versions of WebSphere Process Server and WebSphere Enterprise Service Bus.

Related tasks

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.



Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

“Migrating from an operating system that is no longer supported”
Use the migration tools to migrate an earlier WebSphere Process Server release that is running on an operating system that the newer version does not support. (Stand-alone servers only.)

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.

WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

Installing and configuring WebSphere Process Server

This section describes how to prepare for, install, and configure an installation of IBM WebSphere Process Server. Instructions are provided for Linux, i5/OS, UNIX, and Windows systems.

Migrating from an operating system that is no longer supported

Use the migration tools to migrate an earlier WebSphere Process Server release that is running on an operating system that the newer version does not support. (Stand-alone servers only.)

Before you begin

Note: This procedure is supported for stand-alone servers only.

Make sure that the following conditions are met before you start the migration process:

- Your system meets all the hardware and software requirements for the new version of WebSphere Process Server.
- Sufficient disk space is available for the migrated profile and its backup. See “Premigration considerations” on page 2 for disk space requirements.

Make sure that you have completed the following tasks before you start the migration process:

- Back up the databases that support version 6.0.x WebSphere Process Server components.
- Use the backupConfig command or your own preferred backup utility to back up your existing configuration if you want to be able to restore it to its previous state after migration. See the backupConfig command. Make sure that you note the exact name and location of this backed-up configuration.

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Obtain the WebSphere Process Server CD-ROMs and (if you have DVD media) the WebSphere Process Server DVD for the version to which you are migrating (for example version 6.1).

For WebSphere Process Server 6.1 you will need the following:

- WebSphere Process Server version 6.1 , Disk 1 of 2, *or* WebSphere Process Server version 6.1 DVD
- WebSphere Process Server version 6.1, WebSphere Application Server Network Deployment Supplements

You will need to copy files from the product installation media to a directory on the system hosting your older version (version 6.0.x) of WebSphere Process Server. The files in these directories constitute a special environment that you can use to run the **WBIPreUpgrade** command without installing the newer version of WebSphere Process Server.

2. Create a new directory on your local system hosting your older version (version 6.0.x) of WebSphere Process Server called `migration_copy`. (This directory can be any name you choose. We are using “`migration_copy`” here for explanatory purposes.)
3. Copy the files from the product media into the directory you just made, `migration_copy`, on your local version 6.0.x system. Copy the entire directories listed below:
 - a. On either the WebSphere Process Server version 6.1 , Disk 1 of 2 CD-ROM, *or* WebSphere Process Server version 6.1 DVD, the directory called `migration`.
 - b. On the WebSphere Process Server version 6.1, WebSphere Application Server Network Deployment Supplements CD, the directory called `migration`.

Note: The files in this directory and the files in the `migration` directory from the WebSphere Process Server version 6.1 media should combine into a single migration directory on your system.

- c. On the WebSphere Process Server version 6.1, WebSphere Application Server Network Deployment Supplements CD, the directory called `JDK`.

You should now have the following directory structure on your system:

```
migration_copy/  
  migration/  
    JDK/
```

4. Navigate to the `migration_copy/migration/bin` directory.
5. Save the current configuration using the **WBIPreUpgrade** script from the `migration_copy/bin` directory. Save the configuration in the migration-specific backup directory on the system hosting the older version. See the “WBIPreUpgrade command” on page 12 for details.

Linux **UNIX** On Linux or UNIX:

```
./WBIPreUpgrade.sh /filepath/migration_specific_backup  
currentWebSphereDirectory
```

Windows On Windows:

```
WBIPreUpgrade C:\filepath\migration_specific_backup  
currentWebSphereDirectory
```

See the “WBIPreUpgrade command” on page 12 for parameter descriptions. The **WBIPreUpgrade** command provides status to the screen and to log files in the *migration_specific_backup* directory. ASCII log file names start with the text WBIPreUpgrade and include a date and timestamp.

6. Shut down the older WebSphere Process Server release by stopping all server nodes in the configuration.
7. Compress the backup directory (with a utility such as TAR or ZIP), and use FTP to copy the file to another system.
8. Install the new operating system, keeping the same host name.
If possible, keep the system name and passwords the same as on the old system. Place any database files related to applications that you are migrating in the same path as on the previous system. In general, try to keep paths the same. If you must change paths or names, make any changes with the administrative console before running the **WBIPostUpgrade** command as described in a later step.
9. Use FTP to copy the backup directory from the other system, and extract it.
10. Install the newer version of WebSphere Process Server.
See Installing and configuring WebSphere Process Server.
11. Use the Profile Management Tool or the **manageprofiles** command to create a new profile with the newer version of WebSphere Process Server. You must create the new profile with the same node name, cell name, and augmentation levels as the version 6.0.x profile, and if possible, the same host name. If you use a different host name, you will need to use the administrative console after migration to update the host name manually in other places where it is configured. See Creating profiles for information about creating profiles.
12. Run the **WBIPostUpgrade** command from the newer version *install_root/bin* directory.
Specify the copy of the backup directory that you made in step 9. See “WBIPostUpgrade command” on page 14 for the correct command syntax. For example:

```
install_root\bin\WBIPostUpgrade wbi_installation\migration
```
13. If required, manually update the databases used by WebSphere Process Server. Normally, database changes required by new versions of WebSphere Process Server are made automatically. When the server is first started the database tables are migrated to the new schema version. However, in cases in which the server has insufficient permissions to access the database schema, or other database-specific requirements are not met, you must update the database manually. For more information see “Upgrading databases for migration” on page 27.

Results

You have now migrated your configuration to an operating system that supports WebSphere Process Server.

What to do next

Verify that the migration has been successful. If your server has Business Process Choreographer configured, see “Migration considerations for Business Process Choreographer” on page 105. Finally, perform the checks described in “Postmigration configuration checking” on page 88.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Overview of migrating” on page 1

Migrate from earlier versions of WebSphere Process Server and WebSphere Enterprise Service Bus.

Related tasks

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

“Postmigration configuration checking” on page 88

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

“Migrating to a remote system” on page 79

Use the migration tools to migrate from an older version on one system to a newer version of WebSphere Process Server on a different, remote system. (Stand-alone servers only.)



Creating profiles

Learn how to create new WebSphere Enterprise Service Bus or WebSphere Process Server profiles. You can create profiles from a command line by using the `manageprofiles` command, or interactively by using the Profile Management Tool graphical user interface (GUI).

Related reference

“WBIPreUpgrade command” on page 12

Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.



WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information



Installing and configuring WebSphere Process Server

This section describes how to prepare for, install, and configure an installation of IBM WebSphere Process Server. Instructions are provided for Linux, i5/OS, UNIX, and Windows systems.

Verifying migration

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

Before you begin

Make sure the server that has been migrated has been started.

Procedure

1. Check the migration log files for the `WBIPostUpgrade` command and the `WBIProfileUpgrade.ant` script.
 - a. Check the file `backupDirectory/logs/WBIPostUpgradetimestamp.log` for either of the following messages: (*backupDirectory* is the directory in which migrated data was first stored and later retrieved from during the migration process, as specified in the migration wizard or the `WBIPreUpgrade` or `WBIPostUpgrade` commands.)
 - `MIGR0259I`: The migration has successfully completed.
 - `MIGR0271W`: Migration completed successfully, with one or more warnings.
 - b. Check the file `backupDirectory/logs/WBIProfileUpgrade.anttimestamp.log` for the message `BUILD SUCCESSFUL`.

Both of these log files must indicate success, as described by these messages, for you to consider the migration successful.

2. Check the profile's log files.
 - a. Navigate to the `profile_root/logs/server_name` directory corresponding to the migrated profile.
 - b. Review the `SystemOut.log` file and make sure there are no fatal errors.
 - c. Review the `SystemErr.log` file and make sure there are no fatal errors.
3. Check the log files in the logs directory. For example, check the logs for a stand-alone server in the `/WebSphere/V6R1/AppServer/profiles/default/logs` directory.
4. Check operation with the administrative console.
 - a. Open the administrative console (Integrated Solutions Console).
 - b. Select **Applications > Enterprise Applications** from the navigation panel.
 - c. In the right-hand panel, verify that all of the applications listed have started, shown by the green "started" icon.
 - d. From the navigation panel, select **Resources > JDBC > Business Integration Data Sources**.
 - e. For each WebSphere Process Server data source listed on this panel, select the check box and then select **Test connection**.
 - f. For each data source, you should receive a message similar to the following: "The test connection operation for data source `WPS_DataSource` on server `Dmgr1` at node `Dmgr1Node1` was successful."

What to do next

If migration was successful, you can begin using the server. If the migration was not successful, refer to "Troubleshooting version-to-version migration" on page 107 for troubleshooting information.

Related tasks

“Rolling back your environment”

After migrating to a WebSphere Process Server version 6.1 environment, you can roll back to a version 6.0.x environment. This returns the configuration to the state that it was in before migration. After rolling back the environment, you can restart the migration process.

“Troubleshooting version-to-version migration” on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

Related information

 [Administering enterprise applications](#)

Use the console’s Enterprise Application page (viewed by clicking **Applications** > **Enterprise Applications**) to view and administer enterprise applications installed on the server.

Postmigration configuration checking

After migration, you should check some configuration settings. You might need to change them, or further configure the version 6.1 server.

Before you begin

You should have migrated your server or cluster and verified that the migration was successful.

About this task

Perform the following checks, if applicable to your environment:

- Examine any Lightweight Third Party Authentication (LTPA) security settings that you might have used in version 6.0.x, and make sure that version 6.1 security is set appropriately.
- Check the WBIPostUpgrade.log file in the logs directory for details about any JSP objects that the migration tools did not migrate.
If version 6.1 does not support a level for which JSP objects are configured, the migration tools recognize the objects in the output and log them.
- Review your Java virtual machine settings to verify that you are using the recommended heap sizes. See Java virtual machine settings. The information at this link applies to WebSphere Process Server servers as well as WebSphere Application Server servers.
- Verify the results of the automatic Cloudscape database migration, and manually migrate any Cloudscape databases that are not automatically migrated by the tools. See “Migrating Cloudscape databases” on page 94 for more information.

Rolling back your environment

After migrating to a WebSphere Process Server version 6.1 environment, you can roll back to a version 6.0.x environment. This returns the configuration to the state that it was in before migration. After rolling back the environment, you can restart the migration process.

About this task

Generally, migration does not modify anything in the configuration of the prior release; however, there are cases where minimal changes are made that are reversible—those of a deployment manager and its managed nodes.

The subtopics below provide further information for these cases.

Related concepts

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

Related tasks

“Verifying migration” on page 87

Verify that your migration was successful by checking the log files and checking operation with the administrative console.

Rolling back a deployment cell

You can use the **restoreConfig** and **wsadmin** commands to roll back a migrated WebSphere Process Server version 6.1 deployment cell to version 6.0.x. This returns the configuration to the state that it was in before migration. After rolling back the deployment cell, you can restart the migration process.

Before you begin

When migrating a version 6.0.x deployment cell, you must complete the following if you want to be able to roll it back to its previous state after migration:

1. Back up the databases that support WebSphere Process Server components.
2. Back up your existing configuration using the **backupConfig** command or your own preferred backup utility.
 - Run the **backupConfig** command or your own preferred utility to back up the version 6.0.x deployment manager configuration.

Important: Make sure that you note the exact name and location of this backed-up configuration.

See the **backupConfig** command in the WebSphere Application Server Network Deployment, version 6.1 information center.

- Run the **backupConfig** command or your own preferred utility to back up the version 6.0.x managed node configurations.

Important: Make sure that you note the exact name and location of each of these backed-up configurations.

See the **backupConfig** command in the WebSphere Application Server Network Deployment, version 6.1 information center.

3. Migrate the deployment cell.

Procedure

1. Stop all of the servers that are currently running in the WebSphere Process Server version 6.1 environment.
2. If you chose to disable the previous deployment manager when you migrated to the version 6.1 deployment manager, perform one of the following actions.
 - a. If you backed up your previous deployment manager configuration using the **backupConfig** command or your own preferred backup utility, run the **restoreConfig** command or your own preferred utility to restore the version 6.0.x configuration for the deployment manager.

Important: Make sure that you restore the same backed-up configuration that you created just before you migrated the deployment manager.

See the `restoreConfig` command in the WebSphere Application Server Network Deployment, version 6.1 information center.

- b. If you did not back up your previous deployment manager configuration, use the `wsadmin` command to run the `migrationDisablementReversal.jacl` script from the version 6.0.x `profile_root/bin` directory of the deployment manager that you need to roll back from version 6.1.

Linux In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migrationDisablementReversal.jacl -conntype NONE
```

Tip: If you have trouble running the `migrationDisablementReversal.jacl` script, try to manually perform the steps in the script.

- 1) Go to the following directory:
`profile_root/config/cells/cell_name/nodes/node_name`

where `node_name` is the name of the deployment manager node that you want to roll back.
 - 2) If you see a `serverindex.xml_disabled` file in this directory, perform the following actions:
 - a) Delete or rename the `serverindex.xml` file.
 - b) Rename the `serverindex.xml_disabled` file to `serverindex.xml`.
3. Perform one of the following actions for each of the deployment cell's managed nodes that you need to roll back.
 - a. If you backed up your previous managed node configuration using the `backupConfig` command or your own preferred backup utility, run the `restoreConfig` command or your own preferred utility to restore the version 6.0.x configuration for the managed node.

Important: Make sure that you restore the same backed-up configuration that you created just before you migrated the managed node.

See the `restoreConfig` command in the WebSphere Application Server Network Deployment, version 6.1 information center.

- b. If you did not back up your previous managed node configuration, use the `wsadmin` command to run the `migrationDisablementReversal.jacl` script from the version 6.0.x `profile_root/bin` directory of the managed node.

Linux In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migrationDisablementReversal.jacl -conntype NONE
```

Tip: If you have trouble running the `migrationDisablementReversal.jacl` script, try to manually perform the steps in the script.

- 1) Go to the following directory:
`profile_root/config/cells/cell_name/nodes/node_name`

where `node_name` is the name of the managed node that you want to roll back.
- 2) If you see a `serverindex.xml_disabled` file in in this directory, perform the following actions:
 - a) Delete or rename the `serverindex.xml` file.

- b) Rename the `serverindex.xml_disabled` file to `serverindex.xml`.
- 4. Synchronize the managed nodes if they were ever running when the version 6.1 deployment manager was running.
See Synchronizing nodes with the `wsadmin` tool in the WebSphere Application Server Network Deployment, version 6.1 information center.
- 5. If you chose to keep the installed applications in the same location as the prior release during migration to version 6.1 and any of the version 6.1 applications are not compatible with the prior release, install applications that are compatible.
- 6. Delete the version 6.1 profiles.
See Deleting a profile in the WebSphere Application Server Network Deployment, version 6.1 Information Center.
- 7. Roll back your databases. (For any databases that support WebSphere Process Server components that were upgraded, either automatically with the migration tools or manually, restore the backups that you made before you started the migration process.)
- 8. Start the rolled-back deployment manager and its managed nodes in the version 6.0.x environment.

Results

The configuration should now be returned to the state that it was in before migration.

What to do next

You can now restart the migration process if you want to do so.

Related tasks

“Rolling back a managed node”

You can use the `restoreConfig` and `wsadmin` commands to roll back a migrated WebSphere Process Server version 6.1 managed node to the state that it was in before migration. For each managed node that you want to roll back, you must roll back the managed node itself and the corresponding changes made to the master repository located on the deployment manager.

“Migrating a deployment manager using the migration wizard” on page 41
Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the migration wizard.



Migrating a deployment manager using command-line tools

Migrate a deployment manager from an older version to a newer version of WebSphere Process Server using the command-line tools.

Related information

`restoreConfig` command

`backupConfig` command

Synchronizing nodes with the `wsadmin` tool

Deleting a profile

Rolling back a managed node

You can use the `restoreConfig` and `wsadmin` commands to roll back a migrated WebSphere Process Server version 6.1 managed node to the state that it was in before migration. For each managed node that you want to roll back, you must roll

back the managed node itself and the corresponding changes made to the master repository located on the deployment manager.

Before you begin

When you migrate a version 6.0.x managed node, you must complete the following if you want to be able to roll it back to its previous state after migration:

1. Back up the databases that support WebSphere Process Server components.
2. Back up your existing configuration using the **backupConfig** command or your own preferred backup utility.
 - Run the **backupConfig** command or your own preferred utility to back up the version 6.0.x deployment manager configuration..

Important: Make sure that you note the exact name and location of this backed-up configuration.

See the backupConfig command in the WebSphere Application Server Network Deployment, version 6.1 information center.

- Run the **backupConfig** command or your own preferred utility to back up the version 6.0.x managed node configuration.

Important: Make sure that you note the exact name and location of this backed-up configuration.

See the backupConfig command in the WebSphere Application Server Network Deployment, version 6.1 information center.

3. Migrate the managed node.

If necessary, you can now roll back the managed node that you just migrated.

Important: If you do not have a backup copy of your version 6.1 deployment manager configuration as it was before you migrated the version 6.0.x managed node that you want to roll back, you cannot use the procedure described in this article and you must roll back your whole cell as described in “Rolling back a deployment cell” on page 89.

About this task

You must perform all of the backup and rollback actions for each migrated managed node before you proceed to roll back another managed node.

Procedure

1. Roll back your databases. (For any databases that support WebSphere Process Server components that were upgraded, either automatically with the migration tools or manually, restore the backups that you made before you started the migration process.)
2. Stop all of the servers that are currently running in the version 6.1 environment.
3. Restore your previous configuration.
 - a. Run the **restoreConfig** command or your own preferred utility to restore the version 6.1 deployment manager configuration.

Important: Make sure that you restore the same backed-up configuration that you created just before you migrated the managed node.

See the `restoreConfig` command in the WebSphere Application Server Network Deployment, version 6.1 Information Center.

- b. Perform one of the following actions to restore the version 6.0.x configuration for the managed node.
 - Run the **restoreConfig** command or your own preferred utility to restore the version 6.0.x configuration.
See `restoreConfig` command in the WebSphere Application Server Network Deployment, version 6.1 Information Center.
 - Use the **wsadmin** command to run the `migrationDisablementReversal.jacl` script from the `version 6.0.x/profile_root/bin` directory of the managed node.

Linux In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migrationDisablementReversal.jacl -conntype NONE
```

Tip: If you have trouble running the `migrationDisablementReversal.jacl` script, try to manually perform the steps in the script.

- 1) Go to the following directory:
`profile_root/config/cells/cell_name/nodes/node_name`

where `node_name` is the name of the managed node that you want to roll back.

- 2) If you see a `serverindex.xml_disabled` file in in this directory, perform the following actions:
 - a) Delete or rename the `serverindex.xml` file.
 - b) Rename the `serverindex.xml_disabled` file to `serverindex.xml`.
4. Start the version 6.1 deployment manager.
5. Synchronize the managed node.
See Synchronizing nodes with the `wsadmin` tool in the WebSphere Application Server Network Deployment, version 6.1 information center.
6. If you chose to keep the installed applications in the same location as the prior release during migration to version 6.1 and any of the version 6.1 applications are not compatible with the prior release, install applications that are compatible.
7. Delete the version 6.1 managed profile.
See Deleting a profile in the WebSphere Application Server Network Deployment, version 6.1 Information Center.
8. Start the rolled-back managed node in the version 6.1 environment.

Results

The configuration should now be returned to the state that it was in before migration.

What to do next

You can now restart the migration process if you want to do so.

Related tasks

“Rolling back a deployment cell” on page 89

You can use the **restoreConfig** and **wsadmin** commands to roll back a migrated WebSphere Process Server version 6.1 deployment cell to version 6.0.x. This

returns the configuration to the state that it was in before migration. After rolling back the deployment cell, you can restart the migration process.

“Migrating non-clustered managed nodes using the migration wizard” on page 49

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server using the migration wizard.

 Migrating a managed node using command-line tools

Migrate non-clustered managed nodes from an older version to a newer version of WebSphere Process Server with the command-line tools.

Related information

restoreConfig command

backupConfig command

Synchronizing nodes with the wsadmin tool

Deleting a profile

Migrating Cloudscape databases

After you use the migration tools to migrate to WebSphere Process Server version 6.1, you should verify the results of the automatic Cloudscape database migration and manually migrate any Cloudscape database instances that are not automatically migrated by the tools.

Before you begin

See “Overview of migrating” on page 1 and “Premigration considerations” on page 2.

Tips:

- Before you migrate a Cloudscape database, ensure that any servers hosting applications that are using the Cloudscape database are shut down. Otherwise, the Cloudscape migration will fail.
- Before you run the migration tools, ensure that the *debug migration trace* is active. By default, this trace function is enabled. To reactivate the debug migration trace if it is disabled, set one of the following trace options:
 - all traces*=all
 - com.ibm.ws.migration.WASUpgrade=all

About this task

WebSphere Process Server version 6.1 requires Cloudscape Version 10.1.

Cloudscape Version 10.1 is a pure Java database server that combines the Apache Derby runtime with the opportunity to use the full services of IBM Software Support. For comprehensive information about Cloudscape Version 10.1, see the Cloudscape product web pages.

For help in troubleshooting problems when migrating, see “Troubleshooting version-to-version migration” on page 107.

Procedure

1. Verify the automatic migration of Cloudscape database instances.

When you migrate from WebSphere Process Server version 6.0.x to version 6.1, the migration tools automatically upgrade the database instances that are accessed through the embedded framework by some internal components such as the UDDI registry. The tools also attempt to upgrade Cloudscape instances that your applications access through the embedded framework. You must verify these migration results after running the migration tools.

See “Verifying the Cloudscape v10.1.x automatic migration” on page 96.

2. Manually migrate Cloudscape database instances where necessary.

The version 6.1 migration tools do not attempt to automatically migrate database instances that transact with applications through the Cloudscape Network Server framework. This exclusion eliminates the risk of corrupting third-party applications that access the same database instances as those accessed by WebSphere Process Server

For details on manually migrating database instances that are accessed through the Cloudscape Network Server framework as well as Cloudscape instances that fail the automatic migration, see “Upgrading Cloudscape manually” on page 99.

3. Manually migrate your UDDI registry if it uses a database on the Cloudscape Network Server framework.

See “Migrating the UDDI registry” on page 103.

Related concepts

“How data is handled during migration from earlier versions” on page 19
The WebSphere Process Server version-to-version migration tools will handle different sets of data (enterprise application data, configuration data, and system application data) in different ways.

“Premigration considerations” on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.

“Overview of migrating” on page 1

Migrate from earlier versions of WebSphere Process Server and WebSphere Enterprise Service Bus.

Related tasks

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

“Verifying the Cloudscape v10.1.x automatic migration” on page 96

WebSphere Process Server version 6.1.x requires Cloudscape to run at a minimal version of v10.1.x. (Note that Cloudscape v10.1.x is comprised of the code base from Apache Derby Version 10.1.) During the WebSphere Process Server version 6.1 upgrade, the migration tool automatically upgrades the database instances that are accessed through the embedded framework by some internal components, such as the UDDI registry. The tool also attempts to upgrade Cloudscape instances that your applications access through the embedded framework. You must verify the migration results for these backend databases.

“Migrating the UDDI registry” on page 103

With most scenarios, migration of existing UDDI registries happens automatically when you migrate to the current level of WebSphere Process Server. However, if your existing UDDI registry uses a network Cloudscape database or a DB2 UDDI Version 2 database, there are some manual steps that you must take.

Related information

IBM Cloudscape product Web pages
Cloudscape migration document

Verifying the Cloudscape v10.1.x automatic migration

WebSphere Process Server version 6.1.x requires Cloudscape to run at a minimal version of v10.1.x. (Note that Cloudscape v10.1.x is comprised of the code base from Apache Derby Version 10.1.) During the WebSphere Process Server version 6.1 upgrade, the migration tool automatically upgrades the database instances that are accessed through the embedded framework by some internal components, such as the UDDI registry. The tool also attempts to upgrade Cloudscape instances that your applications access through the embedded framework. You must verify the migration results for these backend databases.

Before you begin

Do not use Cloudscape v10.1.x as a production database. Use it for development and test purposes only.

Learn more: The new version of Cloudscape combines the Derby runtime with additional benefits, such as IBM Quality Assurance (QA) and national language support (NLS). For information about the Cloudscape v10.1.x open source code base, see the Cloudscape product Web pages.

The migration tool attempts to upgrade Cloudscape database instances that are accessed through the embedded framework only. You must manually upgrade Cloudscape instances that transact with servers on the Derby Network Server framework. (See “Upgrading Cloudscape manually” on page 99.) This requirement eliminates the risk of corrupting third party applications that use the Network Server framework to access the same database instances as WebSphere Process Server

Other applications can access Cloudscape on Network Server because the framework provides the database with a foundation of connectivity software; the embedded framework does not. Cloudscape Network Server can transact with multiple Java Virtual Machines (JVM)s (or servers) concurrently, whereas Cloudscape on the embedded framework works with only a single JVM. Clustered or coexistence implementations of WebSphere Process Server require Network Server. For more information, consult the IBM Cloudscape information center.

About this task

For database instances that your applications access through the embedded framework, the automatic migration can succeed completely, fail completely, or succeed with warnings. A migration that produces warning messages does create a Cloudscape v10.1.x database with your data, but does not migrate all of your configured logic and other settings, such as:

- keys
- checks
- views
- triggers
- aliases
- stored procedures

To distinguish between a partially and a completely successful migration, you must verify the auto-migration results by checking both the general post-upgrade log and the individual database logs. Performing these tasks gives you vital diagnostic data to troubleshoot the partially migrated databases as well as those that fail auto-migration completely. Ultimately, you migrate these databases through a manual process.

Procedure

1. Open the post-upgrade log of each new WebSphere Process Server version 6.1.x profile. The path name of the log is *install_root*profiles/*profileName*/logs/WASPostUpgrade.*timestamp*.log.
2. Examine the post-upgrade log for database error messages. These exceptions indicate database migration failures. The following lines are an example of post-upgrade log content, in which the database error code is DSRA7600E. (The migration tool references all database exceptions with the prefix DSRA.)

