Chapter 5 of *WebSphere Application Server V7 Administration and Configuration Guide,* SG24-7615

WebSphere Application Server V7: Administration Consoles and Commands

WebSphere® application server properties are stored in the configuration repository as XML files. It is not a good idea to manually edit any of the configuration files because this bypasses validation of any changes and could lead to synchronization-related problems. Rather, WebSphere application server provides administrative tools that help you administer the environment. These tools manage modifications to the files in the repository.

In this chapter we introduce the administrative consoles and command line administration. We cover the following topics:

Introducing the WebSphere administrative consoles

The WebSphere administrative consoles are graphical, Web-based tools that you use to configure and manage the resources within the scope of the console. With the introduction of the flexible management topologies, there are multiple administrative consoles available in a WebSphere solution:

- Administrative console hosted by an application server or deployment manager:

  This console is used to manage an entire WebSphere cell. It supports the full range of product administrative activities, such as creating and managing resources and applications, viewing product messages, and so on.

  In a stand-alone server environment, the administrative console is located on the application server and can be used to configure and manage the resources of that server only.

  In a distributed server environment, the administrative console is located in the deployment manager server, dmgr. In this case, the administrative console provides centralized administration of multiple nodes. Configuration changes are made to the master repository and pushed to the local repositories on the nodes by the deployment manager.

- Administrative agent console:

  An administrative agent hosts the administrative console for application server nodes that are registered to it.

  When you access the URL for the console, you can select the node type to manage. After your selection is made, you will be directed to the appropriate console where you can log in:

  - Administrative console for the administrative agent:

    This console allows you to manage the administrative agent, including security settings, and registering nodes that it controls with the job manager.

  - Console for an application server:

    This console is the administrative console for the application server.

- Job manager administrative console (referred to as the job manager console):

  The job manager console provides the interface to manage the job manager itself, including security settings and mail resources. Its primary function, though, is to allow you to submit jobs for processing on the nodes that are registered to it.

Starting and accessing the consoles

The way that you access the administrative console is the same whether you have a stand-alone server environment or a distributed server environment. However, the location and how you start the necessary processes will vary.

Finding the URL for the console:

Each server process that hosts the administrative console has two admin ports that are used to access the console. These ports are referred to as:

- WC_adminhost
- WC_adminhost_secure (for SSL communication)

These ports are assigned at profile creation. If you do not know what the port number is for the console, you can look in the following location:

profile_home/properties/portdef.props

You can always use the following URL to access the console:

http://<*hostname*>:*WC_adminhost*/ibm/console

If administrative security is enabled, you will automatically be redirected to the secure port.

Administrative console in a standalone server environment

In a single application server installation, the console is hosted on the application server, so you must start the server in order to reach the console.

To access the administrative console, do the following steps:

1. Make sure that application server, server1, is running by using this command:

   serverStatus.sh -all

2. If the status of server1 is not STARTED, start it with the following command:

   startServer.sh *server_name*

3. Open a Web browser to the URL of the administrative console. For example:

   http://<*hostname*>:9060/ibm/console

   <*hostname*> is your host name for the machine running the application server.

Administrative console in a distributed environment

If you are working with a deployment manager and its managed nodes, the console is hosted on the deployment manager. You must start it in order to use the console. To access the administrative console, do the following steps:

1. Make sure that deployment manager, dmgr, is running by using this command:

   serverStatus.sh -all

2. If the dmgr status is not STARTED, start it with the following command:

   startManager.sh

3. Open a Web browser to the URL of the administrative console. For example:

   `http://<hostname>:9060/admin`

   <hostname> is your host name for the machine running the deployment manager.

Job manager console

To access the job manager administrative console, do the following steps:

1. Make sure that job manager process (jobmgr) is running by using this command:

   `serverStatus.sh -all`

2. If the status of jobmgr is not STARTED, start it with the following command:

   `startServer.sh jobmgr`

3. Open a Web browser to the URL of the administrative console. For example:

   `http://<hostname>:9960/ibm/console`

Administrative agent console

1. Make sure that administrative agent process (adminagent) is running by using this command:

   `serverStatus.sh -all`

2. If the status of adminagent process is not STARTED, start it with the following command:

   `startServer.sh adminagent`

3. Open a Web browser to the URL of the administrative console. For example:

   `http://<hostname>:9060/ibm/console`

   If you have nodes registered with the administrative agent, you will be prompted to select which node you would like to administer which includes the administrative agent.

4. Log in to the selected console.

Logging in to a console

The user ID specified during login is used to track configuration changes made by the user. This allows you to recover from unsaved session changes made under the same user ID, for example, when a session times out or the user closes the Web browser without saving. The user ID for login depends on whether WebSphere administrative security is enabled:

Note: You cannot log on to two instances of administrative consoles that are running on the same machine from a single browser type. For example, if you use Firefox to log in to the deployment manager administrative console, you cannot also log in to a job manager running on the same machine.

There is a limitation that cookies are unique per domain rather than a combination of domain and port. Therefore, the cookies that control the session and authentication data in the first browser tab or window get overwritten when logging into the other console in a new browser tab or window. However, you should be able to log in to two consoles simultaneously from two completely different browsers, for example, Firefox and Internet Explorer®.

- If WebSphere administrative security is not enabled:

  You can enter any user ID, valid or not, to log in to the administrative console. The user ID is used to track changes to the configuration, but is not authenticated. You can also simply leave the User ID field blank and click the **Log In** button.

  Note: Logging in without an ID is not a good idea if you have multiple administrators.

- If WebSphere administrative security is enabled:

  You must enter a valid user ID and password that have been assigned an administrative security role.

  If you enter a user ID that is already in session, you will receive the message `Another user is currently logged in with the same User ID` and you will be prompted to do one of the following actions:

  – Force the existing user ID out of session. You will be allowed to recover changes that were made in the other user's session.

  – Specify a different user ID.

Note: You can also get this message if a previous session ended without a logout. For example, if the user closed a Web browser during a session and did not log out first or if the session timed out.

Recovering from an interrupted session

Until you save the configuration changes you make during a session, the changes do not become effective. If a session is closed without saving the configuration changes made during the session, these changes are remembered and you are given the chance to pick up where you left off.

When unsaved changes for the user ID exist during login, you are prompted to do one of the following actions:

- Work with the master configuration:

  Selecting this option specifies that you want to use the last saved administrative configuration. Changes made to the user's session since the last saving of the administrative configuration will be lost.

- Recover changes made in a prior session:

  Selecting this option specifies that you want to use the same administrative configuration last used for the user's session. It recovers all changes made by the user since the last saving of the administrative configuration for the user's session.

As you work with the configuration, the original configuration file and the new configuration file are stored in a folder at:

`<profile_home>/wstemp`

After you save the changes, these files are removed from the `wstemp` folder.

Changing the administrative console session timeout

You might want to change the session timeout for the administrative console application. The session timeout is the time it takes for the console session to time out after a period of idleness. The default is 15 minutes. To change the session timeout value, see the following page in the Information Center:

- Changing the console session expiration:

  http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.nd.multiplatform.doc/info/ae/isc/cons_sessionto.html

The console application name is isclite. This is true for all administrative consoles (job manager, administrative agent, deployment manager, or standalone application server console).

The graphical interface

The WebSphere administrative consoles have the same layout pattern. In each console, you can find the following main areas:

- Banner
- Navigation tree
- Workspace, including the messages and help display areas

Each area can be resized as desired. The difference in the console types will be in the Navigation tree. The options that you find there will vary depending on the console type.

Figure 1 uses the administrative console hosted on a deployment manager to illustrate the console layout.

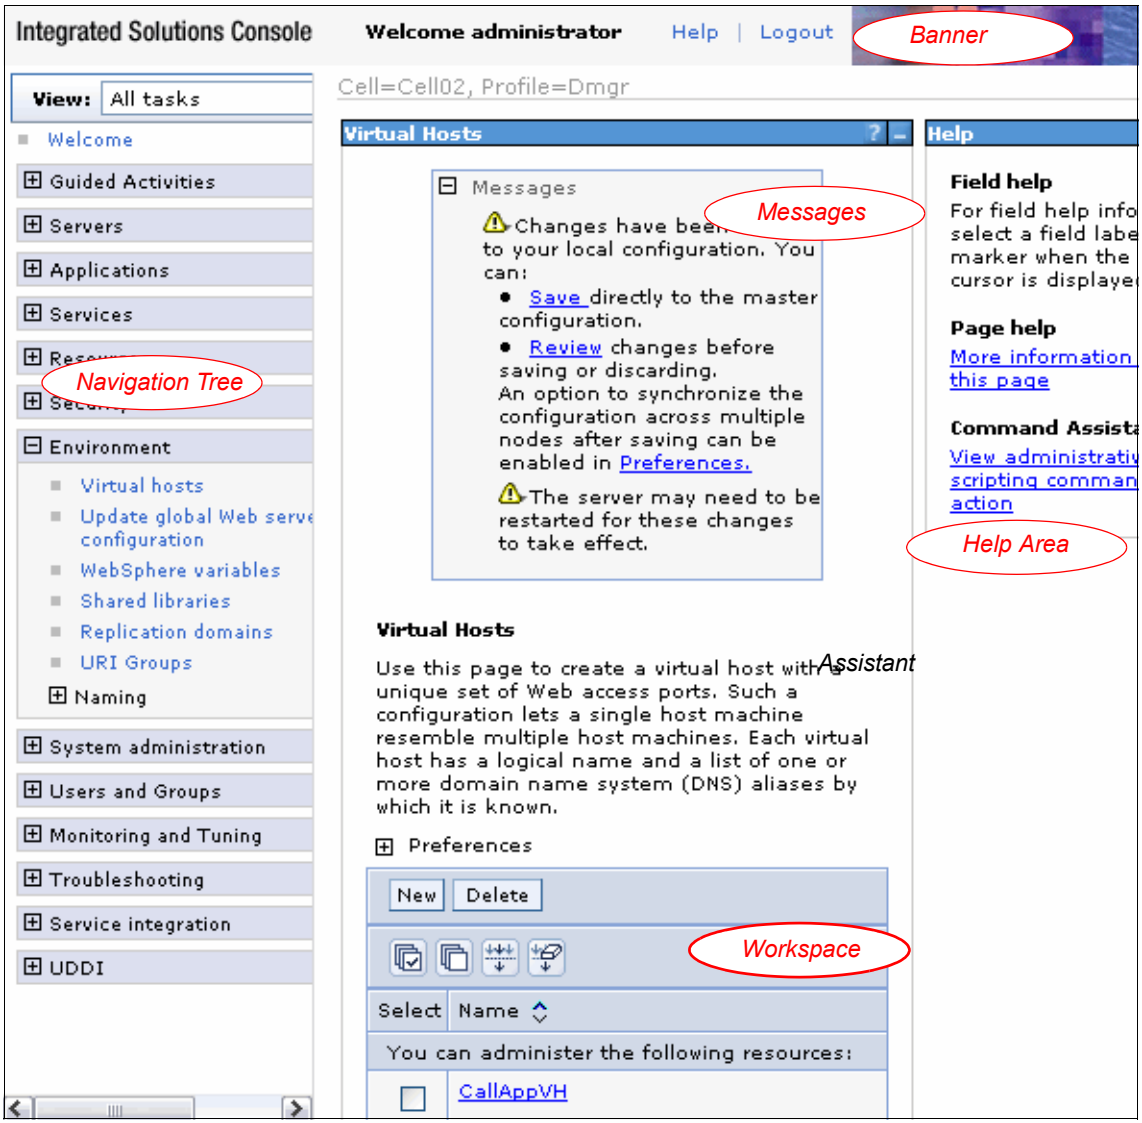

Figure 1 The administrative console graphical interface

Banner

The *banner* is the horizontal bar near the top of the console. The banner provides the following actions:

- Logout logs you out of the administrative console session and displays the Login page. If you have changed the administrative configuration since last saving the configuration to the master repository, the Save page displays before returning you to the Login page. Click **Save** to save the changes, **Discard** to return to the administrative console, or **Logout** to exit the session without saving changes.
- **Help** opens a new Web browser with detailed online help for the administrative console. This is not part of the Information Center.

Console identity (new in V7)

The banner can be customized to show a unique identifier for the console. This can be helpful in cases where administrators log on to multiple administrative consoles. Glancing at the banner lets you know which system you are logged on to. You can add a Console ID from the administrative console (Figure 2 on page 10).

To customize the banner, navigate to **System environment** → **Console Identity**. Select **Custom** and enter the identity string. Save the changes, and log out of the console, then back in. This console identity will be displayed to all users that log in to that console application.

In an administrative agent configuration, the changes are applied to the administrative agent and all of its registered application servers, regardless of where the changes were actually saved.

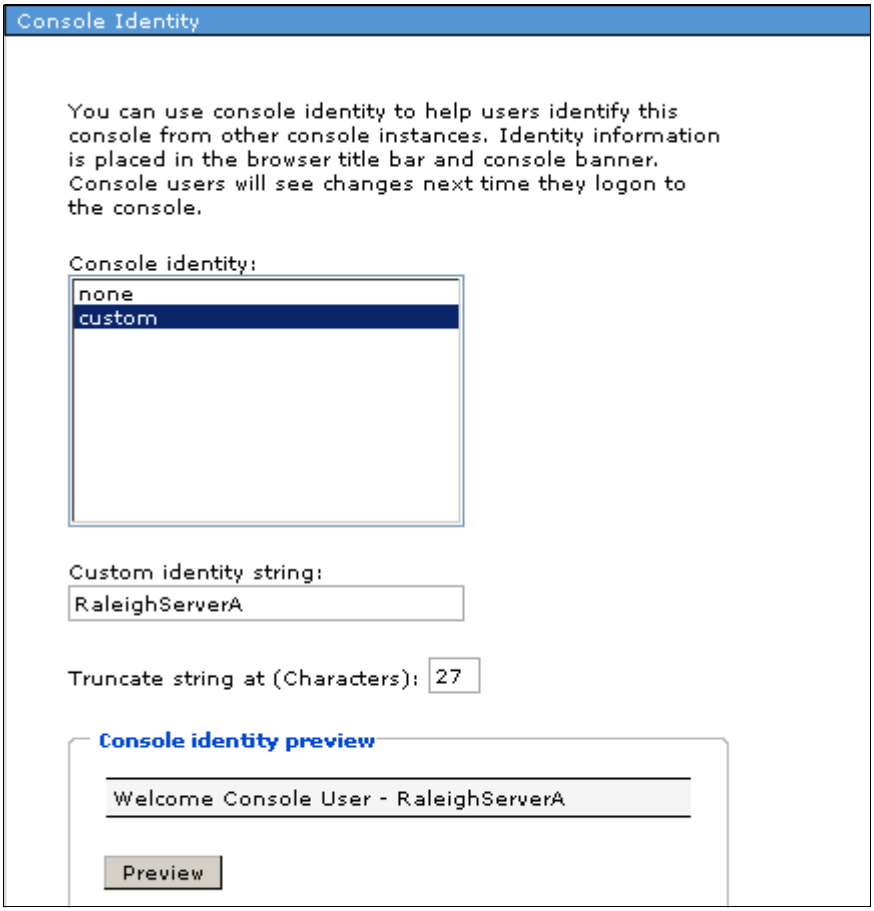

Figure 2 Changing the console identity

After you log back in, you will see the new console identity in the banner (Figure 3).

Figure 3 Banner with a customized console identity

Navigation tree

The navigation tree on the left side of the console offers links for you to view, select, and manage components.

Clicking a **+** beside a tree folder or item expands the tree for the folder or item. Clicking a **-** collapses the tree for the folder or item. Double-clicking an item toggles its state between expanded and collapsed.

The content displayed on the right side of the console, the *workspace*, depends on the folder or item selected in the tree view.

Guided activities

The navigation tree includes a category called “Guided Activities”. This section contains step-by-step assistance for performing some common tasks. These activities can be accomplished by performing each task separately, but using the Guided Activities option provides additional assistance.

Workspace

The workspace, on the right side of the console in Figure 1 on page 8, allows you to work with your administrative configuration after selecting an item from the console navigation tree.

When you click a folder in the tree view, the workspace lists information about instances of that folder type, the collection page. For example, selecting **Servers** → **Server Types** → **WebSphere application servers** shows all the application servers configured in this cell. Selecting an item, an application server in this example, displays the detail page for that item. The detail page can contain multiple tabs. For example, you might have a Runtime tab for displaying the runtime status of the item, and a Configuration tab for viewing and changing the configuration of the displayed item.

Messages

When you perform administrative actions, messages are shown at the top of the workspace to display the progress and results. These messages are limited in nature so if an action fails, review the JVM™ process logs for more detailed information.

When configuration changes have been made, the message area will contain links that you can click to review or save the changes.

Breadcrumb trail

As you navigate into multiple levels of a configuration page, a breadcrumb trail is displayed at the top of the workspace. It indicates how you reached the current page and provides links that allow you to go back to previous pages easily without starting the navigation trail over (Figure 4).

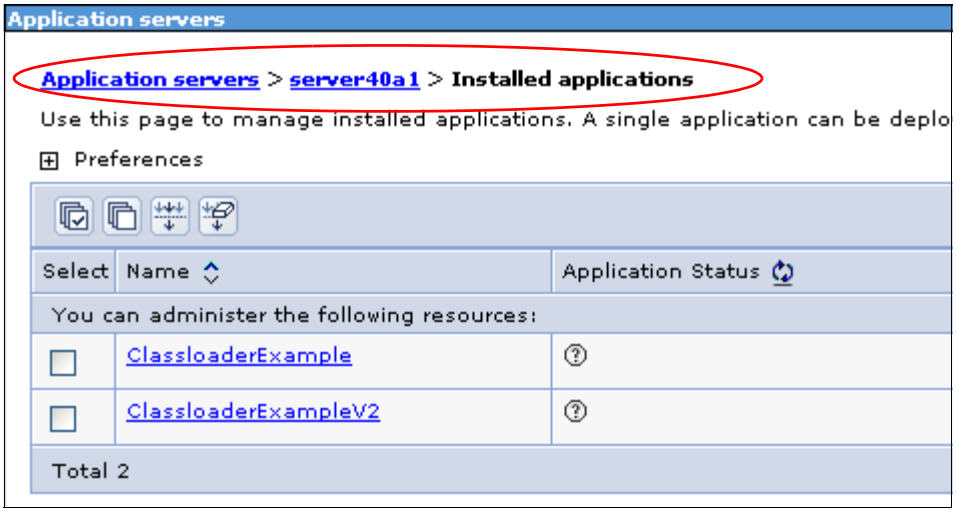

Figure 4 Breadcrumb trail

Help area

As you are working in the administrative console, help is available in multiple ways. As you hover the mouse over a field, help text will be displayed for that field. In addition, most pages will have a **More information about this page** link in the Help area. Clicking the link will open the online help in a separate browser. And finally, many pages will have a **View administrative scripting command for last action** link. Clicking this link will display an equivalent scripting command for the action you just performed.