```
MIGR0344I: Processing configuration file /opt/WebSphere51/AppServer/cloudscape/db2j.properties.
```

```
MIGR0344I: Processing configuration file /opt/WebSphere51/AppServer/config/cells/migr06/applications/MyBankApp.ear/deployments/MyBankApp/deployment.xml.
```

```
DSRA7600E: Cloudscape migration of database instance /opt/WebSphere61/Express/profiles/default/databases/_opt_WebSphere51_AppServer_bin_DefaultDB failed, reason: java.sql.SQLException: Failure creating target db
```

```
MIGR0430W: Cloudscape Database /fvt/temp/51BaseXExpress/PostUpgrade50BaseFVTTest9/testRun/pre/websphere_backup/bin/DefaultDB failed to migrate <new database name>
```

Important: Call IBM WebSphere Process Server Support if you see a migration failure message for a Cloudscape instance that is accessed by a WebSphere internal component (that is, a component that helps comprise WebSphere Process Server rather than one of your applications).

3. Open the individual database migration log that corresponds with each of your backend Cloudscape databases. These logs have the same timestamp as that of the general post-upgrade log. The logs display more detail about errors that are listed in the general post-upgrade log, as well as expose errors that are not documented by the general log.

The path name of each database log is WAS_HOME/profiles/*profileName*/logs/*myFullDbPathName_migrationLogtimestamp*.log.

4. Examine each database migration log for errors. For a completely successful migration, the log displays a message that is similar to the following text:

```
MIGR0429I: Cloudscape Database F:\temp\51BaseXExpress\PostUpgrade50BaseFVTTest2\testRun\pre\websphere_backup\bin\DefaultDB was successfully migrated. See log C:\WebSphere61\Express\profiles\default\logs\DefaultDB_migrationLogSun-Dec-18-13.31.40-CST-2005.log
```

Otherwise, the log displays error messages in the format of the following example:

```
connecting to source db <jdbc:db2j:/fvt/temp/51BaseXExpress/PostUpgrade50BaseFVTTest9/testRun/pre/websphere_backup/bin/DefaultDB>
```

```
connecting to source db <jdbc:db2j:/fvt/temp/51BaseXExpress/PostUpgrade50BaseFVTTest9/testRun/pre/websphere_backup/bin/DefaultDB> took 0.26 seconds
```

```
creating target db <jdbc:derby:/opt/WebSphere61/Express/profiles/default/databases/_opt_WebSphere51_AppServer_bin_DefaultDB>
```

```
ERROR: An error occurred during migration. See debug.log for more details.
```

```
shutting down databases
```

```
shutting down databases took 0.055 seconds
```

- For more data about a migration error, consult the debug log that corresponds with the database migration log. The WebSphere Application Server migration utility triggers a *debug migration trace* by default; this trace function generates the database debug logs. The full path name of a debug log is `WAS_HOME/profiles/profileName/logs/myFullDbPathName_migrationDebugTimestamp.log`.

The following lines are a sample of debug text. The lines display detailed exception data for the error that is referenced in the previous sample of database migration log data.

```
java.sql.SQLException: Database_opt_WebSphere51_AppServer_bin_DefaultDB already exists.
Aborting migration
  at com.ibm.db2j.tools.migration.MigrateFrom51Impl.go(Unknown Source)
  at com.ibm.db2j.tools.migration.MigrateFrom51Impl.doMigrate(Unknown Source)
  at com.ibm.db2j.tools.MigrateFrom51.doMigrate(Unknown Source)
  at com.ibm.ws.adapter.migration.CloudscapeMigrationUtility.migr
```

Results

- The WebSphere Process Server migration utility changes your Cloudscape JDBC configurations whether or not it successfully migrates the database instances that are accessed by your applications. The tool changes Cloudscape JDBC provider class paths, data source implementation classes, and data source helper classes. The following table depicts these changes:

Table 1. New class information

Class type	Old value	New value
JDBC provider class path	<code>\${CLOUDSCAPE_JDBC_DRIVER_PATH}/db2j.jar</code>	<code>\${DERBY_JDBC_DRIVER_PATH}/derby.jar</code> <ul style="list-style-type: none"> Where <code>DERBY_JDBC_DRIVER_PATH</code> is the WebSphere environment variable that defines your Cloudscape JDBC provider Where <code>derby.jar</code> is the base name of the JDBC driver class file (In your environment, reference the JDBC driver class file by the full path name.)
Data source implementation class: Connection pool	<code>com.ibm.db2j.jdbc.DB2JConnectionPool DataSource</code>	<code>org.apache.derby.jdbc.EmbeddedConnection PoolDataSource</code>
Data source implementation class: XA	<code>com.ibm.db2j.jdbc.DB2JXADataSource</code>	<code>org.apache.derby.jdbc.EmbeddedXADataSource</code>
Data source helper class	<code>com.ibm.websphere.rsadapter.Cloudscape DataStoreHelper</code>	<code>com.ibm.websphere.rsadapter.Derby DataStoreHelper</code>

Additionally, the `db2j.properties` file changes:

- The name `WAS_HOME/cloudscape/dbj.properties` changes to `WAS_HOME/derby/derby.properties`
- Within the file, property names change from `db2j.drda.*` to `derby.drda.*`
- A partial or a completely successful database migration changes the location and name of the database according to the following example:
 - Old database name:** `c:\temp\mydb`
 - New database name:** The new name includes a hash code that combines the entire path name of the old database and the migration time stamp. The new name also includes the old database name and time stamp verbatim. Example: `install_root\profiles\profile_name\databases\my_db_hashCode_timestamp`

Note the exact path names: For both partial and failed migrations, the log messages contain the exact old and new database path names that you must use to run the manual migration. Note these new path names precisely.

If you experience a partial migration, attempt to troubleshoot the new v10.1.x database only if you have expert knowledge of Cloudscape. Otherwise, delete the new database. Perform the manual migration procedure on the original database, just as you do for each database that completely fails auto-migration. Consult “Upgrading Cloudscape manually” on page 99 for instructions.

For successfully migrated Cloudscape instances, be aware that new cell-scoped data sources can only be used by nodes that run version 6.0.2 or later of WebSphere Process Server. Earlier versions of the product do not support the new Cloudscape; when applications on pre-version 6.0.2 nodes try to access a Cloudscape 10.1.x data source, the server will issue exceptions at run time.

Related tasks

“Upgrading Cloudscape manually”

During the WebSphere Process Server version 6.1 upgrade, the migration tools attempt to upgrade instances of Cloudscape that are accessed through the embedded framework only. (The new version of Cloudscape is version 10.1.x, which is based on Derby.) The automatic upgrade excludes Cloudscape instances that transact with applications through the Network Server framework. This exclusion eliminates the risk of corrupting third party applications that access the same database instances as WebSphere Process Server. You must manually upgrade database instances that are accessed through the Network Server framework. Do the same for databases that fail the automatic migration.

“Migrating Cloudscape databases” on page 94

After you use the migration tools to migrate to WebSphere Process Server version 6.1, you should verify the results of the automatic Cloudscape database migration and manually migrate any Cloudscape database instances that are not automatically migrated by the tools.

“Migrating the UDDI registry” on page 103

With most scenarios, migration of existing UDDI registries happens automatically when you migrate to the current level of WebSphere Process Server. However, if your existing UDDI registry uses a network Cloudscape database or a DB2 UDDI Version 2 database, there are some manual steps that you must take.

Related information

IBM Cloudscape product Web pages

Cloudscape migration document

Apache Derby

IBM Cloudscape information center

Upgrading Cloudscape manually

During the WebSphere Process Server version 6.1 upgrade, the migration tools attempt to upgrade instances of Cloudscape that are accessed through the embedded framework only. (The new version of Cloudscape is version 10.1.x, which is based on Derby.) The automatic upgrade excludes Cloudscape instances that transact with applications through the Network Server framework. This exclusion eliminates the risk of corrupting third party applications that access the same database instances as WebSphere Process Server. You must manually upgrade database instances that are accessed through the Network Server framework. Do the same for databases that fail the automatic migration.

Before you begin

Do not use Cloudscape v10.1.x as a production database. Use it for development and test purposes only.

Learn more: The new version of Cloudscape combines the Derby runtime with additional benefits, such as IBM Quality Assurance (QA) and national language support (NLS).

- For information about the Cloudscape v10.1.x open source code base, see the Cloudscape product Web pages.
- For information about incompatibilities between Cloudscape v10.1.x and v5.1.60x (plus versions prior to v5.1.60x) see Migrating IBM Cloudscape to Version 10.

For instances of Cloudscape that are accessed through the embedded framework, determine which instances completely failed the automatic upgrade process and which ones were only partially upgraded. The topic “Verifying the Cloudscape v10.1.x automatic migration” on page 96 documents how to uncover database errors and diagnostic data from various migration logs. The log messages contain the exact old and new database path names that you must use to run the manual migration. Note these new path names precisely.

To minimize the risk of migration errors for databases that were only partially upgraded during the automatic migration process, delete the new database. Troubleshoot the original database according to the log diagnostic data, then perform manual migration on the original database.

About this task

The following section consists of steps to migrate Cloudscape instances that are accessed through both frameworks: the embedded as well as the Network Server framework. Steps that apply only to the Cloudscape Network Server framework are marked accordingly. As a migration best practice, ensure that your user ID has one of the following authorities:

- Administrator of the server that accesses the Cloudscape instance
- A umask that can access the database instance

Otherwise, you might see runtime errors about the database instance being read-only.

Procedure

1. **Network Server framework only:** Ensure that every client of the Cloudscape database can support Cloudscape v10.1.x. WebSphere Process Server clients of the database must run versions 6.0.1.x or higher of WebSphere Process Server. In the case of mixed-node cells, remember that only nodes of WebSphere Process Server version 6.0.1.x or later can use data sources that you create postmigration for access to Cloudscape 10.1.x. Earlier versions of the product do not support the new Cloudscape; when applications on WebSphere Process Server pre-version 6.0.1.x nodes try to access a cell-scoped Cloudscape 10.1.x data source, WebSphere Process Server issues run-time exceptions.
2. **Network Server framework only:** Take the database offline. No clients can access it during the migration process.
3. Examine a sample Cloudscape migration script that WebSphere Process Server provides: either `db2jmmigrate.bat`, or `db2jmmigrate.sh`. The path of both scripts is `install_root\derby\bin\embedded\...`. You can modify the script according to the requirements of your environment. Consult Migrating IBM Cloudscape to Version 10 for information about options that you can use with the script. For example, you can use the option `-DB2j.migrate.ddlFile=filename` to specify the DDL file for the new database.
4. To generate database debug logs when you run the migration script, ensure that the *debug migration trace* is active. By default, this trace function is enabled. Reactivate the debug trace if it is disabled.

- a. To set the trace options in the administrative console, click **Troubleshooting > Logging and Tracing** in the console navigation tree.
 - b. Select the server name.
 - c. Click **Change Log Level Details**.
 - d. Optional: If **All Components** has been enabled, you might want to turn it off, and then enable specific components.
 - e. Optional: Select a component or group name. For more information see Log level settings in the WebSphere Application Server Network Deployment, version 6.1 Information Center. If the selected server is not running, you will not be able to see individual component in graphic mode.
 - f. Enter a trace string in the trace string box. In this case, enter one of the following:
 - all traces*=all
 - com.ibm.ws.migration.WASUpgrade=all

For more information on tracing read Working with trace in the WebSphere Application Server Network Deployment, version 6.1 Information Center..
 - g. Select **Apply**, then **OK**.
5. Specify your old database name and the full postmigration path of the new database name when you run the script. For example: E:\WebSphere\ProcServer\derby\bin\embedded>db2jMigrate.bat *myOldDB myNewDB* The logs from the automatic migration provide the exact path names to specify for both the old database and the target database. You must use this target database name to specify the new database, because your migrated Cloudscape data sources (updated by the WebSphere Process Server migration utilities) now point to the target database name. The following sample text demonstrates how log messages display target database names:

```
Cloudscape migration of database instance C:\temp\migration2\profiles\Srv01\
installedApps\ghongellNode01Cell\DynamicQuery.ear\EmployeeFinderDB to
new database instance C:\WebSphere\ProcServer
\profiles\Srv01\databases\C__WAS602_ProcServer_profiles_ProcSrv01_
installedApps_ghongellNode01Cell_DynamicQuery.ear_
EmployeeFinderDB failed, reason: java.sql.SQLException:
Failure creating target db
```

For instances of Cloudscape that are accessed through the Network Server framework, input any name that you want for the new database. Remember to modify your existing data sources to point to the new database name.

6. When the migration process ends, examine the database migration log to verify the results. The path name of each database migration log is *install_root/logs/derby/myFulldbName_migrationLog.log*.

For a successful migration, the database migration log displays a message that is similar to the following text:

```
Check E:\WebSphere\ProcServer\derby\myOldDB_migrationLog.log for progress
Migration Completed Successfully
E:\WebSphere\ProcServer\derby\bin\embedded>
```

Otherwise, the log displays error messages in the format of the following example:

```
Check E:\WebSphere\ProcServer\derby\myOldDB_migrationLog.log for progress
ERROR: An error occurred during migration. See debug.log for more details.
ERROR XMG02: Failure creating target db
java.sql.SQLException: Failure creating target db
    at com.ibm.db2j.tools.migration.MigrationState.getCurrSQLException(Unknown
    Source)
    at com.ibm.db2j.tools.migration.MigrateFrom51Impl.handleException(Unknown
    Source)
```

```
at com.ibm.db2j.tools.migration.MigrateFrom51Impl.go(Unknown Source)
at com.ibm.db2j.tools.migration.MigrateFrom51Impl.main(Unknown Source)
at com.ibm.db2j.tools.MigrateFrom51.main(Unknown Source)
```

7. For more data about a migration error, consult the debug log that corresponds with the database migration log. The full path name of a debug log file is *install_root/logs/derby/myFulldbName_migrationDebug.log*

The following lines are a sample of debug text.

```
sourceDBURL=jdbc:db2j:E:\WebSphere\my01dDB
newDBURL=jdbc:derby:e:\tempo\myNewDB
ddlOnly=false
connecting to source db <jdbc:db2j:E:\WebSphere\my01dDB>
connecting to source db <jdbc:db2j:E:\WebSphere\my01dDB> took 0.611 seconds
creating target db <jdbc:derby:e:\tempo\myNewDB>
creating target db <jdbc:derby:e:\tempo\myNewDB> took 6.589 seconds
initializing source db data structures
initializing source db data structures took 0.151 seconds
recording DDL to create db <E:\WebSphere\my01dDB>
recording DDL to create db <E:\WebSphere\my01dDB> took 5.808 seconds
```

Results

As indicated in the previous steps, the database migration log displays either a Migration Completed Successfully message, or a message containing migration failure exceptions.

What to do next

- For databases that fail migration, troubleshoot according to the logged error data. Then rerun the migration script.
 - To access successfully upgraded databases through the embedded framework, modify your data sources to point to the new database names.
 - To access successfully upgraded databases through the Network Server framework, you can use either the DB2 Universal JDBC driver or the Derby Client JDBC driver.
 - If you want your existing JDBC configurations to continue to use the DB2 Universal JDBC driver, modify your data sources to point to the new database names.
 - If you want to use the Derby Client JDBC driver, which can support XA data sources, modify your JDBC providers to use the new Derby Client JDBC driver class and the new data source implementation classes. Then reconfigure every existing data source to use the correct Derby data source helper class, and to point to the new database name.
- Consult the topic Data source minimum required settings, by vendor in the WebSphere Application Server Network Deployment, version 6.1 Information Center for all of the new class names.
- Run the database upgrade scripts in the *install_root/dbscripts/component_name/Derby* directory to upgrade the database tables and schema to the WebSphere Process Server version 6.1 level. For more information see “Upgrading databases for migration” on page 27.

Related tasks

“Verifying the Cloudscape v10.1.x automatic migration” on page 96
WebSphere Process Server version 6.1.x requires Cloudscape to run at a minimal version of v10.1.x. (Note that Cloudscape v10.1.x is comprised of the code base from Apache Derby Version 10.1.) During the WebSphere Process Server version 6.1 upgrade, the migration tool automatically upgrades the database instances that are accessed through the embedded framework by some

internal components, such as the UDDI registry. The tool also attempts to upgrade Cloudscape instances that your applications access through the embedded framework. You must verify the migration results for these backend databases.

“Migrating the UDDI registry”

With most scenarios, migration of existing UDDI registries happens automatically when you migrate to the current level of WebSphere Process Server. However, if your existing UDDI registry uses a network Cloudscape database or a DB2 UDDI Version 2 database, there are some manual steps that you must take.

“Upgrading databases for migration” on page 27

In conjunction with migration, the database schema of some WebSphere Process Server components must be upgraded. This can occur automatically but in some cases you must upgrade the schema manually.

Related information

IBM Cloudscape product Web pages

Cloudscape migration document

Migrating IBM Cloudscape to Version 10

Log level settings

Working with trace

Data source minimum required settings, by vendor

Migrating the UDDI registry

With most scenarios, migration of existing UDDI registries happens automatically when you migrate to the current level of WebSphere Process Server. However, if your existing UDDI registry uses a network Cloudscape database or a DB2 UDDI Version 2 database, there are some manual steps that you must take.

Before you begin

Migrate your installation of WebSphere Process Server; ensure that you select the option to migrate applications, so that the UDDI registry application will be migrated.

About this task

If your existing UDDI registry uses an Oracle, embedded Cloudscape or DB2 UDDI Version 3 database, you do not need to perform any manual migration; migration happens automatically when you migrate WebSphere Process Server and start the UDDI node for the first time after migration.

If your existing UDDI registry uses a network Cloudscape database or a DB2 UDDI Version 2 database, you must complete some manual steps to migrate the registry.

- If your UDDI registry uses a DB2 UDDI Version 2 database, follow the steps in Migrating to Version 3 of the UDDI registry. and sub-topics.
- If your UDDI registry uses a network Cloudscape database, complete the following steps.
 1. If you have a cluster that contains servers at different levels of WebSphere Process Server, ensure that any UDDI registries are running on servers that are at WebSphere Process Server version 6.1. For example, if you have a cluster that spans two nodes, you can upgrade one node to WebSphere

Process Server version 6.1 while the other node remains at a previous level, provided that any servers that are running a UDDI registry are at version 6.1.

2. Initialize the relevant UDDI node. The initialize process will perform some of the UDDI registry migration.
3. Enter the following commands as the database administrator, from *install_root/cloudscape/lib*.

```
java -cp db2j.jar;db2jtools.jar com.ibm.db2j.tools.ij

connect 'jdbc:db2j:uddi_cloudscape_database_path';

run 'install_root/UDDIReg/databaseScripts/uddi30crt_drop_triggers
_cloudscape.sql';

quit;

cd install_root/derby/migration

java -cp db2j.jar;db2jmigration.jar;../lib/derby.jar
com.ibm.db2j.tools.MigrateFrom51
jdbc:db2j:uddi_cloudscape_database_path
```

where

- *uddi_cloudscape_database_path* is the absolute path of the existing Cloudscape database, for example *install_root/profiles/profile_name/databases/com.ibm.uddi/UDDI30*
- *install_root* is the root directory for the installation of WebSphere Process Server.

Results

The UDDI database and data source are migrated, and the UDDI node is activated.

Note: When you migrate WebSphere Process Server, the post-upgrade log for the profile indicates that the migration of the UDDI database is partially complete, and is missing the steps for triggers, aliases, and stored statements. If you initially enabled the debug function, the debug log for the database indicates that there was a failure creating triggers. Ignore these messages; the UDDI node completes the migration of the database when the UDDI node starts. For more information about these log files, see “Verifying the Cloudscape v10.1.x automatic migration” on page 96. Also refer to this topic if other errors appear in the logs.

If the migration of the UDDI database completes successfully, the following message appears in the server log:

```
CWUDQ0003I: UDDI registry migration has completed
```

If the following error appears, an unexpected error occurred during migration. The UDDI registry node is not activated. Check the error logs for the problem and, if it cannot be fixed, see the IBM software support Web site at <http://www.ibm.com/software/support>.

```
CWUDQ0004W: UDDI registry not started due to migration errors
```

Related tasks

“Verifying the Cloudscape v10.1.x automatic migration” on page 96
WebSphere Process Server version 6.1.x requires Cloudscape to run at a minimal version of v10.1.x. (Note that Cloudscape v10.1.x is comprised of the code base from Apache Derby Version 10.1.) During the WebSphere Process

Server version 6.1 upgrade, the migration tool automatically upgrades the database instances that are accessed through the embedded framework by some internal components, such as the UDDI registry. The tool also attempts to upgrade Cloudscape instances that your applications access through the embedded framework. You must verify the migration results for these backend databases.

“Upgrading Cloudscape manually” on page 99

During the WebSphere Process Server version 6.1 upgrade, the migration tools attempt to upgrade instances of Cloudscape that are accessed through the embedded framework only. (The new version of Cloudscape is version 10.1.x, which is based on Derby.) The automatic upgrade excludes Cloudscape instances that transact with applications through the Network Server framework. This exclusion eliminates the risk of corrupting third party applications that access the same database instances as WebSphere Process Server. You must manually upgrade database instances that are accessed through the Network Server framework. Do the same for databases that fail the automatic migration.

“Migrating Cloudscape databases” on page 94

After you use the migration tools to migrate to WebSphere Process Server version 6.1, you should verify the results of the automatic Cloudscape database migration and manually migrate any Cloudscape database instances that are not automatically migrated by the tools.

Related information

Migrating to Version 3 of the UDDI registry

IBM Software Support web site

Migration considerations for Business Process Choreographer

If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Mixed cell restrictions

If your cell is running nodes at both the version 6.0.x level and version 6.1 levels concurrently during the course of the migration process for the cell, be aware of the following:

- Once a deployment manager has been migrated to version 6.1, you no longer can install, update, or uninstall Business Process Choreographer applications (BPEL applications or human tasks) on nodes in the cell that are still at the version 6.0.x level.
- Once a deployment manager has been migrated to version 6.1, you no longer can configure Business Process Choreographer on nodes in the cell that are still at the version 6.0.x level.
- When you have a version 6.1 cluster with Business Process Choreographer configured on it, you must not create new cluster members on 6.0.x nodes in the same cell.

Postmigration tasks

You might need to perform these tasks, if they apply to your environment, before using your WebSphere Process Server version 6.1 in production.

- If your WebSphere Process Server version 6.0.1 used the Business Process Choreographer Observer sample, remove the sample. See Removing the Business

Process Choreographer Observer Sample Version 6.0.1. This sample is not migrated. Business Process Choreographer Observer for version 6.0.2 or version 6.1 is not a sample.

- If you have written a client that uses Business Process Choreographer APIs without first authenticating the user, you should modify the client to perform a login before using the APIs. After migration, the J2EE roles BPEAPIUser and TaskAPIUser are set to the value Everyone, which maintains backward compatibility by maintaining the 6.0.x behavior of not requiring a login when application security is enabled. After you have fixed your client, you must change these roles to the value AllAuthenticated to prevent unauthenticated users accessing the APIs. For new installations these roles default to the value AllAuthenticated.

To do this:

1. Open the administrative console and select **Applications > Enterprise Applications**.
 2. In the right panel, select BPEContainer_ *name* , where *name* is either *nodeName_serverName* or *clusterName*, depending on whether you configured Business Process Choreographer on a server or on a cluster. (Select the name, not the check box to the left of the name.)
 3. In the right panel, under Detail Properties, select **Security role to user/group mapping**.
 4. Change the mapping for the J2EE BPEAPIUser role from "Everyone" to "All authenticated".
 5. Select **OK**.
 6. Repeat these steps for the TaskAPIUser role of the TaskContainer_ *name* enterprise application.
 7. Save your changes, then restart the server or cluster on which you configured Business Process Choreographer.
- If you have applied any changes to the default XSL transformation files (EverybodyTransformation.xml, LDAPTransformation.xml, SystemTransformation.xml, and UserRegistryTransformation.xml) located in the *install_root/ProcessChoreographer/Staff* directory then you must re-apply your changes to the WebSphere Process Server version 6.1 versions of these files after migration. Custom XSL transformation files located in the *install_root/ProcessChoreographer/Staff* directory will be migrated automatically. Custom XSL transformation files located in other directories may have to be copied manually, depending on the exact value of the transformation file path specified in the version 6.0.x staff plug-in configuration (now known as the people directory configuration in WebSphere Process Server version 6.1).

Related concepts

"Premigration considerations" on page 2

Before you begin the process of migrating to a new version of WebSphere Process Server, you should be aware of these considerations.


Related tasks

"Troubleshooting version-to-version migration" on page 107

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

Related information

 [Removing the Business Process Choreographer Observer Sample Version 6.0.1](#)

 Removing the Business Process Choreographer Observer Version 6.0.1 sample

 Administering enterprise applications

Use the console's Enterprise Application page (viewed by clicking **Applications > Enterprise Applications**) to view and administer enterprise applications installed on the server.

Troubleshooting version-to-version migration

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

- While you are using the version 6.1 migration wizard to create a profile before migrating a configuration, you might see the following profile-creation error messages.