Setting console preferences

The look of the administrative console can be altered by setting console preferences. The preference you see will vary slightly depending on the console type. For example, the preference to synchronize changes with nodes is only applicable to a console on a deployment manager. See Figure 5 on page 13.

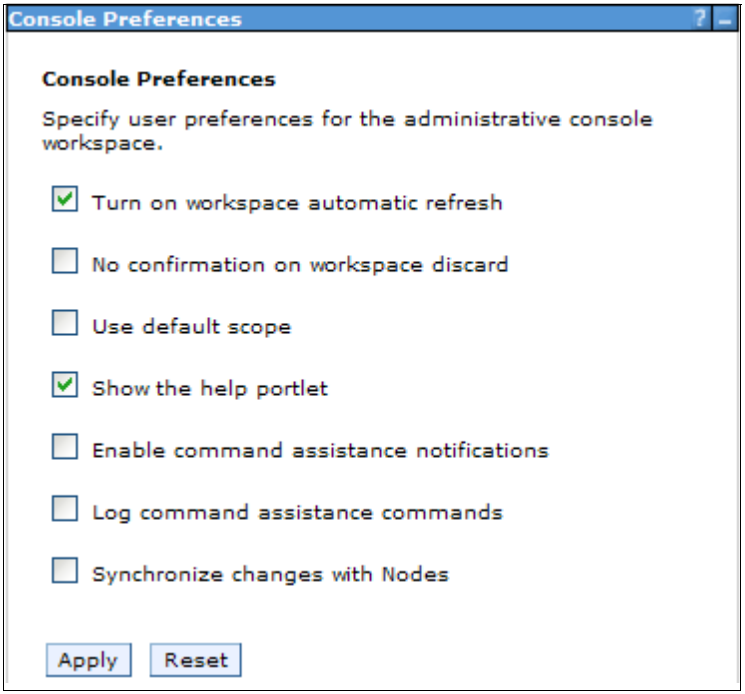

Figure 5 Administrative console preferences

To set console preferences, select **System Administration →Console Preferences** in the navigation tree. You have the following options:

- **Turn on WorkSpace Auto-Refresh** specifies that the view automatically refreshes after a configuration change. If it is not selected, you must re-access the page to see the changes.
- **No Confirmation on Workspace Discard** specifies that a confirmation window be displayed if you elect to discard the workspace. For example, if you have unsaved changes and log out of the console, you will be asked whether you want to save or discard the changes. If this option is not selected and you elect to discard your changes, you will be asked to confirm the discard action.
- **Use default scope** (administrative console node) sets the default scope to the node of the administration console. If you do not enable this setting, the default is all scopes.
- **Show the help portlet** displays the help portlet at right top.

- **Enable command assistance notifications** allows you to send JMX™ notifications that contain command data. These notifications can be monitored in a Rational® Application Developer workspace, providing assistance in creating scripts.
- **Log command assistance commands** specifies whether to log all the command assistance wsadmin data for the current user.

  When you select this option, script commands matching actions you take in the console are logged to the following location:

  profile_root/logs/AssistanceJythonCommands_*<user_name>*.log
- **Synchronize changes with Nodes** synchronizes changes that are saved to the deployment manager profile with all the nodes that are running.

Click the boxes to select which preferences you want to enable and click **Apply**.

Finding an item in the console

To work with items in the console, do the following steps:

1. Select the associated task from the navigation tree. For example, to JDBC™ providers that have been defined, select **Resources** → **JDBC** → **JDBC Providers**. See Figure 6 on page 15.
2. Certain resources are defined at a scope level. If applicable, select the scope from the drop-down.
3. Set the preferences to specify how you would like information to be displayed on the page.

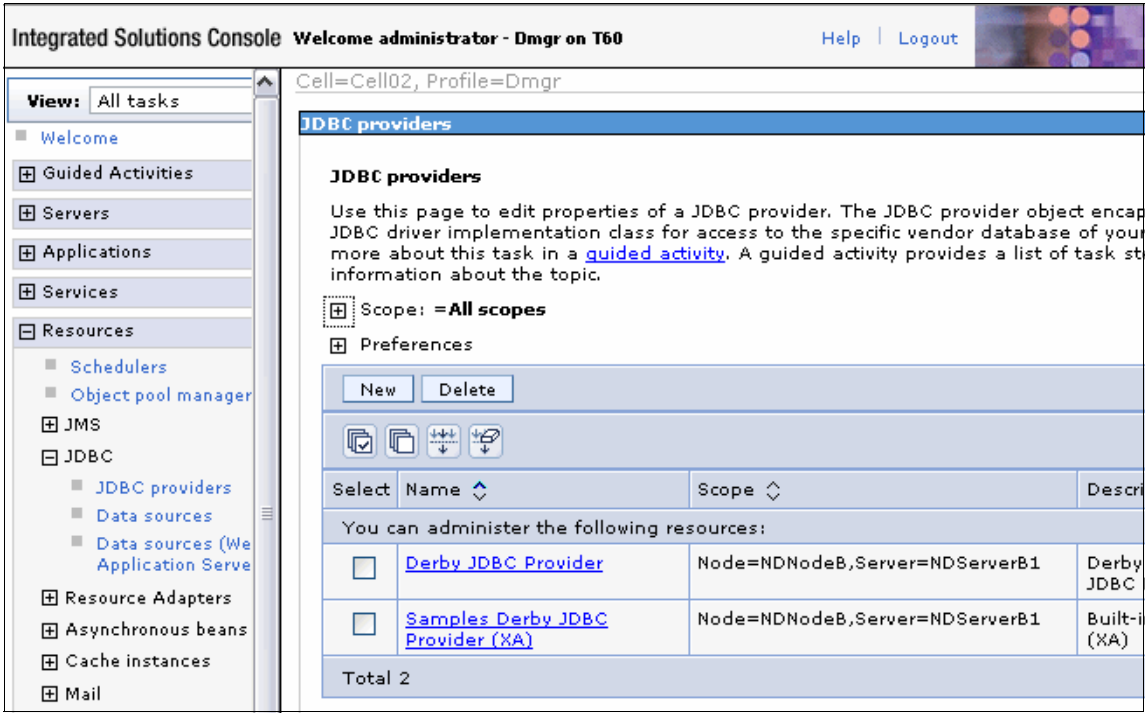

Figure 6 Working with the administrative console

Selecting a scope

The scope level determines which applications or application servers will see and use that configuration. The scope setting is available for all resource types, WebSphere variables, shared libraries, and name space bindings.

Scope levels

Configuration information is defined at the following levels: cell, cluster, node, server, and application. Here, we list these scopes in overriding sequence. Because you see application scope first, anything defined at this scope overrides any conflicting configuration you might find in the higher level scopes:

1. Resources and variables scoped at the *application* level apply only to that application. Resources and variables are scoped at the application level by defining them in an enhanced EAR. They cannot be created from the WebSphere administrative tools, but can be viewed and modified (in the administrative console, navigate to the details page for the enterprise application and select **Application scoped resources** in the References section).

2. Resources scoped at the *server* level apply only to that server. If a node and server combination is specified, the scope is set to that server. Shared libraries configured in an enhance EAR are automatically scoped at the server level.
3. Resources scoped at the *node* level apply to all servers on the node.
4. Resources scoped at the *cluster* level apply to all application servers in the cluster. New cluster members automatically have access to resources scoped at this level. If you do not have any clusters defined, you will not see this option.
5. Resources scoped at the *cell* level apply to all nodes and servers in the cell.

Standalone application servers: Although the concept of cells and nodes is more relevant in a managed server environment, scope is also set when working with standalone application servers. Because there is only one cell, node, and application server, and no clusters, simply let the scope default to the node level.

Configuration information is stored in the repository directory that corresponds to the scope. For example, if you scope a resource at the node level, the configuration information for that resource is in:

`<profile_home>/config/cells/cell_name/nodes/<node>/resources.xml`

If you scoped that same resource at the cell level, the configuration information for that resource is in:

`<profile_home>/config/cells/cell_name/resources.xml`

Setting scope levels in the console

Collection pages that contain items that require a scope level to be identified provide two different options for defining the scope. Setting the scope level both sets the level for any resources you create and limits what is displayed in the collection page.

Selecting the **Show scope selection drop-down list with the all scopes** option provides a drop-down box with all scopes that you can select from, including the “All scopes” option (Figure 7 on page 17). Selecting a scope from the drop-down list changes the scope automatically.

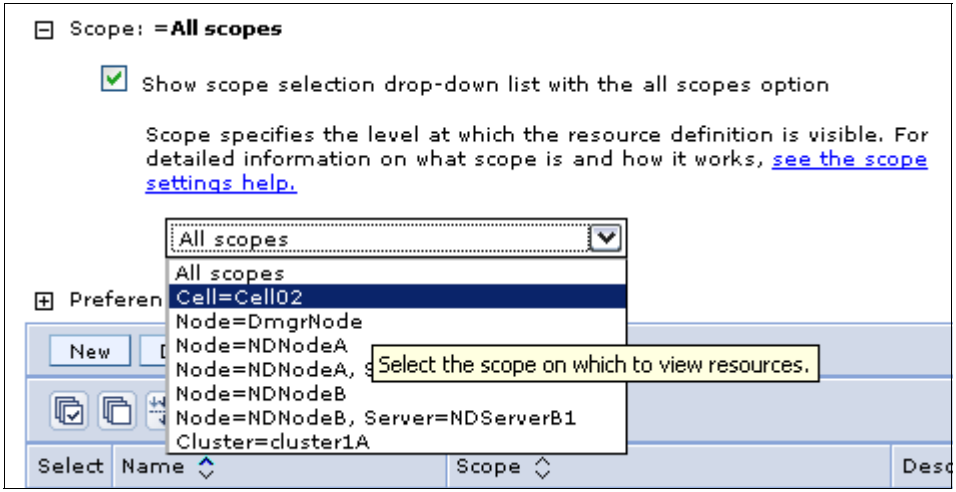

Figure 7 Scopes selected using a drop-down list

(New in V7) The second option for setting the scope is to de-select the **Show scope selection drop-down list with the all scopes.** Instead of a drop-down, you have fields for each scope level where you can browse a list of applicable entries at that scope level. Click **Apply** to complete the selection (Figure 8 on page 18).