```
profileName: profileName cannot be empty  
profilePath: Insufficient disk space
```

These error messages might be displayed if you enter a profile name that contains an incorrect character such as a space. Rerun the migration wizard, and verify that there are no incorrect characters in the profile name such as a space, quotes, or any other special characters.

- If you encounter a problem when you are migrating from a previous version of WebSphere Process Server to version 6.1, check your log files and other available information.

1. Look for the log files, and browse them for clues.

- *migration_backup_directory*/WBIPreUpgrade.time_stamp.log
- *profile_root*/log/WASPostUpgrade.time_stamp.log
- *install_root*/logs/clientupgrade.time_stamp.log
- *profile_root*/logs/bpeupgrade.log
- *migration_backup_directory*/WBIProfileUpgrade.ant.timestamp.log

2. Look for MIGR0259I: The migration has successfully completed. or MIGR0271W: The migration completed with warnings. in the following directories:

- *migration_backup_directory*/WBIPreUpgrade.time_stamp.log
- *profile_root*/logs/WASPostUpgrade.time_stamp.log
- *install_root*/logs/clientupgrade.time_stamp.log

If MIGR0286E: The migration failed to complete. is displayed, attempt to correct any problems based on the error messages that appear in the log file. After correcting any errors, rerun the command from the bin directory of the product installation root.

3. Open the Log and Trace Analyzer built into the Application Server Toolkit (AST) on the service log of the server that is hosting the resource that you are trying to access, and use it to browse error and warning messages.

See Debugging components in the Application Server Toolkit.

4. With WebSphere Process Server, run the dumpNameSpace command and pipe, redirect, or "more" the output so that it can be easily viewed.

This command results in a display of all objects in WebSphere Process Server namespace, including the directory path and object name.

5. If the object a client needs to access does not appear, use the administrative console to verify the following conditions.

- The server hosting the target resource is started.
- The Web module or Enterprise JavaBean container hosting the target resource is running.
- The JNDI name of the target resource is properly specified.

If none of these steps solves the problem, see Troubleshooting and support for additional troubleshooting resources, including information about how to contact IBM Support.

- During the migration process, problems might occur while you are using the WBIPreUpgrade command or the WBIPostUpgrade command.

- Problems can occur when you are using the WBIPreUpgrade command.
 - A "Not found" or "No such file or directory" message is returned.

This problem can occur if you are trying to run the WBIPreUpgrade command from a directory other than the WebSphere Process Server version 6.1 *install_root*/bin directory. Verify that the WBIPreUpgrade script resides in the version 6.1 *install_root*/bin directory, and launch the file from that location.

- The DB2 JDBC driver and DB2 JDBC driver (XA) cannot be found in the drop-down list of supported JDBC providers in the administrative console. The administrative console no longer displays deprecated JDBC provider names. The new JDBC provider names used in the administrative console are more descriptive and less confusing. The new providers will differ only by name from the deprecated ones.

The deprecated names will continue to exist in the *jdbc-resource-provider-templates.xml* file for migration reasons (for existing JAAS scripts for example); however, you are encouraged to use the new JDBC provider names in your JAAS scripts.

- You receive the following message:

```
MIGR0108E: The specified WebSphere directory does not contain a WebSphere version that can be upgraded.
```

This can occur if an incorrect directory was specified with the WBIPreUpgrade command.

See the WBIPreUpgrade command.

- Problems can occur when you are using the WBIPostUpgrade command.
 - A "Not found" or "No such file or directory" message is returned.

This problem can occur if you are trying to run the WBIPostUpgrade command from a directory other than the WebSphere Process Server version 6.1 *install_root*\bin. Verify that the WBIPostUpgrade script resides in the version 6.1 *install_root*\bin directory, and launch the file from that location.

- When you migrate the federated nodes in a cell, you receive the following error messages:

```
MIGR0304I: The previous WebSphere environment is being restored.
com.ibm.websphere.management.exception.RepositoryException:
com.ibm.websphere.management.exception.ConnectorException: ADMC0009E:
The system failed to make the SOAP RPC call: invoke
MIGR0286E: The migration failed to complete.
```

A connection timeout occurs when the federated node tries to retrieve configuration updates from the deployment manager during the WBIPostUpgrade migration step for the federated node. Copying the entire configuration might take more than the connection timeout if the configuration that you are migrating to version 6.1 contains any of the following elements:

- Many small applications
- A few large applications
- One very large application

If this occurs, modify the timeout value before running the `WBIPostUpgrade` command to migrate a federated node.

1. Go to the following location in the version 6.1 directory for the profile to which you are migrating your federated node:
profile_root/properties
2. Open the `soap.client.props` file in that directory and find the value for the `com.ibm.SOAP.requestTimeout` property. This is the timeout value in seconds. The default value is 180 seconds.
3. Change the value of `com.ibm.SOAP.requestTimeout` to make it large enough to migrate your configuration. For example, the following entry would give you a timeout value of a half of an hour:

```
com.ibm.SOAP.requestTimeout=1800
```

Note: Select the smallest timeout value that will meet your needs. Be prepared to wait for at least three times the timeout that you select—once to download files to the backup directory, once to upload the migrated files to the deployment manager, and once to synchronize the deployment manager with the migrated node agent.

4. Go to the following location in the backup directory that was created by the `WBIPreUpgrade` command:
migration_backup_directory/profiles/default/properties
5. Open the `soap.client.props` file in that directory and find the value for the `com.ibm.SOAP.requestTimeout` property.
6. Change the value of `com.ibm.SOAP.requestTimeout` to the same value that you used in the version 6.1 file.

- You receive the "Unable to copy document to temp file" error message. Here is an example:

```
MIGR0304I: The previous WebSphere environment is being restored.
com.ibm.websphere.management.exception.DocumentIOException: Unable to copy
document to temp file:
  cells/sunblade1Network/applications/LARGEApp.ear/LARGEApp.ear
```

Your file system might be full. If your file system is full, clear some space and rerun the `WBIPostUpgrade` command.

- You receive the following message:
MIGR0108E: The specified WebSphere directory does not contain a WebSphere version that can be upgraded.

The following possible reasons for this error exist:

- An incorrect directory might have been specified when launching the `WBIPreUpgrade` command or the `WBIPostUpgrade` .
- The `WBIPreUpgrade` command was not run.

- You receive the following error message:

```
MIGR0253E: The backup directory migration_backup_directory does not exist.
```

The following possible reasons for this error exist:

- The `WBIPreUpgrade` command was not run before the `WBIPostUpgrade` command.
 1. Check to see if the backup directory specified in the error message exists.
 2. If not, run the `WBIPreUpgrade` command.

See WBIPreUpgrade command.

3. Retry the WBIPostUpgrade command.

- An incorrect backup directory might be specified.

For example, the directory might have been a subdirectory of the version 6.0.x tree that was deleted after the WBIPreUpgrade command was run and the older version of the product was uninstalled but before the WBIPostUpgrade command was run.

1. Determine whether or not the full directory structure specified in the error message exists.

2. If possible, rerun the WBIPreUpgrade command, specifying the correct full migration backup directory.

3. If the backup directory does not exist and the older version it came from is gone, rebuild the older version from a backup repository or XML configuration file.

4. Rerun the WBIPreUpgrade command.

- You decide that you need to run WBIPreUpgrade again after you have already run the WBIPostUpgrade command.

During the course of a deployment manager or a managed node migration, WBIPostUpgrade might disable the old environment. If after running WBIPostUpgrade you want to run WBIPreUpgrade again against the old installation, you must run the migrationDisablementReversal.jacl script located in the old *install_root*/bin directory. After running this JACL script, your version 6.0.x environment will be in a valid state again, allowing you to run WBIPreUpgrade to produce valid results.

For more information on scripting, see Getting started with scripting. Scripting, as described there, is available for WebSphere Process Server.

- A federated migration fails with message MIGR0405E.

The migration that has taken place on your deployment manager as part of your federated migration has failed. For a more detailed reason for why this error has occurred, open the folder *your_node_name_migration_temp* located on your deployment manager node under the ...DeploymentManagerProfile/temp directory. For example:

```
/websphere61/procserver/profiles/dm_profile/temp/nodeX_migration_temp
```

The logs and everything else involved in the migration for this node on the deployment manager node are located in this folder. This folder will also be required for IBM support related to this scenario.

- WebSphere Process Server version 6.1 applications are lost during migration.

If any of the version 6.1 applications fail to install during a federated migration, they will be lost during the synchronizing of the configurations. The reason that this happens is that one of the final steps of WBIPostUpgrade is to run a syncNode command. This has the result of downloading the configuration on the deployment manager node and overwriting the configuration on the federated node. If the applications fail to install, they will not be in the configuration located on the deployment manager node. To resolve this issue, manually install the applications after migration. If they are standard version 6.1 applications, they will be located in the *install_root*/installableApps directory.

To manually install an application that was lost during migration, use the wsadmin command to run the *install_application_name.jacl* script that the migration tools created in the backup directory.

Linux

In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migration_backup_directory/install_application_name.jacl  
-conntype NONE
```

See the Wsadmin tool.

- WebSphere Process Server version 6.1 applications fail to install.

Manually install the applications using the wsadmin command after WBIPostUpgrade has completed.

To manually install an application that failed to install during migration, use the wsadmin command to run the `install_application_name.jacl` script that the migration tools created in the backup directory.

Linux

In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migration_backup_directory/install_application_name.jacl  
-conntype NONE
```

See the Wsadmin tool, or see the WBIPostUpgrade command.

- **Solaris** When you use the migration wizard to migrate a profile from WebSphere Process Server version 6.0.x to version 6.1.x on a Solaris x64 processor-based system, the migration might fail during the WBIPostUpgrade step.

You might see messages similar to the following in `profile_root/logs/WASPostUpgrade.time_stamp.log`:

```
MIGR0327E: A failure occurred with stopNode.  
MIGR0272E: The migration function cannot complete the command.
```

WebSphere Process Server version 6.0.x uses a Java virtual machine (JVM) in 32-bit mode. The migration wizard for WebSphere Process Server version 6.1.x calls the WBIPostUpgrade.sh script, which attempts to run the JVM for version 6.0.x in the 64-bit mode when the server stops the version 6.0.x node.

Complete the following actions to remove the incomplete profile and enable WebSphere Process Server to correctly migrate the version 6.0.x profile:

1. On a command line, change to the `install_root/bin` directory.

For example, type the following command:

```
cd /opt/IBM/WebSphere/Procserver/bin
```

2. Locate the WBIPostUpgrade.sh script in the `install_root/bin` directory, and make a backup copy.
3. Open the WBIPostUpgrade.sh or WBIPostUpgrade.bat file in an editor, and perform the following actions:
 - a. Locate the following line of code:

UNIX

Linux

```
"$binDir" /setupCmdLine.sh
```

Windows

```
call "%~dp0setupCmdLine.bat" %*
```

- b. Insert the following line of code after the code that was identified in the previous step:

```
JVM_EXTRA_CMD_ARGS=""
```
 - c. Save the changes.
4. Repeat steps 2 through 4 with the `WASPostUpgrade.sh` or the `WASPostUpgrade.bat` file.

5. Use the following command to delete the incomplete version 6.1.x profile that was created during the migration process:
`install_root/bin/manageprofiles.sh -delete -profileName profile_name`
6. Delete the *profile_root* directory of the version 6.1.x profile that was removed in the previous step.
7. Rerun the migration wizard.

- If you select the option for the migration process to install the enterprise applications that exist in the version 6.0.x configuration into the new version 6.1 configuration, you might encounter some error messages during the application-installation phase of migration.

The applications that exist in the version 6.0.x configuration might have incorrect deployment information—usually, incorrect XML documents that were not validated sufficiently in previous WebSphere Process Server runtimes. The runtime now has an improved application-installation validation process and will fail to install these malformed EAR files. This results in a failure during the application-installation phase of WBIPostUpgrade and produces an "E:" error message. This is considered a "fatal" migration error.

If migration fails in this way during application installation, you can do one of the following:

- Fix the problems in the version 6.0.x applications, and then remigrate.
- Proceed with the migration and ignore these errors.

In this case, the migration process does not install the failing applications but does complete all of the other migration steps.

Later, you can fix the problems in the applications and then manually install them in the new version 6.1 configuration using the administrative console or an install script.

- After migrating to a version 6.1 cell that contains or interoperates with version 6.0.x nodes that are not at WebSphere Process Server version 6.0.1.3 or later, the cluster function might fail.

When starting these version 6.0.x servers, you might see the following problems:

- You might see a first failure data capture (FFDC) log that shows a `ClassNotFoundException` error message. This exception is thrown from the `RuleEtiquette.runRules` method and looks something like the following example:

```
Exception = java.lang.ClassNotFoundException
Source = com.ibm.ws.cluster.selection.SelectionAdvisor.<init>
probeid = 133
Stack Dump = java.lang.ClassNotFoundException: rule.local.server
at java.net.URLClassLoader.findClass(URLClassLoader.java(Compiled Code))
at com.ibm.ws.bootstrap.ExtClassLoader.findClass(ExtClassLoader.java:106)
at java.lang.ClassLoader.loadClass(ClassLoader.java(Compiled Code))
at java.lang.ClassLoader.loadClass(ClassLoader.java(Compiled Code))
at java.lang.Class.forName1(Native Method)
at java.lang.Class.forName(Class.java(Compiled Code))
at com.ibm.ws.cluster.selection.rule.RuleEtiquette.runRules(RuleEtiquette.java:154)at com.ibm.ws.cluster.selection.SelectionAdvisor.handleNotification
(SelectionAdvisor.java:153)
at com.ibm.websphere.cluster.topography.DescriptionFactory$Notifier.run
(DescriptionFactory.java:257)
at com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1462)
```

- You might see a `java.io.IOException` that looks something like the following example:

```
Exception = java.io.IOException
Source = com.ibm.ws.cluster.topography.DescriptionManagerA. update probeid
= 362
```

```

Stack Dump = java.io.IOException
at com.ibm.ws.cluster.topography.ClusterDescriptionImpl.importFromStream
(ClusterDescriptionImpl.java:916)
at com.ibm.ws.cluster.topography.DescriptionManagerA.update
(DescriptionManagerA.java:360)
Caused by: java.io.EOFException
at java.io.DataInputStream.readFully(DataInputStream.java(Compiled Code))
at java.io.DataInputStream.readUTF(DataInputStream.java(Compiled Code))
at com.ibm.ws.cluster.topography.KeyRepositoryImpl.importFromStream
(KeyRepositoryImpl.java:193)

```

During migration, version 6.1 cluster information is distributed throughout the cell. WebSphere Process Server version 6.0.x nodes that are not at version 6.0.1.3 or later fail to read this information.

To avoid this problem, upgrade all version 6.0.x nodes that will be contained in or interoperating with a version 6.1 cell to version 6.0.1.3 or later before migrating your deployment managers to version 6.1.

- After you migrate a managed node to version 6.1, the application server might not start.

When you try to start the application server, you might see errors similar to those in the following example:

```

[5/11/06 15:41:23:190 CDT] 0000000a SystemErr R
com.ibm.ws.exception.RuntimeError:
com.ibm.ws.exception.RuntimeError: org.omg.CORBA.INTERNAL:
CREATE_LISTENER_FAILED_4
vmcid: 0x49421000 minor code: 56 completed: No
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServerImpl.bootServerContainer(WsServerImpl.java:198)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServerImpl.start(WsServerImpl.java:139)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServerImpl.main(WsServerImpl.java:460)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServer.main(WsServer.java:59)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:64)
[5/11/06 15:41:23:197 CDT] 0000000a SystemErr R at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)

```

Change the port number at which the managed node's server is listening. If the deployment manager is listening at port 9101 for ORB_LISTENER_ADDRESS, for example, the server of the managed node should not be listening at port 9101 for its ORB_LISTENER_ADDRESS. To resolve the problem in this example, perform the following steps:

1. On the administrative console, click **Application servers** → *server_name* → **Ports** → **ORB_LISTENER_ADDRESS**.
 2. Change the ORB_LISTENER_ADDRESS port number to one that is not used.
- If synchronization fails when you migrate a managed node to version 6.1, the server might not start.

You might receive messages similar to the following when you migrate a managed node to version 6.1:

```

ADMU0016I: Synchronizing configuration between node and cell.
ADMU0111E: Program exiting with error:
com.ibm.websphere.management.exception.AdminException: ADMU0005E:
Error synchronizing repositories
ADMU0211I: Error details may be seen in the file:
/opt/WebSphere/61AppServer/profiles/AppSrv02/logs/syncNode.log
MIGR0350W: Synchronization with the deployment manager using the SOAP protocol

```

failed.
MIGR0307I: The restoration of the previous WebSphere Application Server environment is complete.
MIGR0271W: Migration completed successfully, with one or more warnings.

These messages indicate the following:

- Your deployment manager is at a version 6.1 configuration level.
- The managed node that you are trying to migrate is at a version 6.1 configuration level on the deployment manager’s repository (including applications).
- The managed node itself is not quite complete given that you did not complete the syncNode operation.

Perform the following actions to resolve this issue:

1. Rerun the syncNode command on the node to synchronize it with the deployment manager.
See the syncNode command .
2. Run the GenPluginCfg command.
See the GenPluginCfg command .

What to do next

If you did not find your problem listed, contact IBM support.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Related tasks

“Verifying migration” on page 87
Verify that your migration was successful by checking the log files and checking operation with the administrative console.

Related reference

“WBIPreUpgrade command” on page 12
Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.



WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

Debugging components in the Application Server Toolkit

Wsadmin tool

syncNode command

GenPluginCfg command



Troubleshooting and support

To help you understand, isolate, and resolve problems with your IBM software, the troubleshooting and support information contains instructions for using the problem-determination resources that are provided with your IBM products.

Getting started with scripting

Chapter 2. Migrating from previous WebSphere products

You can migrate applications and configuration data from certain IBM products that existed before WebSphere Process Server.

Migration from another product to WebSphere Process Server is supported from the following predecessor products:

- WebSphere InterChange Server version 4.2.0 or later. For more information see “Migrating from WebSphere InterChange Server.”
- WebSphere Business Integration Server Foundation versions 5.1 and 5.1.1. For more information see “Migrating source artifacts from WebSphere Studio Application Developer Integration Edition” on page 167.
- WebSphere MQ Workflow version 3.6. For more information, see “Migrating from WebSphere MQ Workflow” on page 167.

Note: You can also migrate to WebSphere Process Server from certain versions of WebSphere Enterprise Service Bus and WebSphere Application Server, as well as from prior versions of WebSphere Process Server itself. For more information about migration from these products, see Chapter 1, “Migrating from previous versions of WebSphere Process Server and WebSphere Enterprise Service Bus,” on page 1.

i5/OS Although these predecessor products were not supported on i5/OS, modules from them can be migrated to WebSphere Process Server version 6.1 on their respective platforms (using the migration tools available such as the `reposMigrate` command) and then deployed to WebSphere Process Server version 6.1 running on the i5/OS operating system.

For migration from a previous product to WebSphere Process Server (for example, WebSphere InterChange Server to WebSphere Process Server), the migration steps include using migration tools to convert source artifacts to the new WebSphere Process Server version of the artifacts.

WebSphere Integration Developer contains migration tools that assist in migrating existing application source artifacts into WebSphere Process Server artifacts. These tools can be accessed through the **File > Import...** wizards of WebSphere Integration Developer. The migration tools designed to assist with migration from WebSphere InterChange Server can also be accessed through the command line of WebSphere Process Server.

You can also find articles that might help you with migration in the IBM developerWorks® “Technical Library” at <http://www.ibm.com/developerworks>.

Migrating from WebSphere InterChange Server

Use the WebSphere Integration Developer wizard or the WebSphere Process Server `reposMigrate` command to migrate from WebSphere InterChange Server 4.2.2 or later to WebSphere Process Server 6.1.

About this task

For this version of WebSphere InterChange Server...	Do this
WebSphere InterChange Server version 4.2.2 or later	Use the Migration wizard from WebSphere Integration Developer to migrate all WebSphere InterChange Server artifacts into WebSphere Process Server deployable artifacts and place them into projects in the active WebSphere Integration Developer workspace. Alternatively, you can use the <code>reposMigrate</code> command to migrate all WebSphere InterChange Server artifacts into WebSphere Process Server deployable artifacts, and optionally deploy them directly to WebSphere Process Server.
WebSphere InterChange Server versions earlier than 4.2.2	Migrate to WebSphere InterChange Server 4.2.2 or later first, and then migrate to WebSphere Process Server.

Related information

Migrating WebSphere InterChange Server using the Migration wizard
 WebSphere Integration Developer information center

Premigration considerations

Consider these guidelines for the development of integration artifacts for WebSphere InterChange Server to ease the migration of WebSphere InterChange Server artifacts to WebSphere Process Server.

These recommendations are meant to be used only as a guide. There may be cases where it is necessary to deviate from these guidelines. In these cases care should be taken to limit the scope of the deviation to minimize the amount of rework required to migrate the artifacts. Note that the guidelines outlined here are not all general recommendations for the development of WebSphere InterChange Server artifacts. They are instead limited in scope to those considerations which may affect the ease in which artifacts can be migrated at a future time.

Related concepts

“Troubleshooting migration from WebSphere InterChange Server” on page 165
 Find solutions to problems you encounter with migration as well as instructions for turning on logging and tracing.

Related reference

“Postmigration considerations” on page 128
 When applications have been migrated from WebSphere InterChange Server to WebSphere Process Server, special attention is required in some areas to enable migrated applications to function in WebSphere Process Server consistently with their intended function due to differences between the architectures of WebSphere Process Server and WebSphere InterChange Server.

Premigration considerations: General development

Follow these recommended practices when developing WebSphere InterChange Server modules to ease future migration to WebSphere Process Server.

Several considerations apply broadly to developing most of the integration artifacts. In general, artifacts that leverage the facilities provided by WebSphere InterChange Server tooling and conform to the metadata models enforced by the

tooling will migrate most smoothly. Also, artifacts with significant extensions and external dependencies are likely to require more manual intervention when migrating.

In general, IBM recommends that you do the following:

- Document the system and component design
- Use the development tooling to edit integration artifacts
- Leverage recommended practices for defining rules with the tooling and Java snippets

It is important for integration solutions to adhere to the programming model and architecture provided by WebSphere InterChange Server. Each of the integration components within WebSphere InterChange Server plays a well-defined role within the architecture. Significant deviations from this model will make it more challenging to migrate content to the appropriate artifacts on WebSphere Process Server.

Another general practice which will improve the success of future migration projects is to document the system design. Be sure to capture the integration architecture and design, including functional design and quality of service requirements, the interdependencies of artifacts shared across projects, and also the design decisions that were made during the deployment. This will assist in system analysis during migration, and will minimize any rework efforts.

For creating, configuring, and modifying artifact definitions, use only the development tooling provided. Avoid manual manipulation of artifact metadata (for example, editing XML files directly), which may corrupt the artifact for migration.

IBM suggests the following when you are developing Java code within collaboration templates, maps, common code utilities, and other components:

- Use only the published APIs.
- Use Activity Editor.
- Use adapters to access EISs.
- Avoid external dependencies in Java snippet code.
- Adhere to J2EE develop practices for portability.
- Do not spawn threads or use thread synchronization primitives. If you must, these will need to be converted to use Asynchronous Beans when you migrate.
- Do not do any disk I/O using java.io.* Use JDBC to store any data.
- Do not perform any functions that may be reserved for an EJB container such as socket I/O, classloading, loading native libraries, and so forth. If you must, these snippets would need manual conversion to use EJB container functions when you migrate.

Use only the APIs published in the WebSphere InterChange Server product documentation for the artifacts. These are outlined in detail in the WebSphere InterChange Server development guides. Compatibility APIs will be provided in WebSphere Process Server for published WebSphere InterChange Server APIs. Although WebSphere InterChange Server has many internal interfaces which you might wish to use, IBM discourages this practice because these interfaces are not guaranteed to be supported in the future.

When designing business logic and transformation rules in maps and collaboration templates, try to avoid field developed common code utility libraries, included as a Java archive (*.jar) file in the classpath of WebSphere InterChange Server, as these will need to be migrated manually.

Use the Activity Editor tool to the greatest extent possible. This will ensure that the logic is described through metadata which can more readily be converted to the new artifacts. For operations that you wish to reuse in the tooling, use the “My Collections” feature of the Activity Editor wherever possible.

In any Java code snippets that may need to be developed, IBM recommends that the code be as simple and atomic as possible. The level of sophistication in the Java code should be on the order of scripting, involving basic evaluations, operations, and computations, data formatting, type conversions, and so forth. If more extensive or sophisticated application logic is required, consider using EJBs running in WebSphere Application Server to encapsulate the logic, and use web service calls to invoke it from WebSphere InterChange Server. Use standard JDK libraries rather than third party or external libraries which would need to be migrated separately. Also, collect all related logic within a single code snippet, and avoid using logic where connection and transaction contexts span multiple code snippets. With database operations, for example, code related to obtaining a connection, beginning and ending a transaction, and releasing the connection should be in one code snippet.

In general, ensure that code which is designed to interface with an Enterprise Information System (EIS) is placed within adapters, and not within maps or collaboration templates. This is generally a recommended practice for architecture design. Also, this will help avoid prerequisites for third party libraries and related considerations within the code, such as connection management and possible Java Native Interface (JNI) implementations.

Make the code as safe as possible by using appropriate exception handling. Also make the code compatible to run within a J2EE application server environment, even though it is currently running within a J2SE environment. Adhere to J2EE development practices, such as avoiding static variables, spawning threads, and disk I/O. While these are generally good practices to adhere to, they can improve portability.

Premigration considerations: Common code utilities

IBM recommends that you avoid the development of common code utility libraries for use across integration artifacts within the WebSphere InterChange Server environment. Where code reuse across integration artifacts is necessary, IBM recommends that you leverage the “My Collections” feature of the Activity Editor tool. Also, consider using EJBs running in WebSphere Application Server to encapsulate the logic, and use web service calls to invoke them from WebSphere InterChange Server.

While it is possible that some common code utility libraries may execute appropriately on WebSphere Process Server, you will be responsible for the migration of the custom utilities.

Premigration considerations: Database connection pools

A WebSphere InterChange Server database connection pool within a map or collaboration template will be rendered as a standard JDBC resource in WebSphere Process Server. However, the way connections and transactions are managed might

differ between WebSphere InterChange Server and WebSphere Process Server, so you should avoid keeping database transactions active across Java snippets.

User-defined database connection pools are useful within maps and collaboration templates for simple data lookups and for more sophisticated state management across process instances. A database connection pool in WebSphere InterChange Server will be rendered as a standard JDBC resource in WebSphere Process Server, and the basic function will be the same. However, the way connections and transactions are managed may differ.

To maximize future portability, avoid keeping database transactions active across Java snippet nodes within a collaboration template or map. For example, code related to obtaining a connection, beginning and ending a transaction, and releasing the connection should be in one code snippet.

Premigration considerations: Preventing database collisions

Prevent database collisions from occurring by scheduling events to occur at least two seconds apart.

If your migrated applications cause multiple events to occur at the same time to WebSphere Business Integration components, this could cause database collisions, or deadlocks. These occur when the WebSphere Process Server Application Scheduler (AppScheduler) schedules multiple events to occur at exactly the same time. When a deadlock occurs, the event that caused it is rolled back and attempted again as soon as possible. This cycle continues until each of the threads attempting to access the database successfully updates it.

For example:

```
AppScheduler E com.ibm.wbiserver.scheduler.AppSchedulerMB process CWLWS0021E:
The AppSchedulerMB.process method has generated an exception.
WSRdbXaResour E DSRA0304E: XAException occurred. XAException contents and
  details are:
The DB2 Error message is : Error executing a XAResource.end(), Server returned
XA_RBDEADLOCK The DB2 Error code is : -4203
The DB2 SQLState is : null
```

To prevent this from occurring, schedule the events to occur far enough apart so that deadlocks do not occur. IBM recommends that you schedule events to occur at least two seconds apart; however, the amount of time you need will vary depending on other factors in your environment that affect performance such as database size, hardware, connection speed and other factors.

Premigration considerations: Business objects

For the development of business objects, use only the tooling provided to configure artifacts, use explicit data types and lengths for data attributes, and use only the documented APIs.

Business objects within WebSphere Process Server are based on Service Data Objects (SDOs). SDOs use data attributes that are strongly typed. For business objects in WebSphere InterChange Server and adapters, data attributes are not strongly typed, and users sometimes specify string data types for non-string data attributes. To avoid issues in WebSphere Process Server, explicitly specify data types.

Because business objects within WebSphere Process Server might be serialized at runtime as they are passed between components, it is important to be explicit with the required lengths for data attributes to minimize utilization of system resources.

For this reason, do not use the maximum 255 character length for a string attribute, for example. Also, do not specify zero length attributes which currently default to 255 characters. Instead, specify the exact length required for attributes.

XSD NCName rules apply to business object attribute names in WebSphere Process Server. Therefore, do not use any spaces or ":" in names for business object attributes. Business object attribute names with spaces or ":" are invalid in WebSphere Process Server. Rename business object attributes before migration.

If using an array in a business object, you cannot rely on the order of the array when indexing into the array in Maps and/or Relationships. The construct that this migrates into in WebSphere Process Server does not guarantee index order, particularly when entries are deleted.

It is important to use only the Business Object Designer tool to edit business object definitions, and to use only the published APIs for business objects within integration artifacts.

Premigration considerations: Collaboration templates

When developing WebSphere InterChange Server collaboration templates, follow these guidelines to ensure the best chance of a smooth migration to WebSphere Process Server.

To ensure processes are described appropriately with metadata, always use the Process Designer tool for the creation and modification of collaboration templates, and avoid editing the metadata files directly. Use the Activity Editor tool wherever possible to maximize the use of metadata to describe the required logic.

To minimize the amount of manual rework that may be required in migration, use only the documented APIs within collaboration templates. Avoid the use of static variables. Instead, use non-static variables and collaboration properties to address the requirements of the business logic. Avoid the use of Java qualifiers final, transient and native in Java snippets. These cannot be enforced in the BPEL Java snippets that are the result of migrating the Collaboration Templates.

To maximize future portability, avoid using explicit connection release calls and explicit transaction bracketing (that is, explicit commits and explicit rollbacks) for User Defined Database Connection Pools. Instead, make use of the container-managed implicit connection clean-up and implicit transaction bracketing. Also, avoid keeping system connections and transactions active across Java snippet nodes within a collaboration template. This applies to any connection to an external system, as well as user-defined database connection pools. Operations with an external EIS should be managed within an adapter, and code related to database operation should be contained within one code snippet. This may be necessary within a collaboration which, when rendered as a BPEL business process component may be selectively deployed as an interruptible flow. In this case, the process may be comprised of several separate transactions, with only state and global variable information passed between the activities. The context for any system connection or related transaction which spanned these process transactions would be lost.

Name collaboration template property names in accordance with W3C XML NCName naming conventions. WebSphere Process Server accepts names conforming to those conventions. Any disallowed characters are invalid in BPEL

property names that they will be migrated into. Rename properties to remove any disallowed characters before migrating to avoid syntactical errors in the BPEL generated by migration.

Do not reference variables using "this." For example, Instead of "this.inputBusObj" use just "inputBusObj"

Use class-level scoping on variables instead of scenario-scoped variables. Scenario-scoping is not carried forward during migration.

Initialize all variables declared in Java snippets with a default value: "Object myObject = null;" for example. Be sure all variables are initialized during declaration before migrating.

Ensure that there are no Java import statements in the user modifiable sections of your collaboration templates. In the definition of the collaboration template, use the import fields to specify Java packages to import.

Do not set incoming business object values to be stored in the *triggeringBusObj* variable. Within WebSphere InterChange Server, the *triggeringBusObj* is read-only and its values cannot be overwritten, so any incoming business object values will not be saved. If the *triggeringBusObj* is used as the receiving variable for an incoming business object on an inbound service call, then after migration the behavior of the inbound service call will be different: within the BPEL process, the incoming value from the inbound service call will overwrite the value stored in *triggeringBusObj*.

Premigration considerations: Maps

When developing WebSphere InterChange Server maps, follow these guidelines to ensure the best chance of a smooth migration to WebSphere Process Server.

To ensure maps are described appropriately with metadata, always use the Map Designer tool for the creation and modification of maps, and avoid editing the metadata files directly. Use the Activity Editor tool wherever possible to maximize the use of metadata to describe the required logic.

When referencing child business objects in a map, use a submap for the child business objects.

Avoid using Java code as the "value" in a SET since that is not valid in WebSphere Process Server. Use constants instead. For example, if the set value is "xml version=" + "1.0" + " encoding=" + "UTF-8" this will not validate in WebSphere Process Server. Instead, change it to "xml version=1.0 encoding=UTF-8" before you migrate.

To minimize the amount of manual rework that may be required in migration, use only the documented APIs within maps. Avoid the use of static variables. Instead, use non-static variables. Avoid the use of Java qualifiers final, transient and native in map custom code.

If using an array in a business object, do not rely on the order of the array when indexing into the array in maps. The construct that this migrates into in WebSphere Process Server does not guarantee index order, particularly when entries are deleted.

To maximize future portability, avoid using explicit connection release calls and explicit transaction bracketing (that is, explicit commits and explicit rollbacks) for User Defined Database Connection Pools. Instead, make use of the container-managed implicit connection clean-up and implicit transaction bracketing. Also, avoid keeping system connections and transactions active in custom map steps across transformation node boundaries. This applies to any connection to an external system, as well as user-defined database connection pools. Operations with an external EIS should be managed within an adapter, and code related to database operation should be contained within one custom step.

Do not use inner classes in your maps. The migration command (`reposMigrate`) does not migrate inner classes and you will receive errors if your maps contain them. In a WebSphere InterChange Server repository, an inner class could be defined in a node and referenced by other nodes within the same collaboration template. In WebSphere Process Server, an inner class defined in a BPEL component cannot be used by other components. Due to this limitation, inner classes are not translated and must be dealt with manually. Recommended changes include packaging the inner class code in a library as an external class, or removing the inner class declaration, resolving any errors, and placing the code as needed throughout the BPEL.

Premigration considerations: Relationships

While relationship definitions can be migrated for use in WebSphere Process Server, the relationship table schema and instance data may be reused by WebSphere Process Server, and shared concurrently between WebSphere InterChange Server and WebSphere Process Server.

For relationships, use only the tooling provided to configure the related components, and use only the published APIs for relationships within integration artifacts.

Use only the Relationship Designer tool to edit relationship definitions. In addition, allow only WebSphere InterChange Server to configure the relationship schema, which is generated automatically upon deployment of relationship definitions. Do not alter the relationship table schema directly with database tools or SQL scripts.

If you must manually modify relationship instance data within the relationship table schema, be sure to use the facilities provided by the Relationship Manager.

Use only the published APIs for relationships within integration artifacts.

Premigration considerations: Access framework clients

Do not develop any new clients adopting the CORBA IDL interface APIs. This is not supported in WebSphere Process Server.

Migrating WebSphere InterChange Server artifacts with the `reposMigrate` command

Migrate WebSphere InterChange Server artifacts to WebSphere Process Server artifacts with the `reposMigrate` command.

Before you begin

Note: The functionality of the `reposMigrate` command is also available from WebSphere Integration Developer with a supporting wizard (graphical user interface). See the WebSphere Integration Developer information center for more information.

The `reposMigrate` command requires as input a WebSphere InterChange Server repository JAR file. This JAR file should be self-contained with respect to the applications being migrated. That is, all artifacts referenced by any of the artifacts in the JAR file must also be contained in the JAR file.

To ensure that the repository JAR file that will be generated is self-contained, run the `repos_copy` command with the `-vr` option before exporting the server repository. This validates the repository. If the repository is valid then `repos_copy` writes the following output to the console: `Validation Succeeded. All Dependencies Resolved`. If the repository is not valid then `repos_copy` prints a list of the dependencies that must be resolved. Resolve the dependencies prior to exporting the repository.

Export the repository artifacts and create the repository JAR file, using the WebSphere InterChange Server `repos_copy` command with the `-o` option (See the WebSphere InterChange Server v4.3 documentation for more details, including how to export individual components).

About this task

The `reposMigrate` command will convert all of the WebSphere InterChange Server artifacts in a JAR file into WebSphere Process Server deployable artifacts. These artifacts are modules created as one or more JAR files. A JAR file is created for each collaboration object and for each connector definition that has been migrated. For other artifacts such as business objects, maps and relationships, a copy of all of these artifacts generated from the input JAR file will be included in each JAR file generated. If no collaboration objects or connectors are migrated, a single JAR file is created containing a module of all the shared artifacts. After the new JAR files are created, you will use the `serviceDeploy` command to generate the EAR files that can be deployed in WebSphere Process Server.

For WebSphere InterChange Server artifacts that have no corresponding artifact in WebSphere Process Server, a Jython script is generated during migration that can be run using the `wsadmin` command to create WebSphere Process Server configuration definitions corresponding to the original WebSphere InterChange Server artifacts.

Procedure

1. Identify the JAR file containing the pre-exported WebSphere InterChange Server artifacts that are to be converted to WebSphere Process Server deployable artifacts.
2. Invoke the `reposMigrate` command from a command-line prompt. Type the command at a command prompt in WebSphere Process Server, with the required arguments and any optional arguments you require. Refer to “`reposMigrate` command” on page 126 for more information.
3. If desired, edit the resulting JAR file.
4. Run `serviceDeploy` to create a deployable EAR file for each JAR file.

Note: The support in the WPS Runtime to handle migrated ICS applications relies on the default naming convention used by the `serviceDeploy` command.

IBM recommends that you do not specify the **serviceDeploy** **-outputApplication** parameter when building migrated projects with the **serviceDeploy** command, so that it will generate its default output filenames. For more information, refer to the WebSphere Process Server **serviceDeploy** command in the *Reference* PDF file.

5. Use the administrative console or the **wsadmin** command to install the EAR files on WebSphere Process Server. Use the **wsadmin** command to execute the `InstallAdministrativeObjects.py` script. This will create resources in the WebSphere Process Server system for all target resources such as JDBC data sources and WBI Scheduler entries.

Example

You can use the **reposMigrate** command to migrate existing WebSphere InterChange Server artifacts directly to a running WebSphere Process Server:

1. Open a command prompt in WebSphere Process Server.
2. Issue the **reposMigrate** command with the following mandatory parameters:
install_root\bin\reposMigrate SourceArtifactJAR OutputArtifactDirectory

The **reposMigrate** command builds the generated artifacts as follows:

- For each WebSphere InterChange Server collaboration object and connector definition in the input JAR file, **reposMigrate** creates a JAR file from the migrated artifacts.
- For other artifacts such as business objects, maps and relationships, a copy of all of these artifacts generated from the input JAR file will be included in each JAR file generated. If no collaboration objects nor connector definitions were in the input, a single JAR file will be created with all the shared artifacts.

The default behavior of the **reposMigrate** command is to log errors for the migration of each individual artifact and continue to migrate the remainder of the artifacts. You should check output messages for errors after the execution completes. To view the output, use the logfile parameter (**-IfLogFileName**) to direct the output into the specified file. To override this default behavior and force **reposMigrate** to end processing when the first artifact that cannot be migrated is encountered, specify the **-fh** (halt at first failure) flag. You can run **reposMigrate** from the beginning to retry after a failed execution.

Related reference

“Postmigration considerations” on page 128

When applications have been migrated from WebSphere InterChange Server to WebSphere Process Server, special attention is required in some areas to enable migrated applications to function in WebSphere Process Server consistently with their intended function due to differences between the architectures of WebSphere Process Server and WebSphere InterChange Server.

Related information

Wsadmin tool

WebSphere InterChange Server v4.3 documentation

WebSphere Integration Developer information center

reposMigrate command

Provides details of the syntax and usage of the **reposMigrate** command.

Purpose

The `reposMigrate` command migrates WebSphere InterChange Server artifacts to WebSphere Process Server deployable artifacts.

Location

The command file is located in the `install_root/bin` directory. The command file is a script named as follows:

- `UNIX` `Linux` For Linux and UNIX-based operating systems: `reposMigrate.sh`
- `Windows` For Windows systems: `reposMigrate.bat`
- `i5/OS` For the i5/OS operating system: `reposMigrate`

Syntax

```
reposMigrate [-es] [-td TemplateDirectory] [-ml] [-fh] [-lv] [-wi] SourceArtifactJAR OutputArtifactDirectory
```

The `SourceArtifactJAR` and `OutputArtifactDirectory` arguments are required.

Parameters

[-es]

Optional parameter. Requests Event Sequencing be enabled for all Asynchronous WSDL methods. If this option is not present, the default is for migration to not enable Event Sequencing on any WSDL methods.

[-td *TemplateDirectory*]

Optional parameter. Requests that all Assembly Editor templates in the specified directory be loaded and used for XML to Java Conversion. The default for this property is for only the Standard Assembly Editor Template v4.3.3 to be used for XML to Java Conversion.

[-ml]

Optional parameter. Requests that all loops present in a Collaboration Template be maintained. If this option is not present, the default is for migration to use loop unraveling. See “Collaboration migration” on page 135 for more information on loops.

[-fh]

Optional parameter. By default, `reposMigrate` will continue processing the remaining artifacts in the JAR file if an error occurs during the processing of a certain artifact. If this option is set, the processing will stop as soon as an error is detected. The artifact with the error and all subsequent artifacts are not processed.

[-lv]

Optional parameter. Sets the log level to verbose.

[-wi]

Optional parameter. By default, migration of an individual artifact will fail if a Java conversion problem is found. If this option is set, all Java conversion problems will be treated as warnings only, and the artifact will be migrated as well as possible.

SourceArtifactJAR

Required parameter. Specifies the WICS repository JAR file that is to undergo migration.

OutputArtifactDirectory

Required parameter. Specifies the output directory in which the generated module jar files will be placed.

Examples

Windows This example on a Windows system migrates existing WebSphere InterChange Server artifacts and places them in the MigratedArtifacts directory. Java conversion warning messages are ignored and the log level is set to verbose.

```
install_root\bin\reposMigrate.bat -wi -lv C:\inputRepos.jar  
C:\IBM\WebSphere\MigratedArtifacts
```

UNIX **Linux** This example on a UNIX-based system migrates existing WebSphere InterChange Server artifacts and places them in the MigratedArtifacts directory. Java conversion warning messages are ignored and the log level is set to verbose.

```
install_root/bin/reposMigrate.sh -wi -lv /inputRepos.jar  
/opt/IBM/WebSphere/MigratedArtifacts
```

This example on an i5/OS system migrates existing WebSphere InterChange Server artifacts and places them in the MigratedArtifacts directory. Java conversion warning messages are ignored and the log level is set to verbose.

```
install_root/bin/reposMigrate -wi -lv /inputRepos.jar  
/home/user_name/MigratedArtifacts
```

where *user_name* is the name of the i5/OS user profile invoking the reposMigrate script.

Postmigration considerations

When applications have been migrated from WebSphere InterChange Server to WebSphere Process Server, special attention is required in some areas to enable migrated applications to function in WebSphere Process Server consistently with their intended function due to differences between the architectures of WebSphere Process Server and WebSphere InterChange Server.

You should be aware of the information described in the following sections if it applies to your application and environment:

“Security” on page 129

“Handling existing database connections, relationships, and scheduled events (InstallAdministrativeObjects.py script)” on page 129

“Handling existing WebSphere InterChange Server database connection pools” on page 130

“Using an existing WebSphere InterChange Server relationship database” on page 130

“Migrating scheduled events” on page 131

“Access Enterprise JavaBean (EJB) support” on page 131

“DynamicSend API configuration” on page 132

“Enabling the BaseCollaboration.dynamicSend method call” on page 132

“Event sequencing migration” on page 134

“Failed events” on page 134

“Map migration” on page 134

“Collaboration migration” on page 135

“BPEL variables must be defined after migration” on page 136

“Enabling logError API e-mail notification on WebSphere Process Server” on page 136

“Handling asynchronous calls in WebSphere Process Server” on page 137

“Enabling AppScheduler to start after network deployment upgrade” on page 137

“Handling correlation values in WebSphere Process Server” on page 137

“Packaging and deploying migrated applications” on page 138

Security

Additional security configuration is required for your applications to have the same security levels set as they had when running in WebSphere InterChange Server. For details on this configuration, refer to “Configuring global security after WebSphere InterChange Server migration” on page 139.

Handling existing database connections, relationships, and scheduled events (InstallAdministrativeObjects.py script)

The Jython script `InstallAdministrativeObjects.py` is generated during migration. This script has three purposes: It allows migration of WebSphere InterChange Server scheduler entries that have no corresponding artifact in WebSphere Process Server; it allows the use of existing `DBConnection` pools; and it allows the use of an existing relationship database. You can run the script with the `wsadmin` command to create WebSphere Process Server configuration definitions corresponding to the original WebSphere InterChange Server artifacts. A copy of `InstallAdministrativeObjects.py` is included wherever the shared artifacts are included. That is, the script is included in every JAR file created by the `reposMigrate` command, and it is put in the shared library project specified during import in WebSphere Integration Developer. An `InstallAdministrativeObjects.py` script is always generated even if there are no artifacts that require it. This script can be modified to add or delete entries before using the `wsadmin` command to execute it.

For more information about using the `wsadmin` command, see `wsadmin` tool .

Handling existing WebSphere InterChange Server database connection pools

To preserve existing WebSphere InterChange Server DataBase Connection pools for use by WebSphere Process Server you can run the `InstallAdministrativeObjects.py` script with the `wsadmin` command to create the connection pools in WebSphere Process Server. If an appropriate JDBC provider is not defined, this script will use default JDBC provider templates to create JDBC providers. A side effect of using these default templates is that WebSphere Process Server creates an empty, sample data source definition. This sample data source is not used; you must delete it to prevent exceptions from occurring during server start because it does not specify all of the information required for a data source.

In the WebSphere InterChange Server environment, resources are defined only once for the entire system. To simulate this in the WebSphere Process Server environment, the `InstallAdministrativeObjects.py` script defines resources at the cell scope. WebSphere variables are predefined at the node scope in the WebSphere Process Server system for use by JDBC providers created from the default JDBC provider templates. These variables are defined at the node scope so that they can be customized for each node. Because of this scoping discrepancy, you will need to do one of the following:

- Define the WebSphere Variables needed by the created JDBC providers at the cell scope.
- Run the `InstallAdministrativeObjects.py` script and then move the JDBC providers to the node scope.

Use the administrative console to examine the JDBC providers that are generated to determine which WebSphere variables are needed. From the administrative console, select **Environment > WebSphere Variables** to create any required variables. For more information, see *Defining WebSphere variables in the WebSphere Application Server Network Deployment, version 6.1 information center*.

Here is an example of what the generated `InstallAdministrativeObjects.py` script might contain to generate the JDBC connector pool:

```
dsName = "sqls"  
create_datasource(dsName, JNDI_PREFIX + dsName, DATASOURCE_DESCRIPTION,  
MS_SQL_JDBC_PROVIDER_NAME, MS_SQL_JDBC_PROVIDER_TYPE, "icsadmin", "icsadmin",  
4, 50, "qaxs17", "1433", "wicsrepos")
```

For more information about the `wsadmin` command, see `wsadmin` tool .

Using an existing WebSphere InterChange Server relationship database

To use an existing WebSphere InterChange Server relationship database in WebSphere Process Server, you can use the `InstallAdministrativeObjects.py` script with the `wsadmin` command to create the data source and relationship configuration information in WebSphere Process Server. Normally, WebSphere Process Server automatically creates the configuration information for the migrated relationships when they are deployed. To be able to use the existing database, the `InstallAdministrativeObjects.py` script has to create the database connection for the existing WebSphere InterChange Server relationship database and the relationship configuration information in WebSphere Process Server. Runs the `InstallAdministrativeObjects.py` script before you deploy the migrated components.

Then, when WebSphere Process Server deploys the relationships it will use the configuration information that was generated by the script.

Here is an example of what the generated `InstallAdministrativeObjects.py` script might contain to generate the relationship database connection:

```
dsName = "ContactR"
create_datasource(dsName, JNDI_PREFIX + dsName, DATASOURCE_DESCRIPTION,
MS_SQL_JDBC_PROVIDER_NAME, MS_SQL_JDBC_PROVIDER_TYPE, "icsadmin", "icsadmin",
-1, -1, "9.26.230.56", "1433", "wicsrepos")

create_relationship("ContactR", "jdbc/wbi60migration/ContactR", "false")
create_role("ContactR", "ID1", "", "null", "", "null")
create_attribute("ContactR", "ID1", "JtextEmployeeID")
create_role("ContactR", "ID2", "", "null", "", "null")
create_attribute("ContactR", "ID2", "EmployeeID")
create_role("ContactR", "ID3", "", "null", "", "null")
create_attribute("ContactR", "ID3", "EmployeeID")
```

For more information about the `wsadmin` command, see `wsadmin` tool .

Migrating scheduled events

Because there is no WebSphere Process Server component that corresponds to WebSphere InterChange Server scheduler entries, migration of WebSphere InterChange Server scheduler entries is accomplished by extracting the pertinent data from the existing WebSphere InterChange Server repository JAR file and creating corresponding entries in the WebSphere Process Server scheduler tables in the WebSphere Process Server Common database. The data is represented in string form in the Jython script. To create the scheduler entries in the WebSphere Process Server database, you can run the `InstallAdministrativeObjects.py` script with the `wsadmin` command.

Here is an example of what the generated `InstallAdministrativeObjects.py` script might contain to generate the scheduler entry:

```
create_scheduler_entry("true", "stop", "JDBCCconnector", "Connector",
"2006-09-07T10:44:29.000PDT", "undefined", 0, 0)
create_scheduler_entry("true", "start", "JTextConnector", "Connector",
"2006-09-07T10:47:06.000PDT", "undefined", 0, 0)
create_scheduler_entry("true", "stop", "jtext_jdbcCollab", "Collaboration",
"2006-09-07T10:48:10.000PDT", "undefined", 0, 0)
create_scheduler_entry("true", "start", "jtext_jdbcCollab", "Collaboration",
"2006-09-07T10:48:10.000PDT", "undefined", 0, 0)
create_scheduler_entry(true, "START", "JDBCCconnector", "Connector",
"2006-10-22T12:34.56.789CDT", "MINUTES", 20, 0):
```

Access Enterprise JavaBean (EJB) support

WebSphere InterChange Server supports the triggering of collaborations by client code via the J2EE EJB (Enterprise JavaBeans™) protocol. Support for this method of triggering collaborations is referred to as "AccessEJB" or "AccessEJB for EJB" support. For backward compatibility, WebSphere Process Server provides support for AccessEJB. The AccessEJB support assumes that the SCA BPEL modules to be invoked were generated by the WebSphere InterChange Server migration tools described in this documentation. The mapping from the collaboration name and port name (that is, the input parameters for the AccessEJB) to the SCA module name, interfaces and business object types assume the conventions used by the migration tools. The AccessEJB support in WebSphere Process Server is delivered in the `AccessEJB.zip` project interchange file. This file is located in the `install_root/HeritageAPI` directory. The AccessEJB support consists of an EJB

(AccessEJB) which references an SCA module project (DynamicRouting) that invokes the SCA BPEL module. This SCA BPEL module is the migrated version of the collaboration that was invoked in WebSphere InterChange Server. The DynamicRouting module uses a selector component to select the correct SCA target based on the collaboration name and port name passed to the AccessEJB. To enable AccessEJB support in WebSphere Process Server, do the following:

1. Import the WebSphere InterChange Server repository containing the collaboration that is the target of the AccessEJB invocation into WebSphere Integration Developer.
2. Import the AccessEJB.zip project interchange file into WebSphere Integration Developer.
3. Open the DynamicRouting project and update the selector table to include the migrated module that is to be invoked via the AccessEJB.
4. Go to the migrated project containing the BPEL component to be invoked via the AccessEJB EJB, and drag the export that references the BPEL module over to the DynamicRouting project.
5. Repeat steps 3 and 4 for each BPEL module that is to be accessible via AccessEJB.
6. Build the project and deploy it to the WebSphere Process Server server.
7. Ensure that any required data handlers are provided in the runtime class path of the WebSphere Process Server server.
8. To enable your Access client to use WebSphere Process Server, ensure that it points to the WebSphere Process Server server and uses the JNDI name `com/crossworlds/access/business/cwsession/CwSession` when looking up the Access EJB.

DynamicSend API configuration

In WebSphere InterChange Server, the DynamicSend API can be used to directly invoke one collaboration from another. The collaboration to be invoked does not have to be predetermined; instead, it can be determined dynamically at runtime. The support for the DynamicSend API in WebSphere Process Server uses the DynamicRouting project described in "Access Enterprise JavaBean (EJB) support" on page 131." Follow the instructions in "Enabling the BaseCollaboration.dynamicSend method call" to enable the DynamicSend API to be able to invoke the specified BPEL modules.

Enabling the BaseCollaboration.dynamicSend method call

To enable the WebSphere InterChange Server BaseCollaboration.dynamicSend method call to work correctly after migration you must modify the DynamicRouting Projects in the AccessEJB Project Interchange file. This requires two main procedures:

1. Migrating the WebSphere InterChange Server repository.
2. Enabling the DynamicSend API.

To migrate the WebSphere InterChange Server repository:

1. Import the WebSphere InterChange Server repository containing the collaboration that invokes DynamicSend API into WebSphere Integration Developer.
2. Import the WebSphere InterChange Server repository containing the collaboration or connector that is the target of the DynamicSend API invocation into WebSphere Integration Developer.

3. Build all, and correct all errors.

To enable the DynamicSend API:

1. Import the AccessEJB.zip project interchange file into WebSphere Integration Developer.
2. Open the DynamicRouting project, and add the WebSphere InterChange Server shared library into the DynamicRouting project's dependencies.
3. Go to the migrated module containing the component to be invoked through the BaseCollaboration.dynamicSend method, and drag the export that references the module over to the DynamicRouting project. Choose **Import with SCA Binding** and then click **OK**.
4. In the DynamicRouting Assembly Diagram window, copy and paste PreRoute_TargetCollab_TargetPort and then rename the newly created copy to PreRoute_ModuleName_ExportName (the name of the copied import will be PreRoute_TargetCollab_TargetPortCopy).
5. On PreRoute_ModuleName_ExportName, left click on the reference, which is the small box attached to the right containing 1.1. Right click and choose **Delete**.
6. Wire the PreRoute_ModuleName_ExportName to the import generated in step 3. Respond with "no" to the Java WSDL reference question.
7. Rename the import to ModuleName_ExportName. Save the changes to the Assembly Diagram.
8. Update the selector table in the DynamicRouting project to include the migrated module that is to be invoked through the DynamicSend API.
 - a. Switch to the Java Perspective Package Explorer View. Expand DynamicRouting/com.ibm and open RoutingSelector.selt with the text editor.
 - b. Copy the OperationSelectionRecord block and paste the entire block immediately following the existing block.
 - c. In the new block, change
 componentName="PreRoute_TargetCollab_TargetPort" to
 componentName="PreRoute_ModuleName_ExportName". Also in the new block, change value="TargetCollab_TargetPort" to value="ModuleName_ExportName".