The scope is set to the lowest level entry you select (a red arrow to the left of the field indicates the current scope). To move to a higher scope, simply clear the lower field. For example, if you select a server as the scope level, and want to change the scope to the node level, clear the server field and click **Apply**.

This option is useful in cells that contain a large number of nodes, servers, or clusters. In those situations, the drop-down list can be difficult to navigate. However, note that the option to view all scopes is not available.

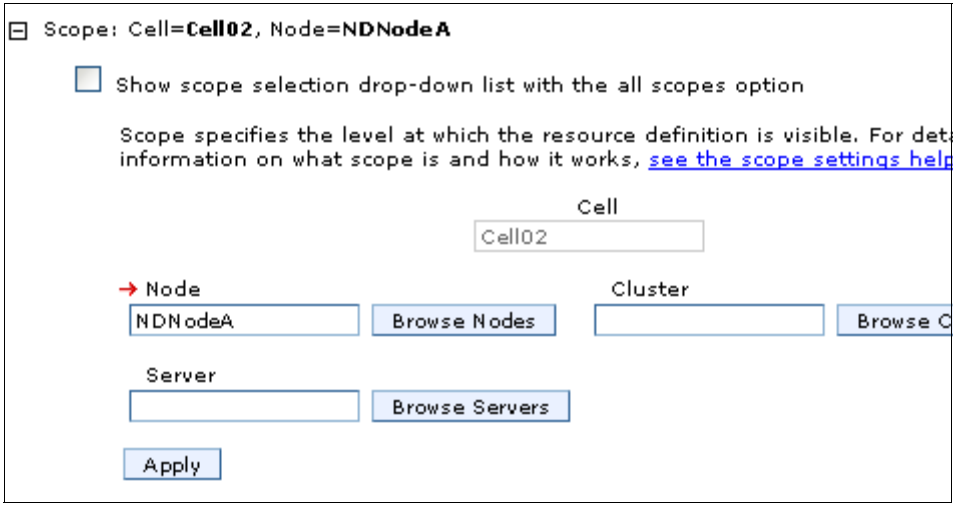

Figure 8 Selecting the scope with individual fields

Set preferences for viewing the console page

After selecting a task and a scope, the administrative console page shows a collection table with all the objects created at that particular scope. You can change the list of items you see in this table by using the filter and preference settings.

The preference settings available will vary by the type of item you are displaying. Many of these settings are new in V7. A list of the preference settings and their use is available in the Information Center:

- Administrative console preference settings:

  `http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.nd.multiplatform.doc/info/ae/ae/rcon_preferences.html`

Figure 9 on page 19 shows the preference settings you would see when displaying a list of JDBC providers.

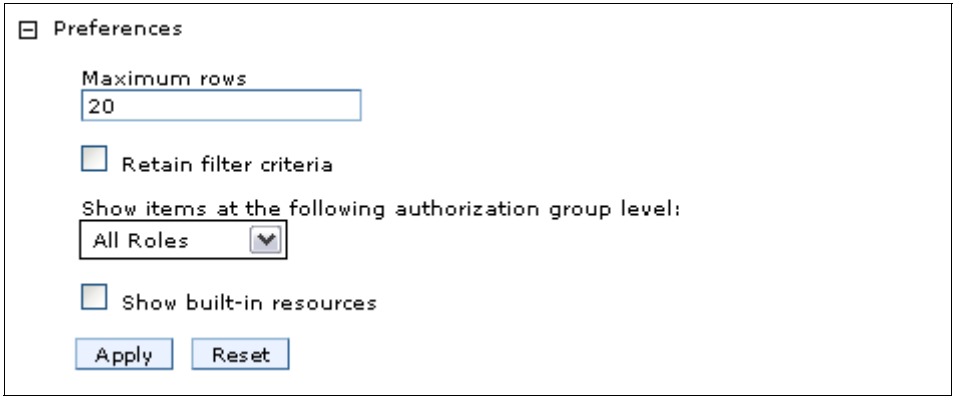

Figure 9 Filter and preference settings

The filter options can be displayed or set by clicking the Show Filter Function icon at the top of the table. See Figure 10.

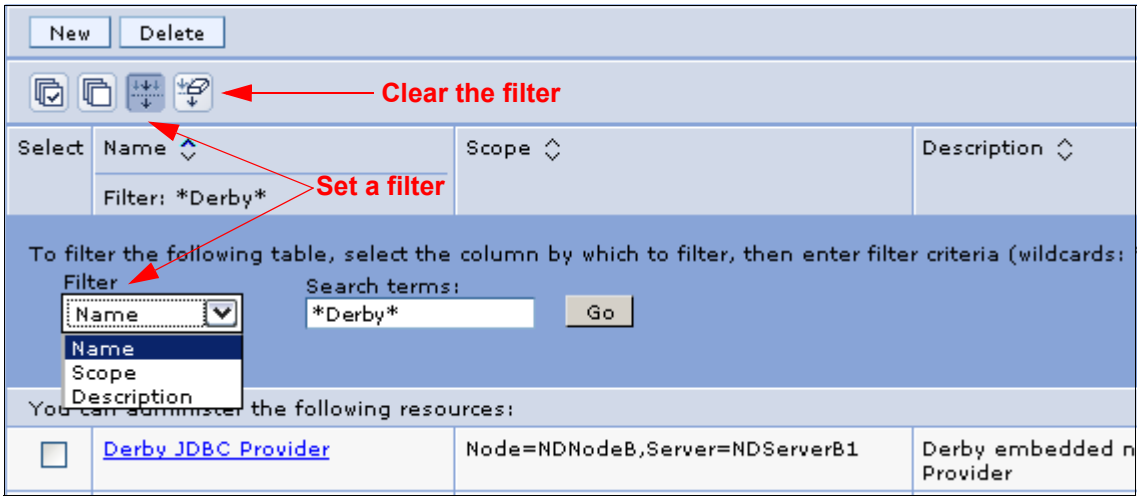

Figure 10 Setting filters and preferences

When you click the icon, a new area appears at the top of the table allowing you to enter filter criteria. To filter entries, do the following steps:

1. Select the column to filter on. For example, in Figure 10, the display table has three columns to choose from. Your options vary depending on the type of item you are filtering.

2. Enter the filter criteria. The filter criteria is case sensitive and wild cards can be used. In our example, to see only providers with names starting with “S”, select the Name column to filter on and enter S* as the filter.
3. Click **Go**.
4. After you have set the filter, click the **Show Filter** Icon again to remove the filter criteria from view. You still have a visual indication that the filter is set at the top of the table.

Setting the filter is temporary and only lasts for as long as you are in that collection. To keep the filter active for that collection, check the **Retain filter criteria** box in the Preferences section and click **Apply**. To clear the filter criteria, click the  icon.

For more help on using the filtering feature, see:

- Administrative console buttons:

  http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.nd.multiplatform.doc/info/ae/ae/rcon_buttons.html

Updating existing items

To edit the properties of an existing item, complete these tasks:

1. Select the category and type in the navigation tree. For example, select **Servers** → **Server Types** → **WebSphere application servers**.
2. A list of the items of that type in the scope specified will be listed in a collection table in the workspace area. Click an item in the table. This opens a detail page for the item.
3. In some cases, you see a Configuration tab and a Runtime tab on this page. In others, you only see a Configuration tab.

   Updates are done under the Configuration tab. Specify new properties or edit the properties already configured for that item. The configurable properties will depend on the type of item selected.

For example, if you select a WebSphere Application Server cluster, this opens a detail page resembling Figure 11.

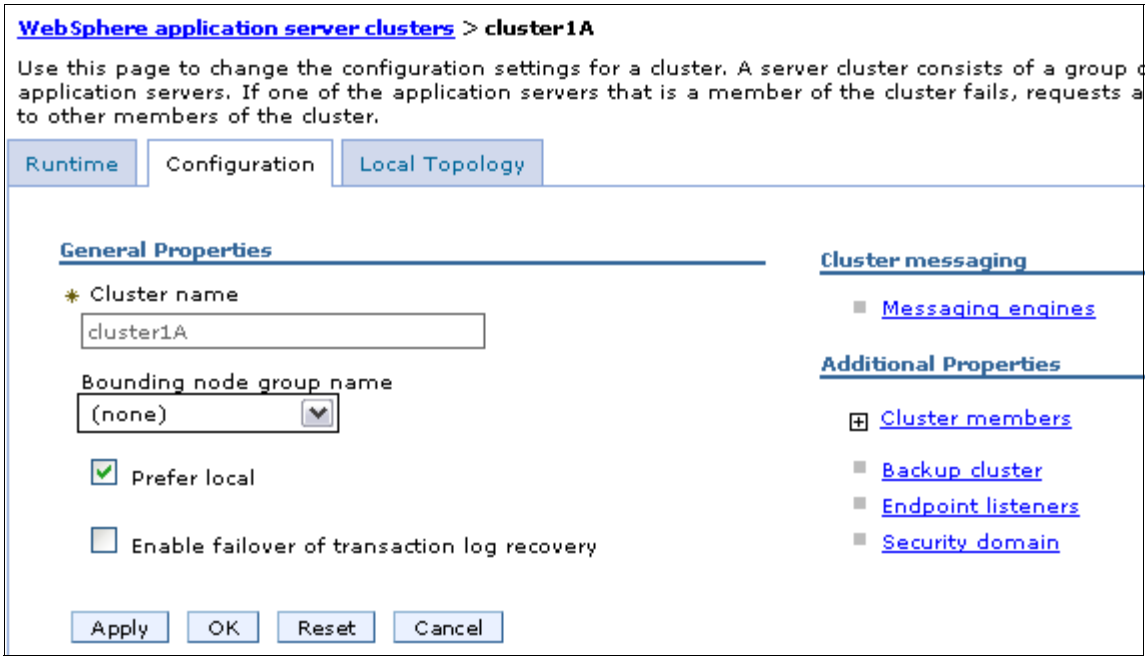

Figure 11 Editing an application server cluster properties

The detail page provides fields for configuring or viewing the more common settings and links to configuration pages for additional settings.

4. Click **OK** to save your changes to the workspace and exit the page. Click **Apply** to save the changes without exiting. The changes are still temporary. They are only saved to the workspace, not to the master configuration. This still needs to be done.

5. As soon as you save changes to your workspace, you will see a message in the Messages area reminding you that you have unsaved changes. See Figure 12.

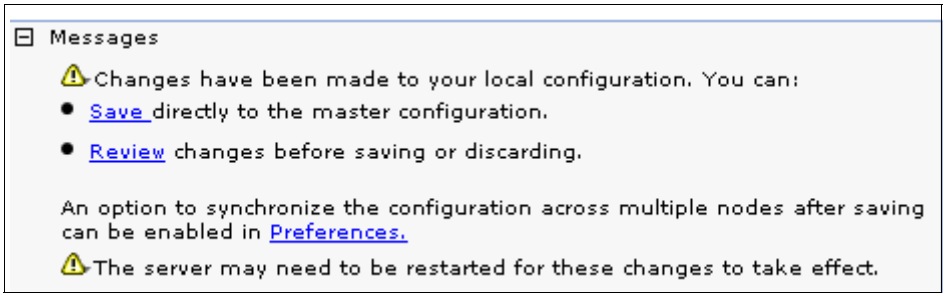

Figure 12 Save changes to the master repository

At intervals during your configuration work and at the end, you should save the changes to the master configuration. You can do this by clicking **Save** in the message, or by selecting **System administration** →**Save Changes to Master Repository** in the navigation tree.

To discard changes, use the same options. These options simply display the changes you have made and give you the opportunity to save or discard.

Adding new items

To create new instances of most item types (Figure 13 on page 23), complete these tasks:

1. Select the category and type in the navigation tree.
2. Select the **Scope**. (To create a new item, you cannot select the **All** option for scope.)
3. Click the **New** button above the collection table in the workspace.

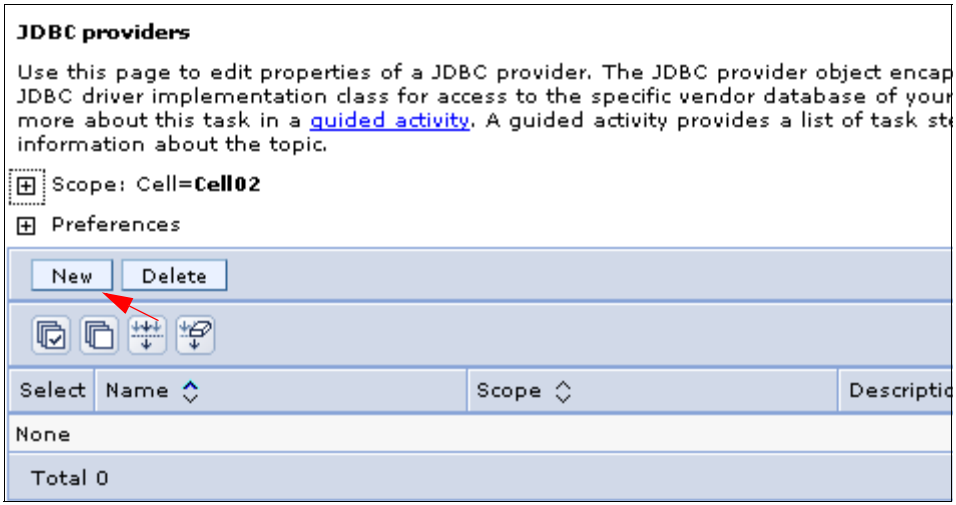

Figure 13 Create a new item

When you click **New** to add an item, one of two things will happen, depending on the type of item you are creating. A wizard will start to guide you through the definitions, or a new details page will open allowing you to fill in the basic details. In the latter case, enter the required information and click **Apply**. This will usually activate additional links to detail pages required to complete the configuration.

Note: In the configuration pages, you can click **Apply** or **OK** to store your changes in the workspace. If you click **OK**, you will exit the configuration page. If you click **Apply**, you will remain in the configuration page. As you are becoming familiar with the configuration pages, we suggest that you always click **Apply** first. If there are additional properties to configure, you will not see them if you click **OK** and leave the page.

4. Click **Save** in the task bar or in the Messages area when you are finished.

Removing items

To remove an item (Figure 14), complete these tasks:

1. Find the item.
2. Select the item in the collection table by checking the box next to it.
3. Click **Delete**.
4. If asked whether you want to delete it, click **OK**.
5. Click **Save** in the Messages area when you are finished.

For example, to delete an existing JDBC provider, select **Resources** → **JDBC** → **JDBC Providers**. Check the provider you want to remove and click **Delete**.

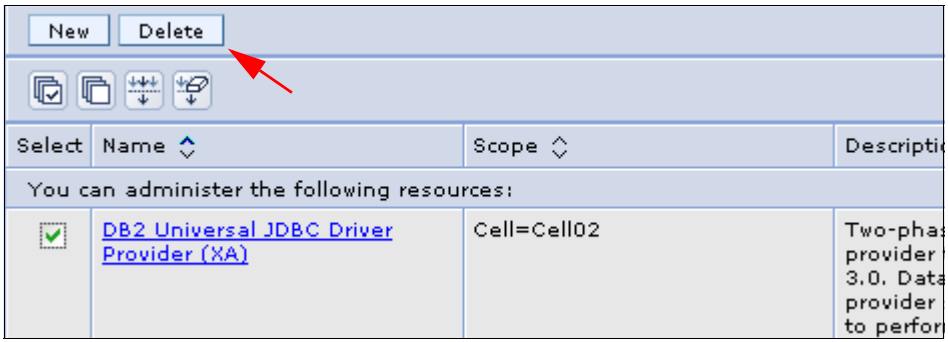

Figure 14 Deleting an item

Starting and stopping items

To start or stop an item using the console:

1. Select the category and type in the navigation tree.
2. Select the item in the collection table by checking the box next to it.
3. Click **Start** or **Stop.** The collection table shows the status of the item. See Figure 15 on page 25.

For example, to start an application server in a distributed server environment, select **Servers** → **Server Types** → **WebSphere application servers.** Place a check mark in the check box beside the application server you want and click **Start**.

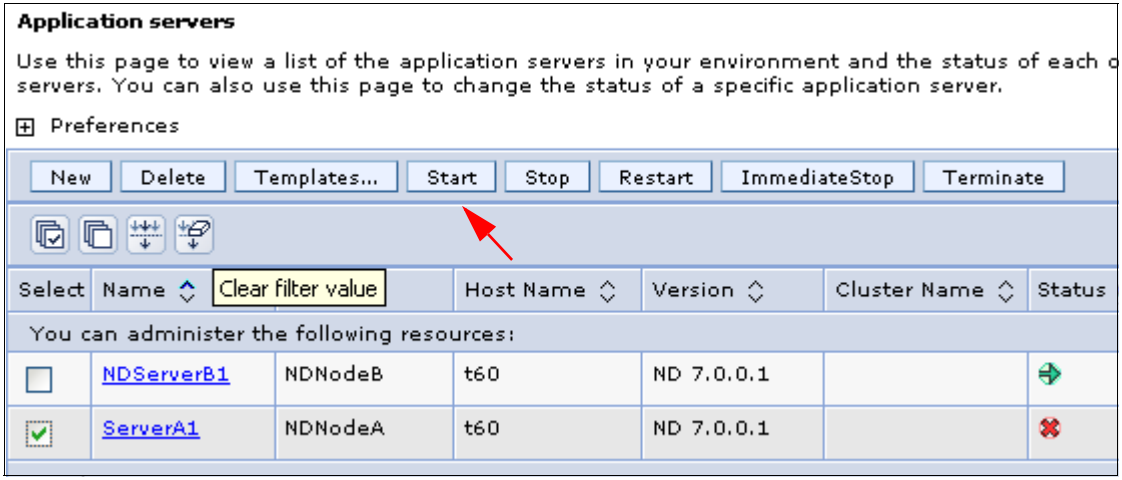

Figure 15 Starting and stopping items

Not all items can be started and stopped from the console. For example, the deployment manager and nodes must be started independently from the console. Also, there can be multiple options for starting and stopping an item (restart, stop immediate, and so forth).