```
<OperationSelectionRecord>
  <SelectionKey>
    <SelectionKeyElement xsi:type="selt:StringSingletonKey" value="
      TargetCollab_TargetPort"/>
  </SelectionKey>
  <SelectionData xsi:type="selt:SCAInternalComponent"
    componentName="PreRoute_TargetCollab_TargetPort"/>
</OperationSelectionRecord>
<OperationSelectionRecord>
  <SelectionKey>
    <SelectionKeyElement xsi:type="selt:StringSingletonKey"
      value="ModuleName_ExportName"/>
  </SelectionKey>
  <SelectionData xsi:type="selt:SCAInternalComponent" componentName
    ="PreRoute_ModuleName_ExportName"/>
</OperationSelectionRecord>
```
 - d. Save and close RoutingSelector.selt.
9. Generate the implementation file.
 - a. Expand **com.ibm.sel** and copy PreRoute_TargetCollab_TargetPortImpl.java and paste it to the same location. Name the newly created Java file PreRoute_ModuleName_ExportNameImpl.java.

- b. Edit `PreRoute_ModuleName_ExportNameImpl.java`. Change the method name of `locateService.TestBOInterfacePartner` to `locateService_InterfaceNamePartner` (`InterfaceName` is the method). Change `TestBOInterfacePartner` to `InterfaceNamePartner`.
 - c. Search for "`locateService_TestBOInterfacePartner`" in `PreRoute_ModuleName_ExportNameImpl.java`, and change its name to `locateService_InterfaceNamePartner`.
10. Switch back to the Business Integration Perspective. Open the `DynamicRouting` Assembly Diagram. Click on **PreRoute_ModuleName_ExportName**. Open **Properties** and select **Implementation**. In the **Class:** field, enter `com.ibm.sel.PreRoute_ModuleName_ExportNameImpl`.
11. Save all changes.
12. Repeat steps 3 to 11 for any other modules that you want to call from the `BaseCollaboration.dynamicSend` Method. There is currently no way to "dynamically look these modules up" if you do not add them to the `DynamicRouting` Table so you can access them during run time.
13. For the project that calls the `dynamicSend` API, do the following
 - a. Copy and paste Interface "`RoutingPacket`" from Module `DynamicRouting`.
 - b. In the component that calls the `dynamicSend` method. Add the newly copied interface "`RoutingPacket`" to `Reference_Partners`, and rename it "`RoutingPacketPartner`."
 - c. Save it.
 - d. Open the Assembly Diagram. Drag "`RoutingInput`" from `DynamicRouting`. Choose "Import with SCA Binding" and click "OK". Rename it from "`Import1`" to "`DynamicRouting`".
 - e. Delete and re-drag the component that calls the `dynamicSend` API to the Assembly Diagram window, wire the Reference "`RoutingPacketPartner`" to "`DynamicRouting`" and re-wire the other references.
14. Save all and build, then correct all errors. Export all the modules to EAR files.

Event sequencing migration

Methods are available for sequencing events with WebSphere Process Server in ways similar to the way you could with WebSphere InterChange Server. Articles on this subject that you might find helpful are available from the IBM developerWorks web site. Search in the "Technical Library" at <http://www.ibm.com/developerworks>.

Failed events

Methods for handling failed events in WebSphere Process Server are described in article(s) that you might find helpful on the IBM developerWorks web site. Search in the "Technical Library" at <http://www.ibm.com/developerworks>.

Map migration

WebSphere InterChange Server migration converts WebSphere InterChange Server maps into WebSphere Process Server maps. Two output maps are generated: the business graph map and the business object map. The business graph map calls the business object map as a submap. All the business graph maps are identical, except for their name and the name of the submap they call. These business graph maps are present only to satisfy the necessary mapping steps that can only be

done at the business graph level. The business object maps are each unique and are the migrated form of the WebSphere InterChange Server map. If the WebSphere InterChange Server input map contains custom messages for the supported WebSphere InterChange Server API log methods, these messages will be converted into a properties file.

Collaboration migration

Collaboration Templates: The WebSphere InterChange Server to WebSphere Process Server migration tools migrate WebSphere InterChange Server Collaboration templates into WebSphere Process Server BPEL files. One BPEL file is created for each triggering port defined in a collaboration template, and its name is based on the following naming convention:

CollaborationTemplateName_TriggeringPortName. Each BPEL file receives a business object type that is based on the business object type associated with the triggering port. For example, if the triggering port takes a business object type of Customer, then the BPEL file that is created will have a “TriggeringBusObj” variable type of Customer.

Maintaining Loops: When migrating a WebSphere InterChange Server repository, you have the option to either maintain loops that might have existed in the collaboration template, or unravel these loops. If you chose to maintain the loops and you view the resulting migration JAR files in WebSphere Integrated Developer, validation errors will appear because the WebSphere Integrated Developer tooling does not currently support loops. (WebSphere Process Server does support loops.) To remove the validation errors in WebSphere Integrated Developer, the BPEL file needs to be manually updated. The drawback to changing the BPEL file is that it will no longer be viewable from within the WebSphere Integrated Developer BPEL editor. For more information on how to perform the specific changes required to remove the validation errors with WebSphere Integrated Developer, see the “Loop Unraveling” section of the IBM developerWorks article Migrating WebSphere InterChange Server artifacts to WebSphere Process Server artifacts, Part 1: Migrating collaboration templates to BPEL .

Not maintaining loops: If you chose to not maintain the loops, the collaboration template loop is converted into a BPEL while loop. A variable initialized to true allows the loop to execute at least once. After the initial execution of the loop, the loop variable is set to false and then the normal loop condition will determine if the loop should continue to execute or terminate. (Note: This option does not cause validation errors within WebSphere Integrated Developer.)

Collaboration Objects: The WebSphere InterChange Server to WebSphere Process Server migration tools migrate collaboration objects into several service component architecture (SCA) components. Currently, migration supports collaboration objects that reference collaboration templates as follows:

- Supported:
 - One or more triggering ports, no correlation sets and no asynchronous in calls
 - Exactly one triggering port, correlation sets and asynchronous in calls
- Unsupported:
 - Migration does not support the case of one or more triggering ports, correlation sets and asynchronous in calls. In this case, the resulting artifacts are migrated as if they were the first case listed above. Additionally you will need to manually create the missing SCA components and wire them together appropriately.

SCA Components:

- Exports: An export is created for every triggering port defined in the collaboration template associated with the collaboration object. The export name is *TriggeringPortName*.
- Export to BPEL: An interface map is generated that maps the data from the export to the BPEL file. The interface map name is *Export_To_BPELname*. When there is exactly one triggering port and the collaboration template has an asynchronous in call, additional SCA components are created. Instead of having just one interface map, the migration results in two interface maps: one for synchronous calls, and another for asynchronous calls. A Java component is used to decide which of these two interface maps to follow.
- BPEL: For every triggering port, the export will be wired to an interface map and that interface map will map to an instance of the BPEL file.
- BPEL to import: Every port, triggering and non-triggering, has an interface map mapping the BPEL file to the import. The interface map name is *BPEL_to_Port*.
- Import: Finally, an import file is created. The import name is *ConnectorName_BONameBG*.

For further detailed information on how collaboration templates are migrated to WebSphere Process Server BPEL files, see the IBM developerWorks article [Migrating WebSphere InterChange Server artifacts to WebSphere Process Server artifacts, Part 1: Migrating collaboration templates to BPEL](#) .

BPEL variables must be defined after migration

Problem: A variable that is not defined in the Ports definitions of the WebSphere InterChange Server Collaboration Template is used to invoke a partner. After migration, the variable is referenced in the business process execution language (BPEL) invoke but has not been set up as a BPEL variable, so it is flagged as an error when executing the serviceDeploy command against the module or after building the module in WebSphere Integration Developer. **Cause:** When invoking a partner from a BPEL process in WebSphere Process Server, any object used in the invoke must be declared as a BPEL variable so that the type of the object being used can be determined. During migration, only the Ports declarations in the Collaboration Template are examined to determine what BPEL variables need to be declared. For global variables, or variables declared in snippets elsewhere in the ICS Collaboration Template definition, the migration code cannot reliably determine the object type, so BPEL variables are not declared for these in the BPEL file generated by migration. **Solution:** After migration, you must define the variable as a BPEL variable for the variable to be referenced during an invoke.

Enabling logError API e-mail notification on WebSphere Process Server

Problem: After migration to WebSphere Process Server, the WebSphere InterChange Server logError API does not send an e-mail to a list of users that has been configured in WebSphere Interchange Server. **Cause:** In WebSphere InterChange Server, you could configure the API call logError to send an error e-mail to a specified list of users. However, this list of users, configured on the server, is not accessible to the migration code so it must be set up manually in WebSphere Process Server. **Solution:**To enable the WebSphere InterChange Server logError e-mail notification functionality in WebSphere Process Server, a new BPEL environment variable called *LOGERROR_EMAIL_LIST* is created in each BPEL file generated by migration. Set this variable with the list of e-mail users needing to receive log error e-mails. Separate the names within the list with a comma.

Handling asynchronous calls in WebSphere Process Server

Problem: Async-in events act as triggering events when both types of events can be received on the same connector. **Cause:** If both async-in and triggering events can be received on the same connector, the migrated application cannot determine which events are which type. By default, all events are treated as triggering events in a migrated application in this scenario. **Solution:** Application-specific logic that can determine if an event is async-in or triggering must be added to the migrated application. Migrated modules that can receive triggering events and async-in events on the same connector will have a component named JavaSelector. The implementation code for the JavaSelector component will contain the AsyncIn() method shown below. This method must be updated with logic to check if events are async-in or triggering. This logic will be specific to each application and will be based on the nature of the events being handled.

```
/** * Method generated to support async inbound service call routing */
public boolean isAsyncIn()
{ //Add custom code here
  //TODO
  return false;
}
```

Enabling AppScheduler to start after network deployment upgrade

Problem: After migrating a WebSphere Process Server 6.0.1.x network deployment configuration to WebSphere Process Server 6.1, the AppScheduler fails to start on WebSphere Process Server 6.0.1.x servers and clusters that have not been upgraded. An Exception similar to the following will be generated:

```
WSVR0040E: addEjbModule failed for WBISchedulerEJB.jar
[class com.ibm.ws.runtime.component.
DeployedEJBModuleImpl] java.lang.NoClassDefFoundError:
com/ibm/wbiserver/scheduler/common/AppSchedulerException
```

Cause: After migrating the WebSphere Process Server 6.0.1.x network deployment configuration to WebSphere Process Server 6.1, the AppScheduler application looks for the AppSchedulerException Class in the WebSphere Process Server 6.0.1.x version of the wbischedulercommon.jar file and fails to find it in the local system's *install_root/lib* directory. It then throws a java.lang.NoClassDefFoundError: com/ibm/wbiserver/scheduler/common/AppSchedulerException exception.

Solution: Replace the WebSphere Process Server 6.0.1.x version of the wbischedulercommon.jar file with the WebSphere Process Server 6.1 or WebSphere Process Server 6.0.2.x version of that JARfile. You can obtain the new JAR file from the WebSphere Process Server 6.1.x *install_root*AppScheduler/lib directory or the WebSphere Process Server 6.0.2 *install_root/lib* directory. Copy the JAR file into the WebSphere Process Server 6.0.1.x lib directory and replace the existing JAR file. Do not rename the existing JAR file and leave it in the lib directory, because WebSphere Process Server picks up all files in the lib directory as JAR files regardless of the extension. Then, restart the server or cluster so that WebSphere Process Server picks up the new JAR file.

Handling correlation values in WebSphere Process Server

Problem: In WebSphere Process Server, new events attempting to use the existing correlation values will fail. In such instances, the error message

```
CWWBE0074E: Correlation violation in activity 'null' for correlation set
'CorrelationSetA'java.sql.
SQLException: Could not insert new row - duplicate value in a UNIQUE INDEX column
```

appears. **Cause:** When a collaboration or process instance completes in WebSphere InterChange Server, data related to that instance is deleted except for cases dealing with failures. In WebSphere Process Server, the persistence of process instance-related data is controlled by the business process execution language (BPEL) option, "Automatically delete the process after completion." BPEL files generated by the WebSphere InterChange Server to WebSphere Process Server migration wizard will not have this option selected. As a result, process instance data will persist, even after the process instance completes, until you clean it up manually. When a process defines a correlation set, the correlation values locked by process instances remain locked as long as the process instance data is persisted, even after the process has completed. As a result, new events attempting to use the same correlation values will fail as long as the previous process instance's data persists. This behavior will be different than in WebSphere InterChange Server, where new events with duplicate correlation set values could be processed as soon as the previous instance was complete. **Solution:** To simulate the behavior of WebSphere InterChange Server with respect to multiple events with duplicate correlation set values, you can choose to select the BPEL option "Automatically delete the process after completion" so that process instance data is deleted, and the correlation value is unlocked, as soon as the process instance completes. Before selecting this option, you should investigate and fully understand the way failures are handled in WebSphere Process Server and to ensure that your failed event strategy does not rely on data that will be automatically deleted when this option is set.

Packaging and deploying migrated applications

After migrating the WebSphere InterChange Server repository using the `reposMigrate` command, you will need to package the resulting JAR files into EAR files in order for them to be deployed to the WebSphere Process Server. To do this, you can either import each migration-generated JAR file into WebSphere Integration Developer and export the modules as EAR files, or you can use the `serviceDeploy` command. The `serviceDeploy` command accepts a JAR file as input and outputs a deployable EAR file. Packaging the migration code into EAR files involves compiling the resulting migrated JAR file. If this produces validation errors, they are most likely due to the use of unsupported WebSphere InterChange Server APIs or third-party APIs that were present in WebSphere InterChange Server but have not yet been included in the WebSphere Process Server class path. Remove the unsupported APIs and add the third-party classes to the WebSphere Process Server class path.

Validation errors may also be caused by not following premigration recommended practices or may indicate postmigration work that still needs to be performed on the artifacts. As with migration errors, each validation error should be handled on a per-error basis. If a recommended premigration practice was not followed, you can update the repository and migrate it again, or you can edit the output artifacts to remove the problem.

Any other validation errors should be resolved as if these artifacts were created from scratch. You should refer to the validator documentation that outlines common artifact errors and their resolutions. Inevitably an automated migration cannot account completely for your program's intent; it can just make best guesses. Therefore, even if there are no validation errors, it is possible the migrated artifacts do not perform as intended. You should review all artifacts to confirm that the intended purpose of each artifact is met by its migrated content.

Related concepts

“Limitations when migrating from WebSphere InterChange Server” on page 163
Some characteristics of WebSphere InterChange Server are not precisely duplicated by WebSphere Process Server. Therefore you might have to modify your applications after migration to get them to perform as they did in WebSphere InterChange server.

“Troubleshooting migration from WebSphere InterChange Server” on page 165
Find solutions to problems you encounter with migration as well as instructions for turning on logging and tracing.

Related tasks

“Migrating WebSphere InterChange Server artifacts with the `reposMigrate` command” on page 124

Migrate WebSphere InterChange Server artifacts to WebSphere Process Server artifacts with the **`reposMigrate`** command.

Related reference

 `serviceDeploy` command

Use the `serviceDeploy` command to package Service Component Architecture (SCA) compliant modules as Java applications that can be installed on a server. The command is useful when performing batch installs through `wsadmin`.

“Premigration considerations” on page 118

Consider these guidelines for the development of integration artifacts for WebSphere InterChange Server to ease the migration of WebSphere InterChange Server artifacts to WebSphere Process Server.

Related information

`Wsadmin` tool

WebSphere Integration Developer information center

IBM developerWorks

Migrating WebSphere InterChange Server artifacts to WebSphere Process Server artifacts, Part 1: Migrating collaboration templates to BPEL

Defining WebSphere variables

Configuring global security after WebSphere InterChange Server migration

Perform these additional security configuration steps to enable projects migrated from WebSphere InterChange Server to run successfully in a WebSphere Process Server environment.

Before you begin

You must first configure security for your WebSphere Process Server as described in *Securing applications and their environment*. In particular, make sure you have completed the steps in *Securing adapters* and *Creating end-to-end security*. In addition, install the EAR file for each module. Refer to *Deploying (installing) secure applications* for details.

About this task

After performing the above tasks, you are ready to complete the configuration steps, as follows:

- Binding the message-driven bean to activation specification
- Mapping the resource references to resources

- Mapping security roles to users or groups (required only when monitoring Common Based Events)
- Mapping RunAs roles (required only when monitoring Common Based Events)

Note: Mapping security roles to users or groups and mapping RunAs roles is possible from the administrative console only if the EJB deployment descriptors for the EJB projects have had a RunAs role defined. See Mapping users to RunAs roles using an assembly tool in the WebSphere Application Server Network Deployment, version 6.1 information center for information about defining RunAs roles with an assembly tool.

Procedure

1. Bind the message-driven bean to activation specification.
 - a. From the administrative console, select **Applications > Enterprise Applications**.
 - b. In the right panel, select the name of the application you just installed. (Select the name, not the check box to the left of the name.)
 - c. In the right panel again, under Enterprise Java Bean Properties, select **Message Driven Bean listener bindings**.
 - d. For each import or export EJB (indicated by an EJB name that starts with "_import" or "_export"), under the Bindings column, specify **SCA_Auth_Alias** in the "ActivationSpec authentication alias" field.
 - e. Select **OK**, then **Save**.
2. Map the resource references to resources.
 - a. From the administrative console, select **Applications > Enterprise Applications**.
 - b. In the right panel, select the name of the application you just installed. (Select the name, not the check box to the left of the name.)
 - c. In the right panel, under References, select **Resource references**.
 - d. In the Specify authentication method: field under javax.jms.ConnectionFactory, select the **Use default method (many to one mapping)** radio button.
 - e. In the **Select authentication data entry** pull down menu, select **SCA_Auth_Alias**.
 - f. Check the check box to select all of the modules.
 - g. Select **Apply**, then **OK**, then **Save**.
3. Map security roles to user groups.
 - a. From the administrative console, select **Applications > Enterprise Applications**.
 - b. In the right panel, select the name of the application you just installed. (Select the name, not the check box to the left of the name.)
 - c. In the right panel, under Detail Properties, select **Security role to user/group mapping**.
 - d. Select the check box to the left of the role you wish to map and then select **Look up users**.
 - e. Select **Search** to display a list of users who are available to map to the role, and move the correct user name to the "Selected:" column.
 - f. Select **OK**. The "Security role to user/group mapping" panel will be redisplayed.

Support for WebSphere Business Integration data handlers

The data handler support API enables certain data handler methods to be invoked from the AccessEJB, a WebSphere Process Server SCA Java component, or WebSphere Process Server bindings.

WebSphere Process Server (version 6.0.2.3 and higher) provides a data handler support Application Programming Interface (API) which enables select WebSphere Business Integration data handler methods to be invoked from the AccessEJB, a WebSphere Process Server SCA Java component, or WebSphere Process Server bindings. Access EJB has been replicated as an EJB that allows JService calls to route the input business object to the appropriate migrated module. The BPEL file in the migrated module will be invoked instead of the original WebSphere InterChange Server target collaboration.

WebSphere Process Server bindings invoke data bindings to perform data transformation. WebSphere Process Server provides several built-in data bindings as well as the capability to provide user-defined data bindings. You can implement a user-defined, or custom data binding to invoke a WebSphere Business Integration data handler.

By providing a custom data binding implementation, it is possible to leverage WebSphere Business Integration data handlers via the data handler support API. The data handler support API provides wrapper methods around existing WebSphere Business Integration data handler interface methods which perform the conversion between WebSphere Business Integration business objects and SDOs.

Data handler support API

By providing a custom data binding implementation, it is possible to leverage WebSphere Business Integration data handlers via the data handler support API. This API defines a set of public methods which can be invoked from a custom data binding or a Java component. It provides a way to invoke a text-based WebSphere Business Integration data handler from a process server binding. The following are the API methods:

```
getSDOFromString(String inputString, String sdoName, String  
metaObjectName, String mimeType) (Returns dataObject)
```