Using variables

WebSphere variables are name and value pairs used to represent variables in the configuration files. This makes it easier to manage a large configuration.

To set a WebSphere variable:

1. Click **Environment** →**WebSphere Variables**. See Figure 16 on page 26.

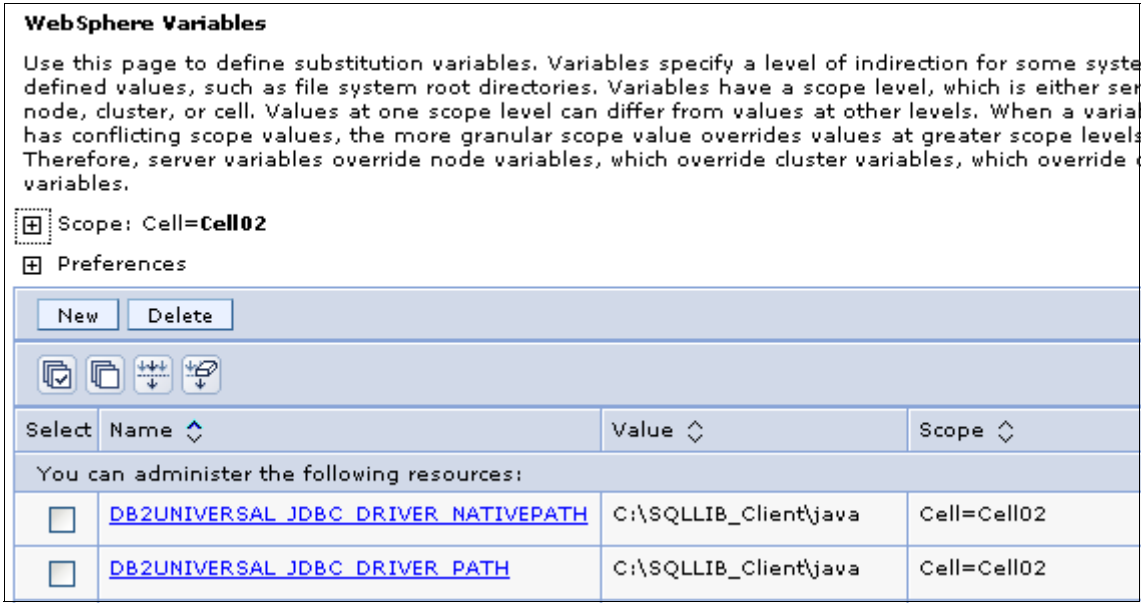

Figure 16 WebSphere variables

2. To add a new variable, click **New**, or click a variable name to update its properties.
3. Enter a name and value and click **Apply**. See Figure 17.

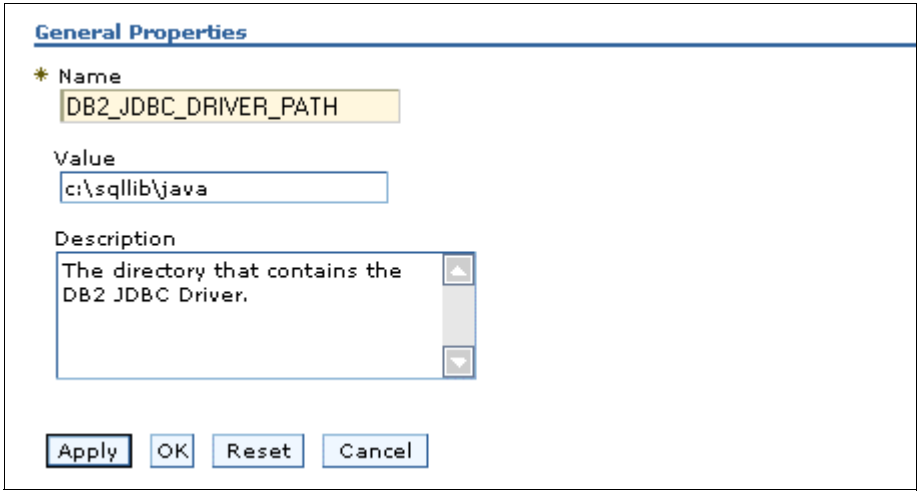

Figure 17 New WebSphere variable

Saving work

As you work with the configuration, your changes are saved to temporary workspace storage. For the configuration changes to take effect, they must be saved to the master configuration. If you have a distributed server environment, a second step is required to *synchronize*, or send, the configuration to the nodes. Consider the following possibilities.

If you work on a page, and click **Apply** or **OK**, the changes are saved in the workspace under your user ID. This allows you to recover changes under the same user ID if you exit the session without saving.

You need to save changes to the master repository to make them permanent. You have several options to save:

- Use the Save window in the Messages area. If this is displayed, it is the quickest method.
- Selecting **System administration →Save Changes to Master Repository**.
- When you log in, if you logged out without saving the changes, you will be given the option to save the changes.

The Save window presents you with the following options:

- Save
- Discard: This reverses any changes made during the working session and reverts to the master configuration.
- Cancel: This does not reverse changes made during the working session. It just cancels the action of saving to the master repository for now.
- Synchronize changes with nodes: This distributes the new configuration to the nodes in a distributed server environment.

Before deciding whether you want to save or discard changes, you can see the changed items by expanding **Total changed documents** in the Save window.

Important: All the changes made during a session are cumulative. Therefore, when you decide to save changes to the master repository, all changes are committed. There is no way to be selective about what changes are saved to the master repository.

Getting help

Help is available to you in several different ways:

- Click **Help** on the administrative console banner. This opens a new Web browser with online help for the administrative console. It is structured by administrative tasks. See Figure 18.

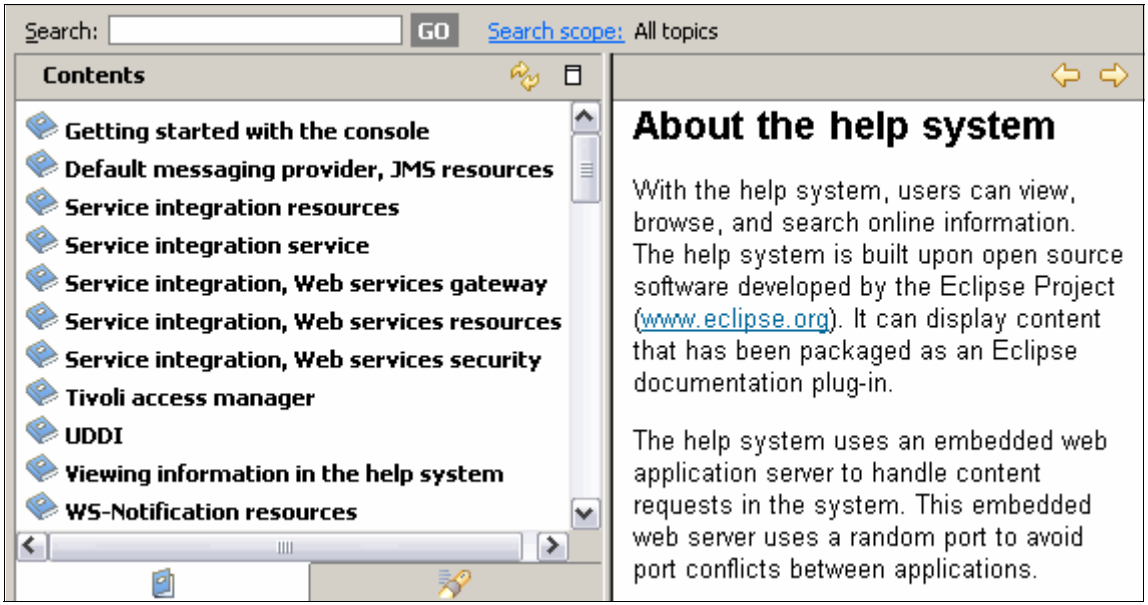

Figure 18 Online help

- With the option **Show the help portlet** enabled, you can see the Help window in the workspace. Click **More information about this page**. This will open the help system to a topic-specific page.
- The Information Center can be viewed online or downloaded from:

  `http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.home.doc/welcome.html`

Securing the console

WebSphere Application Server provides the ability to secure the administrative consoles so only authenticated users can use them. Note that enabling administrative security does not enable application security.

Before enabling any type of security for a production system, you should be very familiar with WebSphere security and have a plan for securing your WebSphere environment. Security encompasses many components, including administrative security, application security, infrastructure security, and specialized resource security options. This section only provides an overview of administrative security.

The first decision you have to make is to select the user registry you will use. If you enable security when you create a profile for distributed systems, a file-based registry is automatically created and populated with one administrative user ID. On z/OS® platforms, you will have the option of using the file-based registry or the z/OS system's SAF-compliant security database.

While a file-based user registry is not a good choice for securing applications, you can federate additional registries to the existing file-based registry to manage users and groups for application security.

If you are using a registry other than the WebSphere Application Server federated user registry, you must create at least one user ID to be used for the WebSphere administrator.

And while you might have heard about security domains that have been introduced in WebSphere Application Server V7, these domains are used for application security (not administrative security).

Before implementing security in a production environment, be sure to consult the *WebSphere Application Server V7 Security Guide*, SG24-7660.

Enabling security after profile creation

You can enable administrative security after profile creation through the administrative console by navigating to **Security** → **Global security**. Doing this allows you more flexibility in specifying security options. You will need to complete the configuration items for authentication, authorization, and realm (user registry). You will also need to populate the chosen user registry with at least one user ID to be used as an administrator ID (Figure 19 on page 30).