```
getStringFromSDO(DataObject sdo, String metaObjectName, String mimeType)  
(Returns String)
```

You can access these methods with the Java class `com.ibm.wbi.datahandler.JavaConnectorUtilDH`. This is the class that IBM recommends that you use from a data binding or a Java component. If you have existing code you can use the `AppSide_Connector.JavaConnectorUtil` class.

Usage

The methods defined in the data handler support API can be invoked from either a WebSphere Process Server binding or a Java component. However, because data is typically transformed in the binding in a WebSphere Process Server environment, IBM strongly recommends that the methods of the data handler support API be invoked from a custom data binding rather than a Java component.

Limitations

The data handler support API has the following limitations:

- Binary conversion methods are not supported. That is, no support is provided for `getBytesFromSDO()`, `getStreamFromSDO()`, `getSDO(byte[])`, and similar calls.
- `setEncoding()`, `setLocale()` and `setOptions()` methods are not exposed via the data handler support API.
- Dynamic child meta objects are not supported.
- You must use WebSphere Business Integration Adapter business object tooling for creation of new objects.

Related reference

“Supported WebSphere InterChange Server APIs”

In addition to the WebSphere InterChange Server source artifact migration tools provided in WebSphere Process Server and WebSphere Integration Developer, WebSphere Process Server also provides support for many of the APIs that were provided in WebSphere InterChange Server. The migration tools work in conjunction with these WebSphere InterChange Server APIs by preserving your custom snippet code as much as possible when migrating.

Related information

IBM WebSphere Business Integration Adapters/IBM WebSphere InterChange Server Data Handler Guide

Supported WebSphere InterChange Server APIs

In addition to the WebSphere InterChange Server source artifact migration tools provided in WebSphere Process Server and WebSphere Integration Developer, WebSphere Process Server also provides support for many of the APIs that were provided in WebSphere InterChange Server. The migration tools work in conjunction with these WebSphere InterChange Server APIs by preserving your custom snippet code as much as possible when migrating.

Note: These APIs are provided only to support migrated WebSphere InterChange Server applications until they can be modified to use new WebSphere Process Server APIs.

The supported WebSphere InterChange Server APIs are listed below. These APIs provide functions in WebSphere Process Server similar to the function that they provide in WebSphere InterChange Server. See the WebSphere InterChange Server v4.3 documentation for a functional description of these APIs.

CwBiDiEngine

AppSide_Connector/

- `BiDiBOTransformation(BusinessObject, String, String, boolean):BusinessObj`
- `BiDiBusObjTransformation(BusObj, String, String, boolean):BusObj`
- `BiDiStringTransformation(String, String, String):String`

JavaConnectorUtil

AppSide_Connector/

- `INFRASTRUCTURE_MESSAGE_FILE`
- `CONNECTOR_MESSAGE_FILE`
- `XRD_WARNING`
- `XRD_TRACE`

- XRD_INFO
- XRD_ERROR
- XRD_FATAL
- LEVEL1
- LEVEL2
- LEVEL3
- LEVEL4
- LEVEL5
- createBusinessObject(String):BusinessObjectInterface
- createBusinessObject(String, Locale):BusinessObjectInterface
- createBusinessObject(String, String):BusinessObjectInterface
- createContainer(String):CxObjectContainerInterface
- generateMsg(int, int, int, int, int, Vector):String
- generateMsg(int, int, int, int, Vector):String
- getBlankValue():String
- getEncoding():String
- getIgnoreValue():String
- getLocale():String
- getSDOFromString(String inputString, String sdoName, String metaObjectName, String mimeType)
- getStringFromSDO(DataObject sdo, String metaObjectName, String mimeType)
- isBlankValue(Object):boolean
- isIgnoreValue(Object):boolean
- isTraceEnabled(int):boolean
- logMsg(String)
- logMsg(String, int)
- traceWrite(int, String)

JavaConnectorUtilDH

datahandler/

wbi/

ibm/

com/

- getSDOFromString(String inputString, String sdoName, String metaObjectName, String mimeType)
- getStringFromSDO(DataObject sdo, String metaObjectName, String mimeType)

BusObj

Collaboration/

- BusObj(DataObject)
- BusObj(String)
- BusObj(String, Locale)
- copy(BusObj)
- duplicate():BusObj
- equalKeys(BusObj):boolean
- equals(Object):boolean
- equalsShallow(BusObj):boolean

- exists(String):boolean
- get(int):Object
- get(String):Object
- getBoolean(String):boolean
- getBusObj(String):BusObj
- getBusObjArray(String):BusObjArray
- getCount(String):int
- getDouble(String):double
- getFloat(String):float
- getInt(String):int
- getKeys():String
- getLocale():java.util.Locale
- getLong(String):long
- getLongText(String):String
- getString(String):String
- getType():String
- getValues():String
- getVerb():String
- isBlank(String):boolean
- isKey(String):boolean
- isNull(String):boolean
- isRequired(String):boolean
- keysToString():String
- set(BusObj)
- set(int, Object)
- set(String, boolean)
- set(String, double)
- set(String, float)
- set(String, int)
- set(String, long)
- set(String, Object)
- set(String, String)
- setContent(BusObj)
- setDefaultAttrValues()
- setKeys(BusObj)
- setLocale(java.util.Locale)
- setVerb(String)
- setVerbWithCreate(String, String)
- setWithCreate(String, boolean)
- setWithCreate(String, BusObj)
- setWithCreate(String, BusObjArray)
- setWithCreate(String, double)
- setWithCreate(String, float)
- setWithCreate(String, int)
- setWithCreate(String, long):

- setWithCreate(String, Object)
- setWithCreate(String, String)
- toString():String
- validData(String, boolean):boolean
- validData(String, BusObj):boolean
- validData(String, BusObjArray):boolean
- validData(String, double):boolean
- validData(String, float):boolean
- validData(String, int):boolean
- validData(String, long):boolean
- validData(String, Object):boolean
- validData(String, String):boolean

BusObjArray Collaboration/

- addElement(BusObj)
- duplicate():BusObjArray
- elementAt(int):BusObj
- equals(BusObjArray):boolean
- getElements():BusObj[]
- getLastIndex():int
- max(String):String
- maxBusObjArray(String):BusObjArray
- maxBusObjs(String):BusObj[]
- min(String):String
- minBusObjArray(String):BusObjArray
- minBusObjs(String):BusObj[]
- removeAllElements()
- removeElement(BusObj)
- removeElementAt(int)
- setElementAt(int, BusObj)
- size():int
- sum(String):double
- swap(int, int)
- toString():String

BaseDLM DLM/

- BaseDLM(BaseMap)
- getDBConnection(String):CwDBConnection
- getDBConnection(String, boolean):CwDBConnection
- getName():String
- getRelConnection(String):DtpConnection
- implicitDBTransactionBracketing():boolean
- isTraceEnabled(int):boolean
- logError(int)

- `logError(int, Object[])`
- `logError(int, String)`
- `logError(int, String, String)`
- `logError(int, String, String, String)`
- `logError(int, String, String, String, String)`
- `logError(int, String, String, String, String, String)`
- `logError(String)`
- `logInfo(int)`
- `logInfo(int, Object[])`
- `logInfo(int, String)`
- `logInfo(int, String, String)`
- `logInfo(int, String, String, String)`
- `logInfo(int, String, String, String, String)`
- `logInfo(int, String, String, String, String, String)`
- `logInfo(String)`
- `logWarning(int)`
- `logWarning(int, Object[])`
- `logWarning(int, String)`
- `logWarning(int, String, String)`
- `logWarning(int, String, String, String)`
- `logWarning(int, String, String, String, String)`
- `logWarning(int, String, String, String, String, String)`
- `logWarning(String)`
- `raiseException(RuntimeEntityException)`
- `raiseException(String, int)`
- `raiseException(String, int, Object[])`
- `raiseException(String, int, String)`
- `raiseException(String, int, String, String)`
- `raiseException(String, int, String, String, String)`
- `raiseException(String, int, String, String, String, String)`
- `raiseException(String, int, String, String, String, String, String)`
- `raiseException(String, String)`
- `releaseRelConnection(boolean)`
- `trace(int, int)`
- `trace(int, int, Object[])`
- `trace(int, int, String)`
- `trace(int, int, String, String)`
- `trace(int, int, String, String, String)`
- `trace(int, int, String, String, String, String)`
- `trace(int, int, String, String, String, String, String)`
- `trace(int, String)`
- `trace(String)`

CwDBConnection
CwDBConnection/
CxCommon/

- beginTransaction()
- commit()
- executePreparedSQL(String)
- executePreparedSQL(String, Vector)
- executeSQL(String)
- executeSQL(String, Vector)
- executeStoredProcedure(String, Vector)
- getUpdateCount():int
- hasMoreRows():boolean
- inTransaction():boolean
- isActive():boolean
- nextRow():Vector
- release()
- rollback()

CwDBConstants

CwDBConnection/

CxCommon/

- PARAM_IN - 0
- PARAM_INOUT - 1
- PARAM_OUT - 2

CwDBStoredProcedureParam

CwDBConnection/

CxCommon/

- CwDBStoredProcedureParam(int, Array)
- CwDBStoredProcedureParam(int, BigDecimal)
- CwDBStoredProcedureParam(int, boolean)
- CwDBStoredProcedureParam(int, Boolean)
- CwDBStoredProcedureParam(int, byte[])
- CwDBStoredProcedureParam(int, double)
- CwDBStoredProcedureParam(int, Double)
- CwDBStoredProcedureParam(int, float)
- CwDBStoredProcedureParam(int, Float)
- CwDBStoredProcedureParam(int, int)
- CwDBStoredProcedureParam(int, Integer)
- CwDBStoredProcedureParam(int, java.sql.Blob)
- CwDBStoredProcedureParam(int, java.sql.Clob)
- CwDBStoredProcedureParam(int, java.sql.Date)
- CwDBStoredProcedureParam(int, java.sql.Struct)
- CwDBStoredProcedureParam(int, java.sql.Time)
- CwDBStoredProcedureParam(int, java.sql.Timestamp)
- CwDBStoredProcedureParam(int, Long)
- CwDBStoredProcedureParam(int, String)
- CwDBStoredProcedureParam(int, String, Object)
- getParamType():int getValue():Object

DataHandler (Abstract Class)

**DataHandlers/
crossworlds/
com/**

- createHandler(String, String, String):DataHandler
- getBO(InputStream, Object):BusinessObjectInterface
- getBO(Object, BusinessObjectInterface, Object)
- getBO(Object, Object):BusinessObjectInterface
- getBO(Reader, BusinessObjectInterface, Object) (Abstract Method)
- getBO(Reader, Object):BusinessObjectInterface (Abstract Method)
- getBO(String, Object):BusinessObjectInterface
- getBOName(InputStream):String
- getBOName(Reader):String
- getBOName(String):String
- getBooleanOption(String):boolean
- getEncoding():String
- getLocale():Locale
- getOption(String):String
- getStreamFromBO(BusinessObjectInterface, Object):InputStream (Abstract Method)
- getStringFromBO(BusinessObjectInterface, Object):String (Abstract Method)
- setConfigMOName(String)
- setEncoding(String)
- setLocale(Locale)
- setOption(String, String)
- traceWrite(String, int)

NameHandler (Abstract Class)

**DataHandlers/
crossworlds/
com/**

- getBOName(Reader, String):String (Abstract Method)

ConfigurationException (extends java.lang.Exception)

**Exceptions/
DataHandlers/
crossworlds/
com/**

MalformedDataException (extends java.lang.Exception)

**Exceptions/
DataHandlers/
crossworlds/
com/**

NotImplementedException (extends java.lang.Exception)

**Exceptions/
DataHandlers/
crossworlds/
com/**

BusinessObjectInterface

CxCommon/

- clone():Object
- dump():String
- getAppText():String
- getAttrCount():int
- getAttrDesc(int):CxObjectAttr
- getAttrDesc(String):CxObjectAttr
- getAttribute(String):Object
- getAttributeIndex(String):int
- getAttributeType(int):int
- getAttributeType(String):int
- getAttrName(int):String
- getAttrValue(int):Object
- getAttrValue(String):Object
- getBusinessObjectVersion():String
- getDefaultAttrValue(int):String
- getDefaultAttrValue(String):String
- getLocale():String
- getName():String
- getParentBusinessObject():BusinessObjectInterface
- getVerb():String
- getVerbAppText(String):String
- isBlank(int):boolean
- isBlank(String):boolean
- isIgnore(int):boolean
- isIgnore(String):boolean
- isVerbSupported(String):boolean
- makeNewAttrObject(int):Object
- makeNewAttrObject(String):Object
- setAttributeWithCreate(String, Object)
- setAttrValue(int, Object)
- setAttrValue(String, Object)
- setDefaultAttrValues()
- setLocale(Locale)
- setLocale(String)
- setVerb(String)

CxObjectAttr

CxCommon/

- BOOLEAN
- BOOLSTRING
- DATE
- DATESTRING
- DOUBLE
- DOUBSTRING

- FLOAT
- FLTSTRING
- INTEGER
- INTSTRING
- INVALID_TYPE_NUM
- INVALID_TYPE_STRING
- LONGTEXT
- LONGTEXTSTRING
- MULTIPLECARDSTRING
- OBJECT
- SINGLECARDSTRING
- STRING
- STRSTRING
- equals(Object):boolean
- getAppText():String
- getCardinality():String
- getDefault():String
- getMaxLength():int
- getName():String
- getRelationType():String
- getTypeName():String
- getTypeNum():String
- hasCardinality(String):boolean
- hasName(String):boolean
- hasType(String):boolean
- isForeignKeyAttr():boolean
- isKeyAttr():boolean
- isMultipleCard():boolean
- isObjectType():boolean
- isRequiredAttr():boolean
- isType(Object):boolean

CxObjectContainerInterface

CxCommon/

- getBusinessObject(int):BusinessObjectInterface
- getObjectCount():int
- insertBusinessObject(BusinessObjectInterface)
- removeAllObjects()
- removeBusinessObjectAt(int)
- setBusinessObject(int, BusinessObjectInterface)

DtpConnection

Dtp/

CxCommon/

- beginTran()
- commit()

- executeSQL(String)
- executeSQL(String, Vector)
- executeStoredProcedure(String, Vector)
- getUpdateCount():int
- hasMoreRows():boolean
- inTransaction():boolean
- isActive():boolean
- nextRow():Vector
- rollback()

DtpDataConversion

Dtp/

CxCommon/

- BOOL_TYPE - 4
- CANNOTCONVERT - 2
- DATE_TYPE - 5
- DOUBLE_TYPE - 3
- FLOAT_TYPE - 2
- INTEGER_TYPE - 0
- LONGTEXT_TYPE - 6
- OKTOCONVERT - 0
- POTENTIALDATALOSS - 1
- STRING_TYPE - 1
- UNKNOWN_TYPE - 999
- getType(double):int
- getType(float):int
- getType(int):int
- getType(Object):int
- isOKToConvert(int, int):int
- isOKToConvert(String, String):int
- toBoolean(boolean):Boolean
- toBoolean(Object):Boolean
- toDouble(double):Double
- toDouble(float):Double
- toDouble(int):Double
- toDouble(Object):Double
- toFloat(double):Float
- toFloat(float):Float
- toFloat(int):Float
- toFloat(Object):Float
- toInteger(double):Integer
- toInteger(float):Integer
- toInteger(int):Integer
- toInteger(Object):Integer
- toPrimitiveBoolean(Object):boolean
- toPrimitiveDouble(float):double

- toPrimitiveDouble(int):double
- toPrimitiveDouble(Object):double
- toPrimitiveFloat(double):float
- toPrimitiveFloat(int):float
- toPrimitiveFloat(Object):float
- toPrimitiveInt(double):int
- toPrimitiveInt(float):int
- toPrimitiveInt(Object):int
- toString(double):String
- toString(float):String
- toString(int):String
- toString(Object):String

DtpDate

Dtp/

CxCommon/

- DtpDate()
- DtpDate(long, boolean)
- DtpDate(String, String)
- DtpDate(String, String, String[], String[])
- addDays(int):DtpDate
- addMonths(int):DtpDate
- addWeekdays(int):DtpDate
- addYears(int):DtpDate
- after(DtpDate):boolean
- before(DtpDate):boolean
- calcDays(DtpDate):int
- calcWeekdays(DtpDate):int
- get12MonthNames():String[]
- get12ShortMonthNames():String[]
- get7DayNames():String[]
- getCWDate():String
- getDayOfMonth():String
- getDayOfWeek():String
- getHours():String
- getIntDay():int
- getIntDayOfWeek():int
- getIntHours():int
- getIntMilliseconds():int
- getIntMinutes():int
- getIntMonth():int
- getIntSeconds():int
- getIntYear():int
- getMaxDate(BusObjArray, String, String):DtpDate
- getMaxDateBO(BusObj[], String, String):BusObj[]
- getMaxDateBO(BusObjArray, String, String):BusObj[]

- getMinDate(BusObjArray, String, String):DtpDate
- getMinDateBO(BusObj[], String, String):BusObj[]
- getMinDateBO(BusObjArray, String, String):BusObj[]
- getMinutes():String
- getMonth():String
- getMSSince1970():long
- getNumericMonth():String
- getSeconds():String
- getShortMonth():String
- getYear():String
- set12MonthNames(String[], boolean)
- set12MonthNamesToDefault()
- set12ShortMonthNames(String[])
- set12ShortMonthNamesToDefault()
- set7DayNames(String[])
- set7DayNamesToDefault()
- toString():String
- toString(String):String
- toString(String, boolean):String

DtpMapService

Dtp/

CxCommon/

- runMap(String, String, BusObj[], CxExecutionContext):BusObj[]

DtpSplitString

Dtp/

CxCommon/

- DtpSplitString(String, String)
- elementAt(int):String
- firstElement():String
- getElementCount():int
- getEnumeration():Enumeration
- lastElement():String
- nextElement():String
- prevElement():String
- reset()

DtpUtils

Dtp/

CxCommon/

- padLeft(String, char, int):String
- padRight(String, char, int):String
- stringReplace(String, String, String):String
- truncate(double):int
- truncate(double, int):double
- truncate(float):int

- truncate(float, int):double
- truncate(Object):int
- truncate(Object, int):double

BusObjInvalidVerbException (extends InterchangeExceptions)

Exceptions/

CxCommon/

- getFormattedMessage()

IdentityRelationship

relationship/

utilities/

crossworlds/

com/

- addMyChildren(String, String, BusObj, String, Object, CxExecutionContext)
- deleteMyChildren(String, String, BusObj, String, CxExecutionContext)
- deleteMyChildren(String, String, BusObj, String, Object, CxExecutionContext)
- foreignKeyLookup(String, String, BusObj, String, BusObj, String, CxExecutionContext)
- foreignKeyXref(String, String, String, BusObj, String, BusObj, String, CxExecutionContext)
- maintainChildVerb(String, String, String, BusObj, String, BusObj, String, CxExecutionContext, boolean, boolean)
- maintainCompositeRelationship(String, String, BusObj, Object, CxExecutionContext)
- maintainSimpleIdentityRelationship(String, String, BusObj, BusObj, CxExecutionContext)
- updateMyChildren(String, String, BusObj, String, String, String, String, CxExecutionContext)

MapExeContext

Dtp/

CxCommon/

- ACCESS_REQUEST - "SUBSCRIPTION_DELIVERY"
- ACCESS_RESPONSE - "ACCESS_RETURN_REQUEST"
- EVENT_DELIVERY - "SUBSCRIPTION_DELIVERY"
- SERVICE_CALL_FAILURE - "CONSUME_FAILED"
- SERVICE_CALL_REQUEST - "CONSUME"
- SERVICE_CALL_RESPONSE - "DELIVERBUSOBJ"
- getConnName():String
- getGenericBO():BusObj
- getInitiator():String
- getLocale():java.util.Locale
- getOriginalRequestBO():BusObj
- setConnName(String)
- setInitiator(String)
- setLocale(java.util.Locale)

**Participant
RelationshipServices/
Server/**

- Participant(String, String, int, BusObj)
- Participant(String, String, int, String)
- Participant(String, String, int, long)
- Participant(String, String, int, int)
- Participant(String, String, int, double)
- Participant(String, String, int, float)
- Participant(String, String, int, boolean)
- Participant(String, String, BusObj)
- Participant(String, String, String)
- Participant(String, String, long)
- Participant(String, String, int)
- Participant(String, String, double)
- Participant(String, String, float)
- Participant(String, String, boolean)
- getBoolean():boolean
- getBusObj():BusObj
- getDouble():double
- getFloat():float
- getInstanceId():int
- getInt():int
- getLong():long
- getParticipantDefinition():String
- getRelationshipDefinition():String
- getString():String INVALID_INSTANCE_ID
- set(boolean)
- set(BusObj)
- set(double)
- set(float)
- set(int)
- set(long)
- set(String)
- setInstanceId(int)
- setParticipantDefinition(String)
- setRelationshipDefinition(String)
- setParticipantDefinition(String)
- setRelationshipDefinition(String)

**Relationship
RelationshipServices/
Server/**

- addMyChildren(String, String, BusObj, String, Object, CxExecutionContext)
- addParticipant(Participant):int
- addParticipant(String, String, boolean):int

- addParticipant(String, String, BusObj):int
- addParticipant(String, String, double):int
- addParticipant(String, String, float):int
- addParticipant(String, String, int):int
- addParticipant(String, String, int, boolean):int
- addParticipant(String, String, int, BusObj):int
- addParticipant(String, String, int, double):int
- addParticipant(String, String, int, float):int
- addParticipant(String, String, int, int):int
- addParticipant(String, String, int, long):int
- addParticipant(String, String, int, String):int
- addParticipant(String, String, long):int
- addParticipant(String, String, String):int
- create(Participant):int
- create(String, String, boolean):int
- create(String, String, BusObj):int
- create(String, String, double):int
- create(String, String, float):int
- create(String, String, int):int
- create(String, String, long):int
- create(String, String, String):int
- deactivateParticipant(Participant)
- deactivateParticipant(String, String, boolean)
- deactivateParticipant(String, String, BusObj)
- deactivateParticipant(String, String, double)
- deactivateParticipant(String, String, float)
- deactivateParticipant(String, String, int)
- deactivateParticipant(String, String, long)
- deactivateParticipant(String, String, String)
- deactivateParticipantByInstance(String, String, int)
- deactivateParticipantByInstance(String, String, int, boolean)
- deactivateParticipantByInstance(String, String, int, BusObj)
- deactivateParticipantByInstance(String, String, int, double)
- deactivateParticipantByInstance(String, String, int, float)
- deactivateParticipantByInstance(String, String, int, int)
- deactivateParticipantByInstance(String, String, int, long)
- deactivateParticipantByInstance(String, String, int, String)
- deleteMyChildren(String, String, BusObj, String, CxExecutionContext)
- deleteMyChildren(String, String, BusObj, String, Object, CxExecutionContext)
- deleteParticipant(Participant)
- deleteParticipant(String, String, boolean)
- deleteParticipant(String, String, BusObj)
- deleteParticipant(String, String, double)
- deleteParticipant(String, String, float)
- deleteParticipant(String, String, int)

- deleteParticipant(String, String, long)
- deleteParticipant(String, String, String)
- deleteParticipantByInstance(String, String, int)
- deleteParticipantByInstance(String, String, int, boolean)
- deleteParticipantByInstance(String, String, int, BusObj)
- deleteParticipantByInstance(String, String, int, double)
- deleteParticipantByInstance(String, String, int, float)
- deleteParticipantByInstance(String, String, int, int)
- deleteParticipantByInstance(String, String, int, long)
- deleteParticipantByInstance(String, String, int, String)
- getNewID(String):int
- maintainCompositeRelationship(String, String, BusObj, Object, CxExecutionContext)
- maintainSimpleIdentityRelationship(String, String, BusObj, BusObj, CxExecutionContext)
- retrieveInstances(String, boolean):int[]
- retrieveInstances(String, BusObj):int[]
- retrieveInstances(String, double):int[]
- retrieveInstances(String, float):int[]
- retrieveInstances(String, int):int[]
- retrieveInstances(String, long):int[]
- retrieveInstances(String, String):int[]
- retrieveInstances(String, String, boolean):int[]
- retrieveInstances(String, String, BusObj):int[]
- retrieveInstances(String, String, double):int[]
- retrieveInstances(String, String, float):int[]
- retrieveInstances(String, String, int):int[]
- retrieveInstances(String, String, long):int[]
- retrieveInstances(String, String, String):int[]
- retrieveInstances(String, String[], boolean):int[]
- retrieveInstances(String, String[], BusObj):int[]
- retrieveInstances(String, String[], double):int[]
- retrieveInstances(String, String[], float):int[]
- retrieveInstances(String, String[], int):int[]
- retrieveInstances(String, String[], long):int[]
- retrieveInstances(String, String[], String):int[]
- retrieveParticipants(String):Participant[]
- retrieveParticipants(String, String):Participant[]
- retrieveParticipants(String, String[]):Participant[]
- retrieveParticipants(String, int):Participant[]
- retrieveParticipants(String, String, int):Participant[]
- retrieveParticipants(String, String[], int):Participant[]
- updateMyChildren(String, String, BusObj, String, String, String, String, CxExecutionContext)
- updateParticipant(String, String, BusObj)

- updateParticipantByInstance(Participant)
- updateParticipantByInstance(String, String, int)
- updateParticipantByInstance(String, String, int, BusObj)

UserStoredProcedureParam

Dtp/

CxCommon/

- UserStoredProcedureParam(int, String, Object, String, String)
- getParamDataTypeJavaObj():String
- getParamDataTypeJDBC():int
- getParamIndex():int
- getParamIOType():String
- getParamName():String
- getParamValue():Object
- setParamDataTypeJavaObj(String)
- setParamDataTypeJDBC(int)
- setParamIndex(int)
- setParamIOType(String)
- setParamName(String)
- setParamValue(Object)
- PARAM_TYPE_IN - "IN"
- PARAM_TYPE_OUT - "OUT"
- PARAM_TYPE_INOUT - "INOUT"
- DATA_TYPE_STRING - "String"
- DATA_TYPE_INTEGER - "Integer"
- DATA_TYPE_DOUBLE - "Double"
- DATA_TYPE_FLOAT - "Float"
- DATA_TYPE_BOOLEAN - "Boolean"
- DATA_TYPE_TIME - "java.sql.Time"
- DATA_TYPE_DATE - "java.sql.Date"
- DATA_TYPE_TIMESTAMP - "java.sql.Timestamp"
- DATA_TYPE_BIG_DECIMAL - "java.math.BigDecimal"
- DATA_TYPE_LONG_INTEGER - "Long"
- DATA_TYPE_BINARY - "byte[]"
- DATA_TYPE_CLOB - "Clob"
- DATA_TYPE_BLOB - "Blob"
- DATA_TYPE_ARRAY - "Array"
- DATA_TYPE_STRUCT - "Struct"
- DATA_TYPE_REF - "Ref"

BaseCollaboration

Collaboration/

- BaseCollaboration(com.ibm.bpe.api.ProcessInstanceData)
- AnyException - "AnyException"
- AppBusObjDoesNotExist - "BusObjDoesNotExist"
- AppLogOnFailure - "AppLogOnFailure"

- AppMultipleHits - "AppMultipleHits"
- AppRequestNotYetSent - "AppRequestNotYetSent"
- AppRetrieveByContentFailed - "AppRetrieveByContent"
- AppTimeOut - "AppTimeOut"
- AppUnknown - "AppUnknown"
- AttributeException - "AttributeException"
- existsConfigProperty(String):boolean
- getConfigProperty(String):String
- getConfigPropertyArray(String):String[]
- getCurrentLoopIndex():int
- getDBConnection(String):CwDBConnection
- getDBConnection(String, boolean):CwDBConnection getLocale():java.util.Locale
- getMessage(int):String
- getMessage(int, Object[]):String
- getName():String
- implicitDBTransactionBracketing():boolean
- isCallerInRole(String):boolean
- isTraceEnabled(int):boolean
- JavaException - "JavaException"
- logError(int)
- logError(int, Object[])
- logError(int, String)
- logError(int, String, String)
- logError(int, String, String, String)
- logError(int, String, String, String, String)
- logError(int, String, String, String, String, String)
- logError(String)
- logInfo(int)
- logInfo(int, Object[])
- logInfo(int, String)
- logInfo(int, String, String)
- logInfo(int, String, String, String)
- logInfo(int, String, String, String, String)
- logInfo(int, String, String, String, String, String)
- logInfo(String)
- logWarning(int)
- logWarning(int, Object[])
- logWarning(int, String)
- logWarning(int, String, String)
- logWarning(int, String, String, String)
- logWarning(int, String, String, String, String)
- logWarning(int, String, String, String, String, String)
- logWarning(String)
- not(boolean):boolean ObjectException - "ObjectException"
- OperationException - "OperationException"

- raiseException(CollaborationException)
- raiseException(String, int)
- raiseException(String, int, Object[])
- raiseException(String, int, String)
- raiseException(String, int, String, String)
- raiseException(String, int, String, String, String)
- raiseException(String, int, String, String, String, String)
- raiseException(String, int, String, String, String, String, String)
- raiseException(String, String)
- ServiceCallException - "ConsumerException"
- ServiceCallTransportException - "ServiceCallTransportException"
- SystemException - "SystemException"
- trace(int, int)
- trace(int, int, Object[])
- trace(int, int, String)
- trace(int, int, String, String)
- trace(int, int, String, String, String)
- trace(int, int, String, String, String, String)
- trace(int, int, String, String, String, String, String)
- trace(int, String)
- trace(String)
- TransactionException - "TransactionException"

CxExecutionContext

CxCommon/

- CxExecutionContext()
- getContext(String):Object
- MAPCONTEXT - "MAPCONTEXT"
- setContext(String, Object)

CollaborationException

Collaboration/

- getMessage():String
- getMsgNumber():int
- getSubType():String
- getText():String
- getType():String
- toString():String

Filter

crossworlds/

com/

- Filter(BaseCollaboration)
- filterExcludes(String, String):boolean
- filterIncludes(String, String):boolean
- recurseFilter(BusObj, String, boolean, String, String):boolean
- recursePreReqs(String, Vector):int

Globals

crossworlds/ com/

- Globals(BaseCollaboration)
- callMap(String, BusObj):BusObj

SmartCollabService

crossworlds/ com/

- SmartCollabService()
- SmartCollabService(BaseCollaboration)
- doAgg(BusObj, String, String, String):BusObj
- doMergeHash(Vector, String, String):Vector
- doRecursiveAgg(BusObj, String, String, String):BusObj
- doRecursiveSplit(BusObj, String):Vector
- doRecursiveSplit(BusObj, String, boolean):Vector
- getKeyValues(BusObj, String):String
- merge(Vector, String):BusObj
- merge(Vector, String, BusObj):BusObj
- split(BusObj, String):Vector

StateManagement

crossworlds/ com/

- StateManagement()
- beginTransaction()
- commit()
- deleteBO(String, String, String)
- deleteState(String, String, String, int)
- persistBO(String, String, String, String, BusObj)
- recoverBO(String, String, String):BusObj
- releaseDBConnection()
- resetData()
- retrieveState(String, String, String, int):int
- saveState(String, String, String, String, int, int, double)
- setDBConnection(CwDBConnection)
- updateBO(String, String, String, String, BusObj)
- updateState(String, String, String, String, int, int)

EventKeyAttrDef

EventManagement/ CxCommon/

- EventKeyAttrDef()
- EventKeyAttrDef(String, String)
- public String keyName
- public String keyValue

EventQueryDef
EventManagement/
CxCommon/

- EventQueryDef()
- EventQueryDef(String, String, String, String, int)
- public String nameConnector
- public String nameCollaboration
- public String nameBusObj
- public String verb
- public int ownerType

FailedEventInfo
EventManagement/
CxCommon/

- FailedEventInfo()
- FailedEventInfo(String x6, int, EventKeyAttrDef[], int, int, String, String, int)
- public String nameOwner
- public String nameConnector
- public String nameBusObj
- public String nameVerb
- public String strTime
- public String strMessage
- public int wipIndex
- public EventKeyAttrDef[] strbusObjKeys
- public int nKeys
- public int eventStatus
- public String expirationTime
- public String scenarioName
- public int scenarioState

Related tasks

“Enabling logging and tracing for supported WebSphere InterChange Server APIs” on page 165

Enable logging and tracing for the supported WebSphere InterChange Server APIs through the administrative console.

Related information

WebSphere InterChange Server v4.3 documentation

Limitations when migrating from WebSphere InterChange Server

Some characteristics of WebSphere InterChange Server are not precisely duplicated by WebSphere Process Server. Therefore you might have to modify your applications after migration to get them to perform as they did in WebSphere InterChange server.

The following sections describe these limitations and possible solutions.

Transaction levels

There is no direct mapping of the levels of transaction between WebSphere InterChange Server collaborations and WebSphere Process Server BPEL files. Therefore, the transaction level specified in the WebSphere InterChange Server collaboration is ignored and the default BPEL transaction level will be used in the migrated application. You should understand BPEL transactions and adjust your migrated applications accordingly to get the desired functionality.

Note: Pending transactions will not be migrated. All transactions should be concluded before starting migration.

Compensation

WebSphere Process Server compensation is different from WebSphere InterChange Server compensation. You should evaluate the new types of compensation offered by WebSphere Process Server and choose the type that best suits your application.

Event Summary and Change Summary not supported when using WebSphere InterChange Server APIs on WebSphere Process Server

Problem: Event Summary and Change Summary do not contain expected information in migrated WebSphere InterChange Server applications. **Cause:** Business Objects (BusObjs) in WebSphere InterChange Server do not support Change Summary and Event Summary. The WebSphere InterChange Server APIs supported in WebSphere Process Server work with the WebSphere InterChange Server type BusObj, so any use of those APIs forces a conversion into a BusObj. When this happens, any Event Summary and Change Summary information contained in a WebSphere Process Server DataObject that is converted into a BusObj is lost. Applications generated by migration from WebSphere InterChange Server will use the WebSphere InterChange Server APIs in WebSphere Process Server, so Event Summary and Change Summary cannot be used with these applications until the code is manually updated to stop using any of the WebSphere InterChange Server APIs. **Solution:** Remove all use of WebSphere InterChange Server APIs or change them to WebSphere Process Server APIs.



Related concepts

“Troubleshooting migration from WebSphere InterChange Server” on page 165
Find solutions to problems you encounter with migration as well as instructions for turning on logging and tracing.

Related reference

“Postmigration considerations” on page 128
When applications have been migrated from WebSphere InterChange Server to WebSphere Process Server, special attention is required in some areas to enable migrated applications to function in WebSphere Process Server consistently with their intended function due to differences between the architectures of WebSphere Process Server and WebSphere InterChange Server.

Related information

-  [Compensation in business processes](#)
-  [Compensation in business processes](#)

Troubleshooting migration from WebSphere InterChange Server

Find solutions to problems you encounter with migration as well as instructions for turning on logging and tracing.

Related concepts

“Limitations when migrating from WebSphere InterChange Server” on page 163
Some characteristics of WebSphere InterChange Server are not precisely duplicated by WebSphere Process Server. Therefore you might have to modify your applications after migration to get them to perform as they did in WebSphere InterChange server.

Related reference

“Postmigration considerations” on page 128

When applications have been migrated from WebSphere InterChange Server to WebSphere Process Server, special attention is required in some areas to enable migrated applications to function in WebSphere Process Server consistently with their intended function due to differences between the architectures of WebSphere Process Server and WebSphere InterChange Server.

“Premigration considerations” on page 118

Consider these guidelines for the development of integration artifacts for WebSphere InterChange Server to ease the migration of WebSphere InterChange Server artifacts to WebSphere Process Server.

Enabling logging and tracing for supported WebSphere InterChange Server APIs

Enable logging and tracing for the supported WebSphere InterChange Server APIs through the administrative console.

About this task

If your migrated application includes any supported WebSphere InterChange Server APIs, you can enable logging and tracing for them for troubleshooting purposes.

Procedure

1. Launch the administrative console.
2. From the left (navigation) panel, select **Troubleshooting > Logs and Trace**.
3. In the right panel, select the name of the server on which you want to enable logging and tracing.
4. In the right panel, under “General properties,” select **Change Log Level Details**.
5. Select the Runtime tab. (Selecting the Runtime tab allows you to make this change in real time without requiring you to restart the server.)
6. Add the name of the package followed by =all to the list of logged packages in the box on the screen. Separate this new entry from any existing entries with a colon. For example, CxCommon=all. In this case, CxCommon is the name of the package for a set of supported WebSphere InterChange Server APIs. Specifying all enables all logging and tracing. See Supported WebSphere InterChange Server APIs for a list of the APIs, including their package names.
7. Select **Apply**.
8. To keep this configuration after the server is restarted, select the **Save runtime changes to configuration as well** check box.
9. Select **OK**.

10. When the next screen appears, select **Save** to save your changes.

Related reference

 Supported WebSphere InterChange Server APIs

In addition to the WebSphere InterChange Server source artifact migration tools provided in WebSphere Process Server and WebSphere Integration Developer, WebSphere Process Server also provides support for many of the APIs that were provided in WebSphere InterChange Server. The migration tools work in conjunction with these WebSphere InterChange Server APIs by preserving your custom snippet code as much as possible when migrating.

Failure trying to serialize an object that is not serializable in a migrated BPEL file

If a serialization failure occurs in a BPEL file generated by the migration, you might be able to modify it to prevent the failure from occurring.

Problem: A serialization failure occurs in a custom snippet node of a business process execution language (BPEL) file generated by migration because an attempt is made to serialize an object that is not serializable.

Cause: In WebSphere InterChange Server, a Collaboration Template is compiled into a single Java class. In WebSphere Process Server, each node in a BPEL file might compile into a separate Java class. In WebSphere InterChange Server, a variable can be declared once and shared throughout the various steps of a Collaboration Template. To simulate that behavior in the migrated BPEL file, each variable used in a code snippet must be retrieved at the start of the snippet and saved at the end of the snippet. Variables defined in WebSphere InterChange Server Port definitions become BPEL variables. These are retrieved into BusObj variables at the beginning of each snippet (if referenced in the snippet) and saved back to the BPEL variables at the end of each snippet. For example, a retrieval at the beginning of snippets looks like this:

```
BusObj tempBusObj = null;if (tempBusObj_var != null) { tempBusObj =
    new BusObj(tempBusObj_var); };
```

and a save at the end of snippets looks like this:

```
if (tempBusObj == null) { tempBusObj_var = null; } else { tempBusObj_var =
    tempBusObj.getBusinessGraph(); }
```

Other variables used in the WebSphere InterChange Server snippet code are serialized and stored as a String in a BPEL variable named *CollabTemplateName_var*. These variables are deserialized at the beginning of each BPEL snippet, and then serialized and saved at the end of each BPEL Snippet that they are referenced in. For example, objects are retrieved like this:

```
BusObj tempBusObj = (BusObj)BaseCollaboration.deserialize
    (FrontEndCollab_var.getString("tempBusObj"));
```

and objects are saved like this:

```
FrontEndCollab_var.setString("tempBusObj", BaseCollaboration.serialize(tempBusObj));
```

If the type of the object being serialized is not serializable, then using `serialize` and `deserialize` will fail when the BPEL is run.

Solution: After migration, modify the BPEL file as follows:

- For any variable that is not Java-serializable, update the BPEL snippets to remove the serialization and deserialization statements. If the variable needs to

be shared across snippets (instead of being re-created in each snippet) another method must be used to preserve the value of the variable across snippets.

- Manually define BPEL variables for variables of type BusObj that are not declared in the WebSphere InterChange Server Port definitions but are used on Partner Invokes. This is a manual step because variables used during invokes in WebSphere Process Server must be strongly typed, and the migration tools cannot accurately determine that type from the WebSphere InterChange Server snippets.

Note: The naming convention used by the migration tools is to add `_var` to the name of the variable in the snippet code when naming the BPEL variables. For example, for a variable called `tempBusObj` in the snippet code, the migration tools will create a BPEL variable named `tempBusObj_var`.

- For variables that must be declared manually as BPEL variables, change the BPEL snippet code so that it uses the "deserialize/serialize" method of preserving these variables rather than the "retrieve from/store into BPEL variable" method of preserving these variables.

Migrating source artifacts from WebSphere Studio Application Developer Integration Edition

To migrate source artifacts from WebSphere Studio Application Developer Integration Edition use the tools available in WebSphere Integration Developer.

About this task

Use the migration wizard or command line available from WebSphere Integration Developer to migrate WebSphere Application Server Developer Integration Edition service projects into projects in the active WebSphere Integration Developer workspace. Refer to the WebSphere Integration Developer information center for more information.

Related information

WebSphere Integration Developer information center

Migrating from WebSphere MQ Workflow

To migrate from WebSphere MQ Workflow, use the WebSphere Integration Developer migration wizard or a special utility to migrate from WebSphere MQ Workflow 3.6 to WebSphere Process Server.

About this task

For this version of WebSphere MQ Workflow...	Do this
WebSphere MQ Workflow 3.6	Use either the WebSphere Integration Developer migration wizard or the FDL2BPEL utility to migrate all WebSphere MQ Workflow artifacts into WebSphere Integration Developer deployable artifacts.
WebSphere MQ Workflow 3.5 or earlier	You must first migrate to WebSphere MQ Workflow version 3.6.

See the WebSphere Integration Developer information center for more information.

Related information

[WebSphere Integration Developer information center](#)

Chapter 3. Deprecated features

This section summarizes deprecated features in the product offerings comprising WebSphere Process Server version 6.0 and 6.1, and WebSphere Business Integration Server Foundation version 5.1. Deprecated features from other WebSphere Application Server version 5.1 and 6.x product offerings are described in the documentation for those products.

Deprecation list

This topic describes the deprecated features in the following versions and releases:

- “Deprecated features in WebSphere Process server version 6.1”
- “Deprecated features in WebSphere Process Server version 6.0.2” on page 173
- “Deprecated features in WebSphere Process Server version 6.0.1” on page 175
- “Deprecated features in WebSphere Process Server version 6.0” on page 175
- “Deprecated features in WebSphere Business Integration Server Foundation version 5.1.1” on page 178
- “Deprecated features in WebSphere Business Integration Server Foundation version 5.1” on page 178

The following tables summarize what is deprecated, by version and release. Each table reflects the version and release where the deprecation took effect and lists what is being deprecated, such as features, APIs, scripting interfaces, tools, wizards, publicly exposed configuration data, naming identifiers, and constants. Where possible, a recommended migration action is provided.

Deprecated features in WebSphere Process server version 6.1

Container Manager Persistence over Anything (CMP/A)
<p>The CMP/A support included with WebSphere Process Server is deprecated. This includes the runtime support for applications which have been customized to use CMP/A, the <code>cmpdeploy.bat/.sh</code> command line tool, and the following public APIs:</p> <ul style="list-style-type: none">• <code>com.ibm.websphere.rsadapter.WSProceduralPushDownHelper</code>• <code>com.ibm.websphere.rsadapter.WSPushDownHelper</code>• <code>com.ibm.websphere.rsadapter.WSPushDownHelperFactory</code>• <code>com.ibm.websphere.rsadapter.WSRelationalPushDownHelper</code>
<p>Recommended migration action:</p> <p>Convert CMP Entity Beans to use a relational data source, or have the CMP entity bean replaced by a different supported data persistence model.</p> <p>You can also use WebSphere Adapters to replace your existing CMP/A applications. The Adapter tools use a ‘Create, Retrieve, Update, and Delete’ architecture for creating service interfaces that is very similar to the architecture that CMP/A uses.</p>
JACL scripts (deprecated in WebSphere Application Server version 6.1)

Container Manager Persistence over Anything (CMP/A)

JACL script files are deprecated in WebSphere Process Server to maintain consistency with the deprecation of JACL scripts in WebSphere Application Server.

Recommended migration action:

Use the corresponding .bat/.sh files or wsadmin commands to perform the same functions.

Note: The following Business Process Choreographer JACL scripts are not deprecated:

1. `<install_root>\ProcessChoreographer\admin\bpcTemplates.jacl`
2. `<install_root>\ProcessChoreographer\config\bpeconfig.jacl`
3. `<install_root>\ProcessChoreographer\config\bpeunconfig.jacl`
4. `<install_root>\ProcessChoreographer\config\bpeupgrade.jacl`
5. `<install_root>\ProcessChoreographer\config\clientconfig.jacl`

IBM Web Services Client for C++

The IBM Web Services Client for C++ is a standalone application with its own installer, but which is distributed on the WebSphere Process Server media. The product does not use or have a dependency on this software, however the IBM Message Service Client for C/C++ which is also distributed with the product does.

Recommended migration action:

Use one of the other freely available tools, such as gSOAP (<http://www.cs.fsu.edu/~engelen/soap.html>) which is an open source product distributed under the GPL license, which will provide the same functions.

Business Process Choreographer

Generic Business Process EJB API

- The `getAutoDelete()` function from `ProcessTemplateData` is deprecated.

Recommended migration action:

Use method `getAutoDeletionMode()` to query how auto deletion is handled for the corresponding process template.

- The exception `SpecificFaultReplyException` is deprecated.

Recommended migration action:

No action is required. This exception is only needed to handle WSIF messages, which are no longer supported.

Generic Business Process Webservice API - XML schema types

Element `autoDelete` of the complex type `ProcessTemplateType` is deprecated

```
<xsd:element name="ProcessTemplate" type="tns:ProcessTemplateType"/>
<xsd:complexType name="ProcessTemplateType">
  <xsd:sequence>
    ...
    <xsd:element name="autoDelete" type="xsd:boolean" minOccurs="0"/>
  ...</xsd:sequence></xsd:complexType>
```

Recommended migration action:

Use element `autoDeletionMode` of type `ProcessTemplateType`

```
<xsd:element name="ProcessTemplate" type="tns:ProcessTemplateType"/>
<xsd:complexType name="ProcessTemplateType">
  <xsd:sequence>
    ...
    <xsd:element name="autoDeletionMode" type="xsd:string" minOccurs="0"/>
  ...</xsd:sequence></xsd:complexType>
```


Container Manager Persistence over Anything (CMP/A)

Deprecation of Observer DB Cleanup Methods of the ProcessContainer MBean

The following methods are deprecated:

- `public String observerForceRemoveInstanceData(String dataSourceName, String state, String templateName, String validFrom, String completedBefore)`
- `public String observerRemoveDeletedInstancesData(String dataSourceName, String completedBefore)`
- `public String observerRemoveInstanceDataOfTemplate(String dataSourceName, String templateName, String validFrom)`

Recommended migration action:

Use the following new methods (with the same name and an additional Parameter 'dbSchemaName')

- `public String observerForceRemoveInstanceData(String dataSourceName, String dbSchemaName, String state, String templateName, String validFrom, String completedBefore)`
- `public String observerRemoveDeletedInstancesData(String dataSourceName, String dbSchemaName, String completedBefore)`
- `public String observerRemoveInstanceDataOfTemplate(String dataSourceName, String dbSchemaName, String templateName, String validFrom)`

LDAP staff resolution plug-in

The attribute evaluation specification for staff queries of the LDAP staff resolution plug-in is deprecated:

```
<ldap:attribute name="attribute name"
                objectclass="LDAP object class"
                usage="simple">
</ldap:attribute>
```

Recommended migration action:

Use the result object evaluation specification supporting multiple attributes per LDAP object. The attributes "objectclass" and "attribute" of the "user" query will be replaced by a full result object evaluation specification that supports multiple result attributes per person.

Generic Human Task Manager EJB API

- The following fields from interface Task are deprecated:
 - STATE_FAILING
 - STATE_SKIPPED
 - STATE_STOPPED
 - STATE_TERMINATING
 - STATE_WAITING
 - STATE_PROCESSING_UNDO

Recommended migration action:

Use retrieve the staff activity associated with the inline human task for inline human tasks, and check the activity state using the `getExecutionState()` method on the `ActivityInstanceData` interface in the Generic Business Process EJB API.

- The field `KIND_WPC_STAFF_ACTIVITY` from interface Task is deprecated.

Recommended migration action:

Use the method `isInline()` on the Task interface to determine whether a human task is associated with a human task (staff) activity in a business process,

Container Manager Persistence over Anything (CMP/A)

Deprecation of e-mail people assignment criteria

The e-mail receiver people assignment criteria (staff verbs) used for escalations with escalation action "e-mail" are deprecated, as they are not required anymore in version 6.1. This applies to the following people assignment criteria:

- Email Address for Department Members
- Email Address for Group Members
- Email Address for Group Members without Filtered Users
- Email Address for Group Search
- Email Address for Role Members
- Email Address for Users
- Email Address for Users by user ID

Recommended migration action:

E-mail addresses and the preferred language are resolved together with the user ID by the standard set of people assignment criteria with version 6.1. This deprecation information is thus especially important for those who write custom XSLT people assignment criteria mapping (staff verb) files. If you do not intend to deploy version 6.0.2 task definitions, you do not need to support the deprecated people assignment criteria. Note that with version 6.1, the people assignment criteria "User Records by user ID" has been introduced and has to be supported by custom XSLT files, since it resolves e-mail addresses as fallback.

You can eliminate the deprecated e-mail people assignment criteria in your existing human task definitions by initiating source artifact migration in WebSphere Integration Developer 6.1. To do this, import your version 6.0.2 task definition into WebSphere Integration Developer 6.1, make a minor change (like adding a blank to the task description and deleting it again), and then save it again.

Deprecation of MQ as JMS Provider for BPC-internal messaging (Configuration of business process container and human task container)

Configuring the business process container and human task container to use MQSeries® as the JMS provider is deprecated. The business process container and human task container use JMS for their internal messaging — specifically, to navigate long-running process instances.

Recommended migration action:

During configuration of the business process container and human task container, use the default JMS messaging provider.

Business Objects

The following Business Object methods are deprecated:

- `com.ibm.websphere.bo.BOFactory.createClass(java.lang.Class interfaceClass);`
- `com.ibm.websphere.bo.BOType.getTypeByClass(java.lang.Class className);`

Recommended migration action:

These methods will raise "function not supported" exceptions if they are called in version 6.1

Common Event Infrastructure

Creation and editing of user-visible Common Base Events are deprecated.

Recommended migration action:

You can now use the tools to specify the Business Object data that is to be included in monitored emitted events

zOS

Container Manager Persistence over Anything (CMP/A)
<p>The requirement to bind a String object into JNDI at esb/messageLogger/qualifier is deprecated.</p> <p>Recommended migration action:</p> <p>The Message Logger primitives will now store message information within the CommonDB database. Where necessary, during the profile augmentation phase, a WebSphere variable called ESB_MESSAGE_LOGGER_QUALIFIER will now be created and its value set to that of the chosen CommonDB schema qualifier.</p>
WebSphere InterChange Server
<p>The APIs (application programming interfaces) listed in Supported WebSphere InterChange Server APIs are no longer deprecated.</p> <p>Note: These APIs were formerly deprecated in WebSphere Process Server version 6.0.2.</p> <p>Recommended migration action:</p> <p>You should use these APIs only for applications with migrated WebSphere InterChange Server components. In all other cases, you should use the Service Data Objects for WebSphere Process Server.</p>
WebSphere Enterprise Service Bus (WESB)
<p>The current method to identify an SSL repertoire to be used when WESB communicates with a secured WSRR instance has been deprecated.</p> <p>Recommended migration action:</p> <p>A new property has been added to WSRR definitions to allow the specification of such a repertoire.</p>

Deprecated features in WebSphere Process Server version 6.0.2

Human Task Manager
<p>The task context variable %htm:task.clientDetailURL% is no longer required, and thus has been deprecated.</p> <p>Recommended migration action:</p> <p>No action is required.</p>
<p>The standard e-mail implementation used for all escalation e-mails in TEL has been deprecated, and replaced by native support for defining e-mails in TEL.</p> <p>Recommended migration action:</p> <p>Use the customizable e-mail feature for escalations.</p>
<p>The following Task object methods that were deprecated in version 6.0 are no longer deprecated:</p> <pre> getInputMessageType() getOutputMessageType() </pre> <p>Recommended migration action:</p> <p>You can now use these methods.</p>
Business Process Choreographer

The method `getProcessAdministrators()` in the Generic Business Process EJB API interfaces `ActivityInstanceData`, `ProcessInstanceData`, and `ProcessTemplateData` are deprecated:

Recommended migration action:

Use these corresponding methods:

- `getProcessAdminTaskID()` in combination with method `getUsersInRole()` of the `HumanTaskManagerService` interface, as follows:
`htm.getUsersInRole(actInstData.getProcessAdminTaskID(), WorkItem.REASON_ADMINISTRATOR)`
- `getAdminTaskID()` in combination with method `getUsersInRole()` of the `HumanTaskManagerService` interface, as follows:
`htm.getUsersInRole(procInstData.getAdminTaskID(), WorkItem.REASON_ADMINISTRATOR)`
- `getAdminTaskTemplateID()` in combination with method `getUsersInRole()` of the `HumanTaskManagerService` interface, as follows:
`htm.getUsersInRole(procTemplData.getAdminTaskTemplateID(), WorkItem.REASON_ADMINISTRATOR)`

The following methods are deprecated for the `BusinessFlowManagerService` interface in Generic Business Process EJB API and the `HumanTaskManagerService` interface in the Generic Task EJB API:

- `query(String storedQueryName, Integer skipTuples)`
- `query(String storedQueryName, Integer skipTuples, Integer threshold)`

Recommended migration action:

Use these corresponding methods:

- `query(String storedQueryName, Integer skipTuples, List parameters)`
- `query(String storedQueryName, Integer skipTuples, Integer threshold, List parameters)`

The following JACL scripts are deprecated:

- `deleteAuditLog.jacl`
- `deleteInvalidProcessTemplate.jacl`
- `deleteInvalidTaskTemplate.jacl`
- `queryNumberOfFailedMessages.jacl`
- `replayFailedMessages.jacl`
- `cleanupUnusedStaffQueryInstances.jacl`
- `refreshStaffQuery.jacl`

Recommended migration action:

For each deprecated JACL scripts, a corresponding Jython script is now provided. Use the Jython scripts (*.py), which can be found in the `<install_root>/ProcessChoreographer/admin` directory.

SCA Admin Commands

The following commands (used with `wsadmin`) are deprecated:

- `configSCAForServer`
- `configSCAForCluster`

Recommended migration action:

Use these two commands in place of `configSCAForServer` for equivalent function. :

- `configSCAAsyncForServer`
- [Optional; use only if required] `configSCAJMSForServer`

Use these two commands in place of `configSCAForCluster` for equivalent function:

- `configSCAAsyncForCluster`
- [Optional; use only if required] `configSCAJMSForCluster`

WebSphere InterChange Server

Note: These APIs are no longer deprecated in version 6.1.

The APIs (application programming interfaces) listed in Supported WebSphere InterChange Server APIs are deprecated.

Recommended migration action:

Code written for WebSphere Process Server should not use these interfaces.

IBM WebSphere InterChange Server Access for Enterprise JavaBeans (EJB) support is deprecated.

Recommended migration action:

Applications developed for use with WebSphere Process Server should not use Access for Enterprise JavaBeans.

Deprecated features in WebSphere Process Server version 6.0.1

WebSphere Process Server version 6.0.1 has no deprecated features.

Deprecated features in WebSphere Process Server version 6.0

Application programming model and container support features

The BRBeans component is deprecated, and is being replaced with business rules.

Recommended migration action:

You must manually remove all usages of BRBeans and move to business rules.

Some BPEL business process modeling constructs have been syntactically changed in version 6. Only the syntax is supported by WebSphere Integration Developer version 6.0. Migration is available for these constructs.

Recommended migration action:

Use the migration wizard provided by WebSphere Integration Developer to migrate WebSphere Business Integration Server Foundation version 5.1 service projects (including process definitions) to WebSphere Process Server version 6.0. After the migration wizard has finished, you must carry out some manual steps to complete the migration. For more information about migrating service projects, see the information center for WebSphere Integration Developer version 6.0.

In WebSphere Business Integration Server Foundation version 5.1, there is an option for the input of an undo service to implicitly provide a message that results from the merge of the input data of the compensable service overlaid by its output data. Given the enhanced compensation support provided by BPEL this functionality is deprecated.

Recommended migration action:

Use BPEL compensation for business processes.

Because of changes in the Business Flow Manager functionality In WebSphere Process Server version 6.0, the following methods are deprecated in the generic process API:

- The WorkList object has been renamed to StoredQuery; consequently, the following methods are deprecated on the BusinessFlowManager bean, and, if applicable, the methods you would use WebSphere Process Server version 6.0 are given:
 - newWorkList(String workListName, String selectClause, String whereClause, String orderByClause, Integer threshold, TimeZone timezone)
Replace with: createStoredQuery(String storedQueryName, String selectClause, String whereClause, String orderByClause, Integer threshold, TimeZone timezone)
 - getWorkListNames()
Replace with: getStoredQueryNames()
 - deleteWorkList(String workListName)
Replace with: deleteStoredQuery(String storedQueryName)
 - getWorkList(String workListName)
Replace with: getStoredQuery(String storedQueryName)
 - executeWorkList(String workListName)
Replace with: query(String storedQueryName, Integer skipTuples)
 - getWorkListActions()
not supported.
- The WorkListData object is deprecated.
Use StoredQueryData instead.
- The following methods of the ProcessTemplateData object are no longer supported:
 - getInputMessageTypeSystemName()
 - getOutputMessageTypeSystemName()
- The following methods of the ProcessInstanceData object are no longer supported:
 - getInputMessageTypeSystemName()
 - getOutputMessageTypeSystemName()
- The following methods of the ActivityInstanceData object are no longer supported:
 - getInputMessageTypeSystemName()
 - getOutputMessageTypeSystemName()
- The following methods of the ActivityServiceTemplateData object are no longer supported:
 - getInputMessageTypeSystemName()

Recommended migration action:

Use the replacement methods, if any, that are given.

Because of changes in the Human Task Manager functionality In WebSphere Process Server version 6.0, the following methods are deprecated in the generic process API

- The following methods are deprecated on the HumanTaskManager bean, and their replacements for use in WebSphere Process Server version 6.0 are given:
 - createMessage(TKIID tkiid, String messageTypeName)
Use the specific methods createInputMessage(TKIID tkiid), createOutputMessage(TKIID tkiid), createFaultMessage(TKIID tkiid) instead.
 - createMessage(String tkiid, String messageTypeName)
Use the specific methods createInputMessage(String tkiid), createOutputMessage(String tkiid), createFaultMessage(String tkiid) instead.
- For the Task object, the following methods are no longer supported:
 - getInputMessageTypeNames()
 - getOutputMessageTypeNames()

Recommended migration action:

Use the replacement methods, if any, that are given.

The following database views are deprecated:

- DESCRIPTION
- CUSTOM_PROPERTY

Recommended migration action:

Use the TASK_DESC view for the DESCRIPTION view and the TASK_CPROP view for the CUSTOM_PROPERTY view.

Programming Model of Java Code Snippets

- In WebSphere Business Integration Server Foundation version 5.1, access to BPEL variables within inline Java code snippets (activities and conditions) is provided through getter and setter methods. These methods are not supported. The WSIFMessage method that is used to represent BPEL variables in Java code snippets is also not supported.
- Methods <typeOfP> getCorrelationSet<cs> Property<p>() are not supported, because they do not consider correlation sets declared at the scope level; they can only be used to access correlation sets declared at the process level.
- The WebSphere Business Integration Server Foundation version 5.1 methods to access custom properties within Java snippet activities are not supported.
- The following getPartnerLink methods are not supported. Because they do not consider partner links declared on the scope level, they can only be used to access partner links declared at the process level.
 - EndpointReference getPartnerLink();
 - EndpointReference getPartnerLink (int role);
 - void setPartnerLink (EndpointReference epr);

Recommended migration action:

Use the migration wizard provided by WebSphere Integration Developer 6.0 to migrate WebSphere Business Integration Server Foundation version 5.1 service projects (including process definitions) to WebSphere Process Server version 6.0. After the migration wizard has finished, you must carry out some manual steps to complete the migration. For more information about migrating service projects, see the information center for WebSphere Integration Developer version 6.0.

Application services features

The Extended Messaging Service feature and all of the EMS/CMM APIs and SPIs are deprecated:

com/ibm/websphere/ems/CMMCorrelator
com/ibm/websphere/ems/CMMException
com/ibm/websphere/ems/CMMReplyCorrelator
com/ibm/websphere/ems/CMMRequest
com/ibm/websphere/ems/CMMResponseCorrelator
com/ibm/websphere/ems/ConfigurationException
com/ibm/websphere/ems/FormatException
com/ibm/websphere/ems/IllegalStateException
com/ibm/websphere/ems/InputPort
com/ibm/websphere/ems/OutputPort
com/ibm/websphere/ems/transport/jms/JMSRequest
com/ibm/websphere/ems/TimeoutException
com/ibm/websphere/ems/TransportException
com/ibm/ws/spi/ems/CMMFactory
com/ibm/ws/spi/ems/format/cmm/CMMFormatter
com/ibm/ws/spi/ems/format/cmm/CMMParser
com/ibm/ws/spi/ems/format/Formatter
com/ibm/ws/spi/ems/format/Parser
com/ibm/ws/spi/ems/transport/CMMReceiver
com/ibm/ws/spi/ems/transport/CMMReplySender
com/ibm/ws/spi/ems/transport/CMMSender
com/ibm/ws/spi/ems/transport/MessageFactory

Recommended migration action:

Instead of using the Extended Messaging Service and its associated tools, you will need to use standard JMS APIs, or equivalent messaging technologies.

Deprecated features in WebSphere Business Integration Server Foundation version 5.1.1

WebSphere Business Integration Server Foundation version 5.1.1 has no deprecated features.

Deprecated features in WebSphere Business Integration Server Foundation version 5.1

Installation and migration tools

Business processes modeled with WebSphere Studio Application Developer Integration Edition version 5.0 or earlier are deprecated.

Recommended migration action:

Use the **Migrate** option provided with WebSphere Studio Application Developer Integration Edition version 5.1 to migrate business process to a BPEL-related process.

Several Business Process Choreographer API interfaces and methods used for business processes created with WebSphere Studio Application Developer Integration Edition version 5.0 or earlier.

Recommended migration action:

See the Javadoc provided with Business Process Choreographer for a detailed list of these API interfaces and methods.

Application programming model and container support features

Business Rule Bean programming interfaces that include the following public classes, methods, and attributes are deprecated:

- Public classes:
 - com.ibm.websphere.brb.RuleImporter
 - com.ibm.websphere.brb.RuleExporter
- Public method:
 - getLocalRuleManager() on class com.ibm.websphere.brb.TriggerPoint
- Protected attribute:
 - ruleMgr on class com.ibm.websphere.brb.TriggerPoint

Recommended migration action:

No action is required.

The com.ibm.websphere.scheduler class programming interface scheduler.Scheduler methods are deprecated:

```
public BeanTaskInfo createBeanTaskInfo();  
public MessageTaskInfo createMessageTaskInfo();
```

Recommended migration action:

Use the following methods:

```
public Object createTaskInfo(Class taskInfoInterface) throws TaskInfoInvalid;  
BeanTaskInfo ti = (BeanTaskInfo) Scheduler.createTaskInfo(BeanTaskInfo.class);
```

The Web Services gateway customization API is deprecated.

Recommended migration action:

No action is required. However, where possible, use the Java API for XML-based Remote Procedure Call (JAX-RPC) handlers rather than Web Services gateway-specific interfaces, such as filters. The Web Services gateway API will be replaced in a future release. For more information, see the topic 'JAX-RPC handlers - An alternative to gateway filters' in the WebSphere Business Integration Server Foundation information center.

Chapter 4. Troubleshooting migration

If you encounter problems during migration, the information described here could help.

Troubleshooting version-to-version migration

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

- While you are using the version 6.1 migration wizard to create a profile before migrating a configuration, you might see the following profile-creation error messages.

```
profileName: profileName cannot be empty
profilePath: Insufficient disk space
```

These error messages might be displayed if you enter a profile name that contains an incorrect character such as a space. Rerun the migration wizard, and verify that there are no incorrect characters in the profile name such as a space, quotes, or any other special characters.

- If you encounter a problem when you are migrating from a previous version of WebSphere Process Server to version 6.1, check your log files and other available information.
 1. Look for the log files, and browse them for clues.
 - *migration_backup_directory*/WBIPreUpgrade.time_stamp.log
 - *profile_root*/log/WASPostUpgrade.time_stamp.log
 - *install_root*/logs/clientupgrade.time_stamp.log
 - *profile_root*/logs/bpeupgrade.log
 - *migration_backup_directory*/WBIProfileUpgrade.ant.timestamp.log
 2. Look for MIGR0259I: The migration has successfully completed. or MIGR0271W: The migration completed with warnings. in the following directories:
 - *migration_backup_directory*/WBIPreUpgrade.time_stamp.log
 - *profile_root*/logs/WASPostUpgrade.time_stamp.log
 - *install_root*/logs/clientupgrade.time_stamp.logIf MIGR0286E: The migration failed to complete. is displayed, attempt to correct any problems based on the error messages that appear in the log file. After correcting any errors, rerun the command from the bin directory of the product installation root.
 3. Open the Log and Trace Analyzer built into the Application Server Toolkit (AST) on the service log of the server that is hosting the resource that you are trying to access, and use it to browse error and warning messages.
See Debugging components in the Application Server Toolkit.
 4. With WebSphere Process Server, run the dumpNameSpace command and pipe, redirect, or "more" the output so that it can be easily viewed.
This command results in a display of all objects in WebSphere Process Server namespace, including the directory path and object name.
 5. If the object a client needs to access does not appear, use the administrative console to verify the following conditions.
 - The server hosting the target resource is started.

- The Web module or Enterprise JavaBean container hosting the target resource is running.
- The JNDI name of the target resource is properly specified.

If none of these steps solves the problem, see Troubleshooting and support for additional troubleshooting resources, including information about how to contact IBM Support.

- During the migration process, problems might occur while you are using the `WBIPreUpgrade` command or the `WBIPostUpgrade` command.
 - Problems can occur when you are using the `WBIPreUpgrade` command.
 - A "Not found" or "No such file or directory" message is returned.
This problem can occur if you are trying to run the `WBIPreUpgrade` command from a directory other than the WebSphere Process Server version 6.1 `install_root/bin` directory. Verify that the `WBIPreUpgrade` script resides in the version 6.1 `install_root/bin` directory, and launch the file from that location.
 - The DB2 JDBC driver and DB2 JDBC driver (XA) cannot be found in the drop-down list of supported JDBC providers in the administrative console. The administrative console no longer displays deprecated JDBC provider names. The new JDBC provider names used in the administrative console are more descriptive and less confusing. The new providers will differ only by name from the deprecated ones.
The deprecated names will continue to exist in the `jdbc-resource-provider-templates.xml` file for migration reasons (for existing JACL scripts for example); however, you are encouraged to use the new JDBC provider names in your JACL scripts.
 - You receive the following message:
MIGR0108E: The specified WebSphere directory does not contain a WebSphere version that can be upgraded.
This can occur if an incorrect directory was specified with the `WBIPreUpgrade` command.

See the `WBIPreUpgrade` command.

- Problems can occur when you are using the `WBIPostUpgrade` command.
 - A "Not found" or "No such file or directory" message is returned.
This problem can occur if you are trying to run the `WBIPostUpgrade` command from a directory other than the WebSphere Process Server version 6.1 `install_root\bin`. Verify that the `WBIPostUpgrade` script resides in the version 6.1 `install_root\bin` directory, and launch the file from that location.
 - When you migrate the federated nodes in a cell, you receive the following error messages:
MIGR0304I: The previous WebSphere environment is being restored.
com.ibm.websphere.management.exception.RepositoryException:
com.ibm.websphere.management.exception.ConnectorException: ADMC0009E:
The system failed to make the SOAP RPC call: invoke
MIGR0286E: The migration failed to complete.
A connection timeout occurs when the federated node tries to retrieve configuration updates from the deployment manager during the `WBIPostUpgrade` migration step for the federated node. Copying the entire configuration might take more than the connection timeout if the configuration that you are migrating to version 6.1 contains any of the following elements:
 - Many small applications

- A few large applications
- One very large application

If this occurs, modify the timeout value before running the `WBIPostUpgrade` command to migrate a federated node.

1. Go to the following location in the version 6.1 directory for the profile to which you are migrating your federated node:
profile_root/properties
2. Open the `soap.client.props` file in that directory and find the value for the `com.ibm.SOAP.requestTimeout` property. This is the timeout value in seconds. The default value is 180 seconds.
3. Change the value of `com.ibm.SOAP.requestTimeout` to make it large enough to migrate your configuration. For example, the following entry would give you a timeout value of a half of an hour:
`com.ibm.SOAP.requestTimeout=1800`

Note: Select the smallest timeout value that will meet your needs. Be prepared to wait for at least three times the timeout that you select—once to download files to the backup directory, once to upload the migrated files to the deployment manager, and once to synchronize the deployment manager with the migrated node agent.

4. Go to the following location in the backup directory that was created by the `WBIPreUpgrade` command:
migration_backup_directory/profiles/default/properties
5. Open the `soap.client.props` file in that directory and find the value for the `com.ibm.SOAP.requestTimeout` property.
6. Change the value of `com.ibm.SOAP.requestTimeout` to the same value that you used in the version 6.1 file.

- You receive the "Unable to copy document to temp file" error message. Here is an example:

```
MIGR0304I: The previous WebSphere environment is being restored.
com.ibm.websphere.management.exception.DocumentIOException: Unable to copy
document to temp file:
  cells/sunblade1Network/applications/LARGEApp.ear/LARGEApp.ear
```

Your file system might be full. If your file system is full, clear some space and rerun the `WBIPostUpgrade` command.

- You receive the following message:
MIGR0108E: The specified WebSphere directory does not contain a WebSphere version that can be upgraded.

The following possible reasons for this error exist:

- An incorrect directory might have been specified when launching the `WBIPreUpgrade` command or the `WBIPostUpgrade` .
- The `WBIPreUpgrade` command was not run.

- You receive the following error message:

```
MIGR0253E: The backup directory migration_backup_directory does not exist.
The following possible reasons for this error exist:
```

- The `WBIPreUpgrade` command was not run before the `WBIPostUpgrade` command.
 1. Check to see if the backup directory specified in the error message exists.
 2. If not, run the `WBIPreUpgrade` command.
See `WBIPreUpgrade` command.

- 3. Retry the WBIPostUpgrade command.
- An incorrect backup directory might be specified.

For example, the directory might have been a subdirectory of the version 6.0.x tree that was deleted after the WBIPreUpgrade command was run and the older version of the product was uninstalled but before the WBIPostUpgrade command was run.

1. Determine whether or not the full directory structure specified in the error message exists.
2. If possible, rerun the WBIPreUpgrade command, specifying the correct full migration backup directory.
3. If the backup directory does not exist and the older version it came from is gone, rebuild the older version from a backup repository or XML configuration file.
4. Rerun the WBIPreUpgrade command.

- You decide that you need to run WBIPreUpgrade again after you have already run the WBIPostUpgrade command.

During the course of a deployment manager or a managed node migration, WBIPostUpgrade might disable the old environment. If after running WBIPostUpgrade you want to run WBIPreUpgrade again against the old installation, you must run the migrationDisablementReversal.jacl script located in the old *install_root*/bin directory. After running this JACL script, your version 6.0.x environment will be in a valid state again, allowing you to run WBIPreUpgrade to produce valid results.

For more information on scripting, see Getting started with scripting. Scripting, as described there, is available for WebSphere Process Server.

- A federated migration fails with message MIGR0405E.

The migration that has taken place on your deployment manager as part of your federated migration has failed. For a more detailed reason for why this error has occurred, open the folder *your_node_name_migration_temp* located on your deployment manager node under the ...DeploymentManagerProfile/temp directory. For example:

```
/websphere61/procserver/profiles/dm_profile/temp/nodeX_migration_temp
```

The logs and everything else involved in the migration for this node on the deployment manager node are located in this folder. This folder will also be required for IBM support related to this scenario.

- WebSphere Process Server version 6.1 applications are lost during migration.

If any of the version 6.1 applications fail to install during a federated migration, they will be lost during the synchronizing of the configurations. The reason that this happens is that one of the final steps of WBIPostUpgrade is to run a syncNode command. This has the result of downloading the configuration on the deployment manager node and overwriting the configuration on the federated node. If the applications fail to install, they will not be in the configuration located on the deployment manager node. To resolve this issue, manually install the applications after migration. If they are standard version 6.1 applications, they will be located in the *install_root*/installableApps directory.

To manually install an application that was lost during migration, use the wsadmin command to run the *install_application_name.jacl* script that the migration tools created in the backup directory.

Linux

In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migration_backup_directory/install_application_name.jacl  
-conntype NONE
```

See the Wsadmin tool.

- WebSphere Process Server version 6.1 applications fail to install.

Manually install the applications using the wsadmin command after WBIPostUpgrade has completed.

To manually install an application that failed to install during migration, use the wsadmin command to run the `install_application_name.jacl` script that the migration tools created in the backup directory.

Linux

In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migration_backup_directory/install_application_name.jacl  
-conntype NONE
```

See the Wsadmin tool, or see the WBIPostUpgrade command.

- **Solaris** When you use the migration wizard to migrate a profile from WebSphere Process Server version 6.0.x to version 6.1.x on a Solaris x64 processor-based system, the migration might fail during the WBIPostUpgrade step.

You might see messages similar to the following in `profile_root/logs/WASPostUpgrade.time_stamp.log`:

```
MIGR0327E: A failure occurred with stopNode.  
MIGR0272E: The migration function cannot complete the command.
```

WebSphere Process Server version 6.0.x uses a Java virtual machine (JVM) in 32-bit mode. The migration wizard for WebSphere Process Server version 6.1.x calls the WBIPostUpgrade.sh script, which attempts to run the JVM for version 6.0.x in the 64-bit mode when the server stops the version 6.0.x node.

Complete the following actions to remove the incomplete profile and enable WebSphere Process Server to correctly migrate the version 6.0.x profile:

1. On a command line, change to the `install_root/bin` directory.
For example, type the following command:

```
cd /opt/IBM/WebSphere/Procserver/bin
```
2. Locate the WBIPostUpgrade.sh script in the `install_root/bin` directory, and make a backup copy.
3. Open the WBIPostUpgrade.sh or WBIPostUpgrade.bat file in an editor, and perform the following actions:
 - a. Locate the following line of code:

UNIX

Linux

```
"$binDir" /setupCmdLine.sh
```

Windows

```
call "%~dp0setupCmdLine.bat" %*
```

- b. Insert the following line of code after the code that was identified in the previous step:

```
JVM_EXTRA_CMD_ARGS=""
```
 - c. Save the changes.
4. Repeat steps 2 through 4 with the `WASPostUpgrade.sh` or the `WASPostUpgrade.bat` file.

5. Use the following command to delete the incomplete version 6.1.x profile that was created during the migration process:
`install_root/bin/manageprofiles.sh -delete -profileName profile_name`
6. Delete the *profile_root* directory of the version 6.1.x profile that was removed in the previous step.
7. Rerun the migration wizard.

- If you select the option for the migration process to install the enterprise applications that exist in the version 6.0.x configuration into the new version 6.1 configuration, you might encounter some error messages during the application-installation phase of migration.

The applications that exist in the version 6.0.x configuration might have incorrect deployment information—usually, incorrect XML documents that were not validated sufficiently in previous WebSphere Process Server runtimes. The runtime now has an improved application-installation validation process and will fail to install these malformed EAR files. This results in a failure during the application-installation phase of WBIPostUpgrade and produces an "E:" error message. This is considered a "fatal" migration error.

If migration fails in this way during application installation, you can do one of the following:

- Fix the problems in the version 6.0.x applications, and then remigrate.
- Proceed with the migration and ignore these errors.

In this case, the migration process does not install the failing applications but does complete all of the other migration steps.

Later, you can fix the problems in the applications and then manually install them in the new version 6.1 configuration using the administrative console or an install script.

- After migrating to a version 6.1 cell that contains or interoperates with version 6.0.x nodes that are not at WebSphere Process Server version 6.0.1.3 or later, the cluster function might fail.

When starting these version 6.0.x servers, you might see the following problems:

- You might see a first failure data capture (FFDC) log that shows a `ClassNotFoundException` error message. This exception is thrown from the `RuleEtiquette.runRules` method and looks something like the following example:

```
Exception = java.lang.ClassNotFoundException
Source = com.ibm.ws.cluster.selection.SelectionAdvisor.<init>
probeid = 133
Stack Dump = java.lang.ClassNotFoundException: rule.local.server
at java.net.URLClassLoader.findClass(URLClassLoader.java:Compiled Code)
at com.ibm.ws.bootstrap.ExtClassLoader.findClass(ExtClassLoader.java:106)
at java.lang.ClassLoader.loadClass(ClassLoader.java:Compiled Code)
at java.lang.ClassLoader.loadClass(ClassLoader.java:Compiled Code)
at java.lang.Class.forName1(Native Method)
at java.lang.Class.forName(Class.java:Compiled Code)
at com.ibm.ws.cluster.selection.rule.RuleEtiquette.runRules(RuleEtiquette.java:154)
at com.ibm.ws.cluster.selection.SelectionAdvisor.handleNotification(SelectionAdvisor.java:153)
at com.ibm.websphere.cluster.topography.DescriptionFactory$Notifier.run(DescriptionFactory.java:257)
at com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1462)
```

- You might see a `java.io.IOException` that looks something like the following example:

```
Exception = java.io.IOException
Source = com.ibm.ws.cluster.topography.DescriptionManagerA.update probeid
= 362
```



```

Stack Dump = java.io.IOException
at com.ibm.ws.cluster.topography.ClusterDescriptionImpl.importFromStream
(ClusterDescriptionImpl.java:916)
at com.ibm.ws.cluster.topography.DescriptionManagerA.update
(DescriptionManagerA.java:360)
Caused by: java.io.EOFException
at java.io.DataInputStream.readFully(DataInputStream.java(Compiled Code))
at java.io.DataInputStream.readUTF(DataInputStream.java(Compiled Code))
at com.ibm.ws.cluster.topography.KeyRepositoryImpl.importFromStream
(KeyRepositoryImpl.java:193)

```

During migration, version 6.1 cluster information is distributed throughout the cell. WebSphere Process Server version 6.0.x nodes that are not at version 6.0.1.3 or later fail to read this information.

To avoid this problem, upgrade all version 6.0.x nodes that will be contained in or interoperating with a version 6.1 cell to version 6.0.1.3 or later before migrating your deployment managers to version 6.1.

- After you migrate a managed node to version 6.1, the application server might not start.

When you try to start the application server, you might see errors similar to those in the following example:

```

[5/11/06 15:41:23:190 CDT] 0000000a SystemErr R
com.ibm.ws.exception.RuntimeError:
com.ibm.ws.exception.RuntimeError: org.omg.CORBA.INTERNAL:
CREATE_LISTENER_FAILED_4
vmcid: 0x49421000 minor code: 56 completed: No
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServerImpl.bootServerContainer(WsServerImpl.java:198)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServerImpl.start(WsServerImpl.java:139)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServerImpl.main(WsServerImpl.java:460)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
com.ibm.ws.runtime.WsServer.main(WsServer.java:59)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr R at
sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:64)
[5/11/06 15:41:23:197 CDT] 0000000a SystemErr R at
sun.reflect.DelegatingMethodAccessorImpl.invoke
(DelegatingMethodAccessorImpl.java:43)

```

Change the port number at which the managed node's server is listening. If the deployment manager is listening at port 9101 for ORB_LISTENER_ADDRESS, for example, the server of the managed node should not be listening at port 9101 for its ORB_LISTENER_ADDRESS. To resolve the problem in this example, perform the following steps:

1. On the administrative console, click **Application servers** → *server_name* → **Ports** → **ORB_LISTENER_ADDRESS**.
 2. Change the ORB_LISTENER_ADDRESS port number to one that is not used.
- If synchronization fails when you migrate a managed node to version 6.1, the server might not start.

You might receive messages similar to the following when you migrate a managed node to version 6.1:

```

ADMU0016I: Synchronizing configuration between node and cell.
ADMU0111E: Program exiting with error:
com.ibm.websphere.management.exception.AdminException: ADMU0005E:
Error synchronizing repositories
ADMU0211I: Error details may be seen in the file:
/opt/WebSphere/61AppServer/profiles/AppSrv02/logs/syncNode.log
MIGR0350W: Synchronization with the deployment manager using the SOAP protocol

```

failed.
MIGR0307I: The restoration of the previous WebSphere Application Server environment is complete.
MIGR0271W: Migration completed successfully, with one or more warnings.

These messages indicate the following:

- Your deployment manager is at a version 6.1 configuration level.
- The managed node that you are trying to migrate is at a version 6.1 configuration level on the deployment manager’s repository (including applications).
- The managed node itself is not quite complete given that you did not complete the syncNode operation.

Perform the following actions to resolve this issue:

1. Rerun the syncNode command on the node to synchronize it with the deployment manager.
See the syncNode command .
2. Run the GenPluginCfg command.
See the GenPluginCfg command .

What to do next

If you did not find your problem listed, contact IBM support.

Related concepts

“Migration considerations for Business Process Choreographer” on page 105
If your servers run Business Process Choreographer, you should be aware of some restrictions and additional tasks you might need to perform.

Related tasks

“Verifying migration” on page 87
Verify that your migration was successful by checking the log files and checking operation with the administrative console.

Related reference

“WBIPreUpgrade command” on page 12
Use the WBIPreUpgrade command for WebSphere Process Server to save the configuration of a previously installed version of WebSphere Process Server into a migration-specific backup directory.



WBIPostUpgrade command

Use the WBIPostUpgrade command for WebSphere Process Server to retrieve the saved configuration that was created by the WBIPreUpgrade command from the *backupDirectory* that you specified. The WBIPostUpgrade command for WebSphere Process Server reads the configuration from this directory to migrate to the more recent version of WebSphere Process Server and adds all migrated applications into the *profile_root/installedApps* directory for the new installation.

Related information

Debugging components in the Application Server Toolkit

Wsadmin tool

syncNode command

GenPluginCfg command



Troubleshooting and support

To help you understand, isolate, and resolve problems with your IBM software, the troubleshooting and support information contains instructions for using the problem-determination resources that are provided with your IBM products.

Getting started with scripting

Troubleshooting migration from WebSphere InterChange Server

Find solutions to problems you encounter with migration as well as instructions for turning on logging and tracing.

Related concepts

“Limitations when migrating from WebSphere InterChange Server” on page 163
Some characteristics of WebSphere InterChange Server are not precisely duplicated by WebSphere Process Server. Therefore you might have to modify your applications after migration to get them to perform as they did in WebSphere InterChange server.

Related reference

“Postmigration considerations” on page 128

When applications have been migrated from WebSphere InterChange Server to WebSphere Process Server, special attention is required in some areas to enable migrated applications to function in WebSphere Process Server consistently with their intended function due to differences between the architectures of WebSphere Process Server and WebSphere InterChange Server.

“Premigration considerations” on page 118

Consider these guidelines for the development of integration artifacts for WebSphere InterChange Server to ease the migration of WebSphere InterChange Server artifacts to WebSphere Process Server.

Enabling logging and tracing for supported WebSphere InterChange Server APIs

Enable logging and tracing for the supported WebSphere InterChange Server APIs through the administrative console.

About this task

If your migrated application includes any supported WebSphere InterChange Server APIs, you can enable logging and tracing for them for troubleshooting purposes.

Procedure

1. Launch the administrative console.
2. From the left (navigation) panel, select **Troubleshooting > Logs and Trace**.
3. In the right panel, select the name of the server on which you want to enable logging and tracing.
4. In the right panel, under “General properties,” select **Change Log Level Details**.
5. Select the Runtime tab. (Selecting the Runtime tab allows you to make this change in real time without requiring you to restart the server.)
6. Add the name of the package followed by =all to the list of logged packages in the box on the screen. Separate this new entry from any existing entries with a colon. For example, CxCommon=all. In this case, CxCommon is the name of the package for a set of supported WebSphere InterChange Server APIs.

Specifying all enables all logging and tracing. See Supported WebSphere InterChange Server APIs for a list of the APIs, including their package names.

7. Select **Apply**.
8. To keep this configuration after the server is restarted, select the **Save runtime changes to configuration as well** check box.
9. Select **OK**.
10. When the next screen appears, select **Save** to save your changes.

Related reference



Supported WebSphere InterChange Server APIs

In addition to the WebSphere InterChange Server source artifact migration tools provided in WebSphere Process Server and WebSphere Integration Developer, WebSphere Process Server also provides support for many of the APIs that were provided in WebSphere InterChange Server. The migration tools work in conjunction with these WebSphere InterChange Server APIs by preserving your custom snippet code as much as possible when migrating.

Failure trying to serialize an object that is not serializable in a migrated BPEL file

If a serialization failure occurs in a BPEL file generated by the migration, you might be able to modify it to prevent the failure from occurring.

Problem: A serialization failure occurs in a custom snippet node of a business process execution language (BPEL) file generated by migration because an attempt is made to serialize an object that is not serializable.

Cause: In WebSphere InterChange Server, a Collaboration Template is compiled into a single Java class. In WebSphere Process Server, each node in a BPEL file might compile into a separate Java class. In WebSphere InterChange Server, a variable can be declared once and shared throughout the various steps of a Collaboration Template. To simulate that behavior in the migrated BPEL file, each variable used in a code snippet must be retrieved at the start of the snippet and saved at the end of the snippet. Variables defined in WebSphere InterChange Server Port definitions become BPEL variables. These are retrieved into BusObj variables at the beginning of each snippet (if referenced in the snippet) and saved back to the BPEL variables at the end of each snippet. For example, a retrieval at the beginning of snippets looks like this:

```
BusObj tempBusObj = null;if (tempBusObj_var != null) { tempBusObj =
    new BusObj(tempBusObj_var); };
```

and a save at the end of snippets looks like this:

```
if (tempBusObj == null) { tempBusObj_var = null; } else { tempBusObj_var =
    tempBusObj.getBusinessGraph(); }
```

Other variables used in the WebSphere InterChange Server snippet code are serialized and stored as a String in a BPEL variable named *CollabTemplateName_var*. These variables are deserialized at the beginning of each BPEL snippet, and then serialized and saved at the end of each BPEL Snippet that they are referenced in. For example, objects are retrieved like this:

```
BusObj tempBusObj = (BusObj)BaseCollaboration.deserialize
    (FrontEndCollab_var.getString("tempBusObj"));
```

and objects are saved like this:

```
FrontEndCollab_var.setString("tempBusObj", BaseCollaboration.serialize(tempBusObj));
```

If the type of the object being serialized is not serializable, then using `serialize` and `deserialize` will fail when the BPEL is run.

Solution: After migration, modify the BPEL file as follows:

- For any variable that is not Java-serializable, update the BPEL snippets to remove the serialization and deserialization statements. If the variable needs to be shared across snippets (instead of being re-created in each snippet) another method must be used to preserve the value of the variable across snippets.
- Manually define BPEL variables for variables of type `BusObj` that are not declared in the WebSphere InterChange Server Port definitions but are used on Partner Invokes. This is a manual step because variables used during invokes in WebSphere Process Server must be strongly typed, and the migration tools cannot accurately determine that type from the WebSphere InterChange Server snippets.

Note: The naming convention used by the migration tools is to add `_var` to the name of the variable in the snippet code when naming the BPEL variables. For example, for a variable called `tempBusObj` in the snippet code, the migration tools will create a BPEL variable named `tempBusObj_var`.

- For variables that must be declared manually as BPEL variables, change the BPEL snippet code so that it uses the "deserialize/serialize" method of preserving these variables rather than the "retrieve from/store into BPEL variable" method of preserving these variables.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
577 Airport Blvd., Suite 800
Burlingame, CA 94010
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: (c) (your company name) (year). Portions of

this code are derived from IBM Corp. Sample Programs. (c) Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

Warning: Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

Trademarks and service marks

IBM, the IBM logo, developerWorks, DB2, i5/OS, Informix, MQSeries, OS/390, RACF, WebSphere, and z/OS are registered trademarks and Cloudscape, DB2 Universal Database, and MVS are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

This product includes software developed by the Eclipse Project (<http://www.eclipse.org>).



IBM WebSphere Process Server for Multiplatforms, Version 6.1.0



Printed in USA