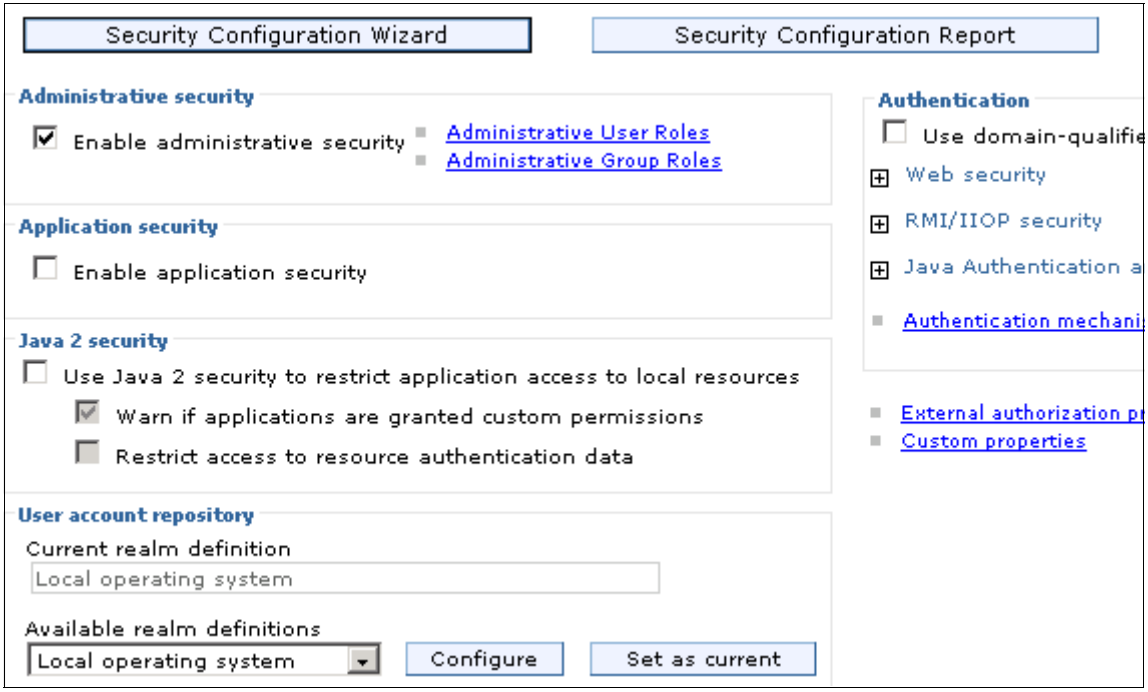

Figure 19 Enabling administrative security

Attention: Be aware that when you check the box to enable administrative security, the application security and Java™ 2 security check boxes are enabled automatically. If you are not prepared to use Java 2 or application security at this time, be sure to uncheck the boxes.

Tip: If you enable administrative security, and then find you cannot login, you can disable the security manually by editing the security.xml profile. This allows you to go back through the security configuration to see where things went wrong.

1. Open the security.xml file at *dmgr_profile_home*/config/cells/*cell_name*
2. Edit the second line, changing enabled=”true” to enabled=”false”.

```
<security:Security xmi:version="2.0"
xmlns:xmi="http://www.omg.org/XMI"
xmlns:orb.securityprotocol="http://www.ibm.com/websphere/appserve
r/schemas/5.0/orb.securityprotocol.xmi"
xmlns:security="http://www.ibm.com/websphere/appserver/schemas/5.
0/security.xmi" xmi:id="Security_1" useLocalSecurityServer="true"
useDomainQualifiedUserNames="false" enabled="false"
cacheTimeout="600" issuePermissionWarning="false"
activeProtocol="BOTH" enforceJava2Security="false"
enforceFineGrainedJCASecurity="false" appEnabled="false"
dynamicallyUpdateSSLConfig="true" allowBasicAuth="true"
activeAuthMechanism="LTPA_1"
activeUserRegistry="WIMUserRegistry_1"
defaultSSLSettings="SSLConfig_1">
```

After saving the configuration, you must restart the application server in a stand-alone server environment or the deployment manager in a distributed server environment.

The next time you log in to the administrative console, you must authenticate with the user ID that was identified as having an administrative role. Entering commands from a command window will also prompt you for a user ID and password. You can add additional administrative users and assign authorization levels from the administrative console.

Administrative security roles

Administrative security is based on identifying users or groups that are defined in the active user registry and assigning roles to each of those users. When you log in to the administrative console or issue administrative commands, you must use a valid administrator user ID and password. The role of the user ID determines the administrative actions the user can perform:

Fine-grained administrative security (new in V7):

In releases prior to WebSphere Application Server Version 6.1, users granted administrative roles could administer all of the resource instances under the cell. With V6.1, administrative roles are now per resource instance rather than to the entire cell. Resources that require the same privileges are placed in a group called the authorization group. Users can be granted access to the authorization group by assigning to them the required administrative role within the group.

A cell-wide authorization group for backward compatibility: Users assigned to administrative roles in the cell-wide authorization group can still access all of the resources within the cell.

- Administrator:

  The administrator role has operator permissions, configurator permissions, and the permission required to access sensitive data, including server password, Lightweight Third Party Authentication (LTPA) password and keys, and so on.

- Configurator:

  The configurator role has monitor permissions and can change the WebSphere Application Server configuration.

- Operator:

  The operator role has monitor permissions and can change the runtime state. For example, the operator can start or stop services.

- Monitor:

  The monitor role has the least permissions. This role primarily confines the user to viewing the WebSphere Application Server configuration and current state.

- Deployer:

  The deployer role is only available for **wsadmin** users, not for administrative console users. Users granted this role can perform both configuration actions and runtime operations on applications.

- AdminSecurityManager:

  The AdminSecurityManager role is only available for **wsadmin** users, not for administrative console users. When using **wsadmin**, users granted this role can map users to administrative roles. When fine grained administrative security is used, users granted this role can manage authorization groups.

- Iscadmins:

  The iscadmins role has administrator privileges for managing users and groups from within the administrative console only.

Assigning administrative roles to users and groups

If you are using a file-based repository, you can add users and groups through the console by navigating to **Users and groups** → **Manage Users** or **Users and groups** → **Manage Groups**. Otherwise, the users and groups must be added to the user registry using the tools provided by the registry product.

Role assignments for users and groups are managed through the administrative console. Navigate to **Users and groups** →**Administrative User Roles** or **Users and groups** →**Administrative Group Roles**. Use these panels to assign an administrative role to a user or group.

(New in V7) Fine-grained security

In releases prior to WebSphere Application Server version V7, users granted administrative roles could administer all of the resource instances in the cell. WebSphere Application Server is now more fine-grained, meaning that access can be granted to each user per resource instance. For example, users can be granted configurator access to a specific instance of a resource only (an application, an application server or a node). The administrative roles are now per resource instance rather than to the entire cell.

To achieve the instance-based security or fine-grained security, resources that require the same privileges are placed in a group called the administrative authorization group or authorization group. Users can be granted access to the authorization group by assigning to them the required administrative role.

You can define groups of resources that will be treated collectively by navigating to **Security** → **Administrative Authorization Groups**. The resource instances which are added to an authorization group can be the following types:

- Cluster
- Node
- Servers, including application servers and Web servers
- Applications, including business level applications
- Node groups
- Assets

After the authorization group has been created, you can assign users or groups an administrative role for the authorization group.

Many administrative console pages have a preference setting that allows you to restrict the items that you can see to those that are valid for your authorization group level. The roles that you can choose from depend on the role of the user ID you are logged into the console with.

Job manager console

The job manager console has many of the basic options that you find in the administrative console, including global security settings, the option to add users and groups to the federated user repository, WebSphere variable settings, and others that are common to any administrative environment.

What is unique to the job manager console is the ability to submit jobs to nodes registered to it.

Figure 20 on page 35 shows a job manager console. The option selected in the Navigation tree is **Jobs** → **Nodes**. You can see in this example, that one application server node has been registered from an admin agent. Two deployment manager nodes are registered as well.

Figure 20 Job manager console - list of nodes

Groups of nodes: You can create “groups of nodes” that contain the nodes you will work with from the job manager (select **Jobs** → **Groups of nodes**). A group of nodes can be used as the target of administrative jobs.

When you submit a job, you can select one or more groups from a drop-down. The alternative is to type in the name of the node or use the Find feature to select each node. Using the Find feature takes several steps.

So, even if you do not plan to use multiple nodes as the target of a job, creating a group for each node allows you to easily select a node rather than typing it in or searching for it.

If you include multiple nodes in the group, beware that all the nodes have to have a common user ID and password. When you submit a job you only have one place where you can enter the user ID and password.

Submitting a job with the job manager

The job manager provides the following job types:

- Run a wsadmin script
- Manage applications
  - distributeFile
  - collectFile
  - removeFile
  - startApplication
  - stopApplication
  - installApplication
  - updateApplication
  - uninstallApplication
- Manage servers
  - createApplicationServer
  - deleteApplicationServer
  - createProxyServer
  - deleteProxyServer
  - createCluster
  - deleteCluster
  - createClusterMember
  - deleteClusterMember
  - configureProperties
- Manage the server runtime
  - startServer
  - stopServer
  - startCluster
  - stopCluster

Details on each of these job types can be found at:

http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/index.jsp?topic=/com.ibm.websphere.nd.doc/info/ae/ae/rxml_7jobtypes.html

Follow these steps:

1. Start the job manager and log in to the job manager console:

   `http://<job_manager_host>:9960/ibm/console`

2. To submit jobs, nodes must have already been registered with the job manager. To verify which nodes have been registered, expand **Jobs** in the navigation window, and select **Nodes**. If this is the first time using the job manager, you might not see all the nodes displayed. To refresh the view, enter * as the value for Node name and click **Find** (Figure 21).

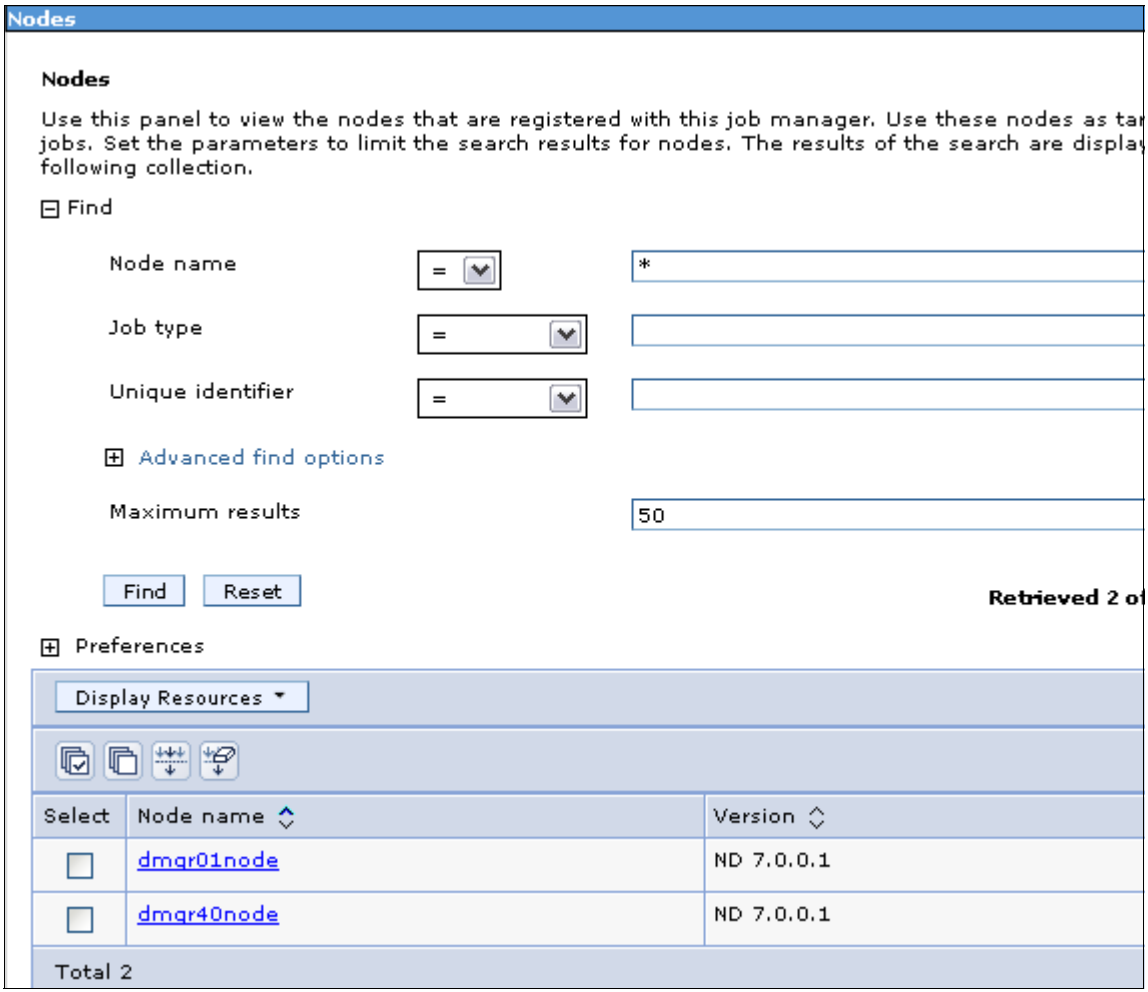

Figure 21 List of all nodes registered with job manager

3. Select **Jobs** → **Submit** to select the type of job to submit (Figure 22) and click **Next**.

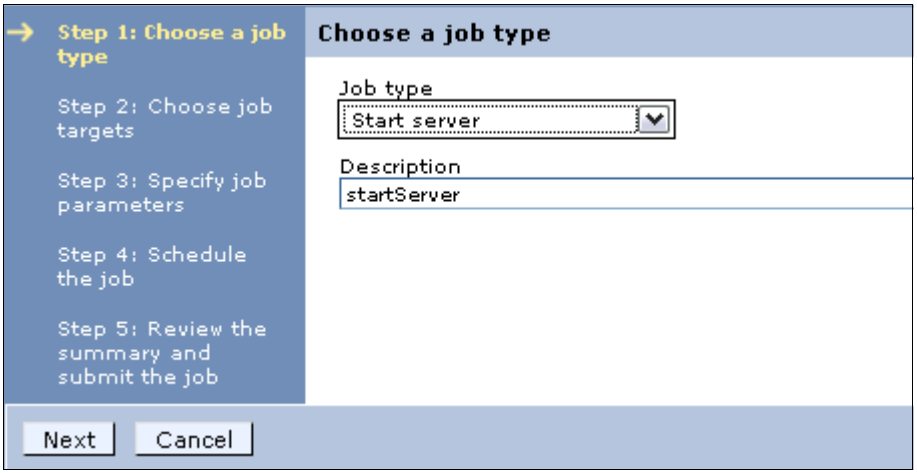

Figure 22 Select the job type

4. Select the node that you want to run the job on. You can select from a node group by using the drop-down next to the **Groups of nodes** field. Or you can select specific nodes.

Enter the user ID and password for the node that you will run the job against (Figure 23).

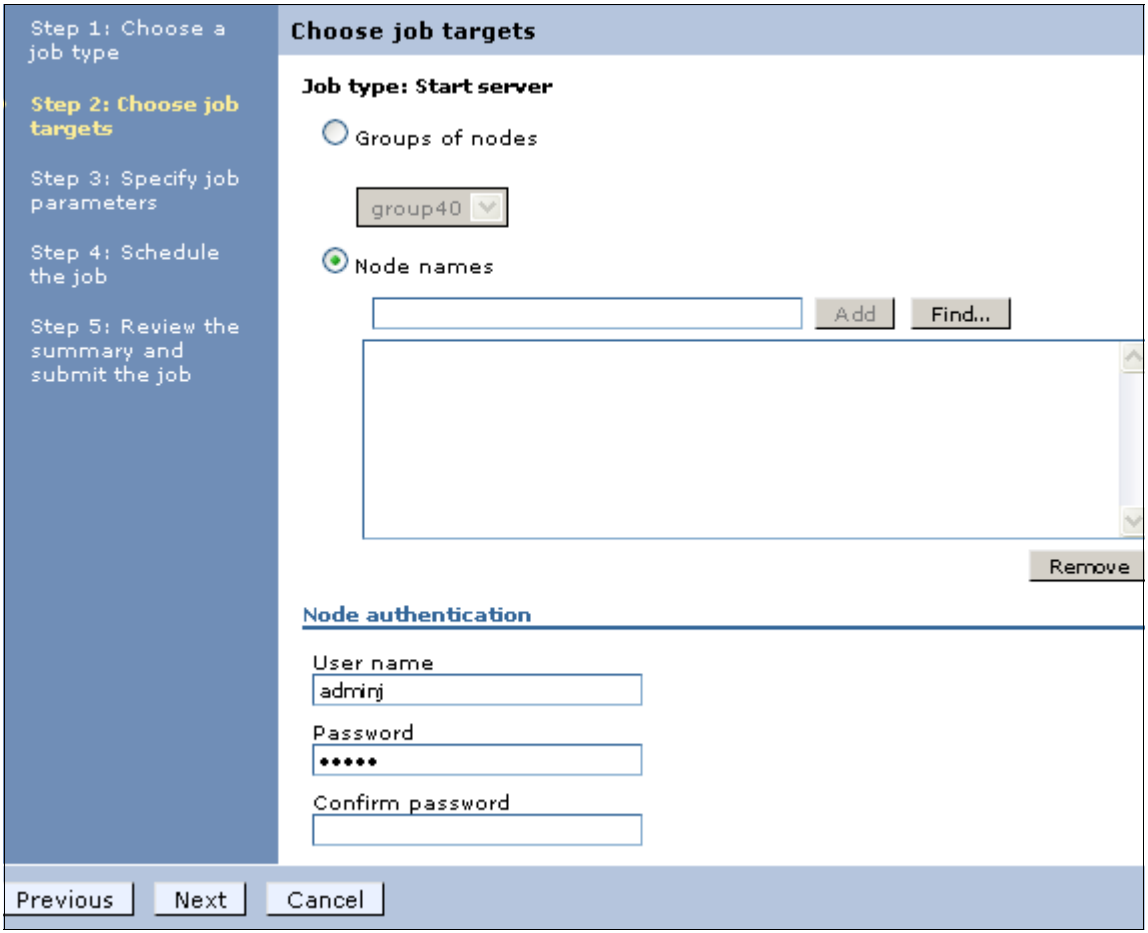

Figure 23 Select the job target

To use a specific node, select **Node names** and either enter the node name and click **Add**, or click **Find**. Using the Find option opens a new panel where you can search and select nodes (Figure 24 on page 40).

The simplest method of searching is to enter an "*" in the Node name field and click **Find**.

The list of nodes is shown in the Excluded nodes box. Select the nodes you want and use the arrow button to move them to the Chosen nodes box. You can hold the shift key down to select multiple nodes, or move them one at a time.

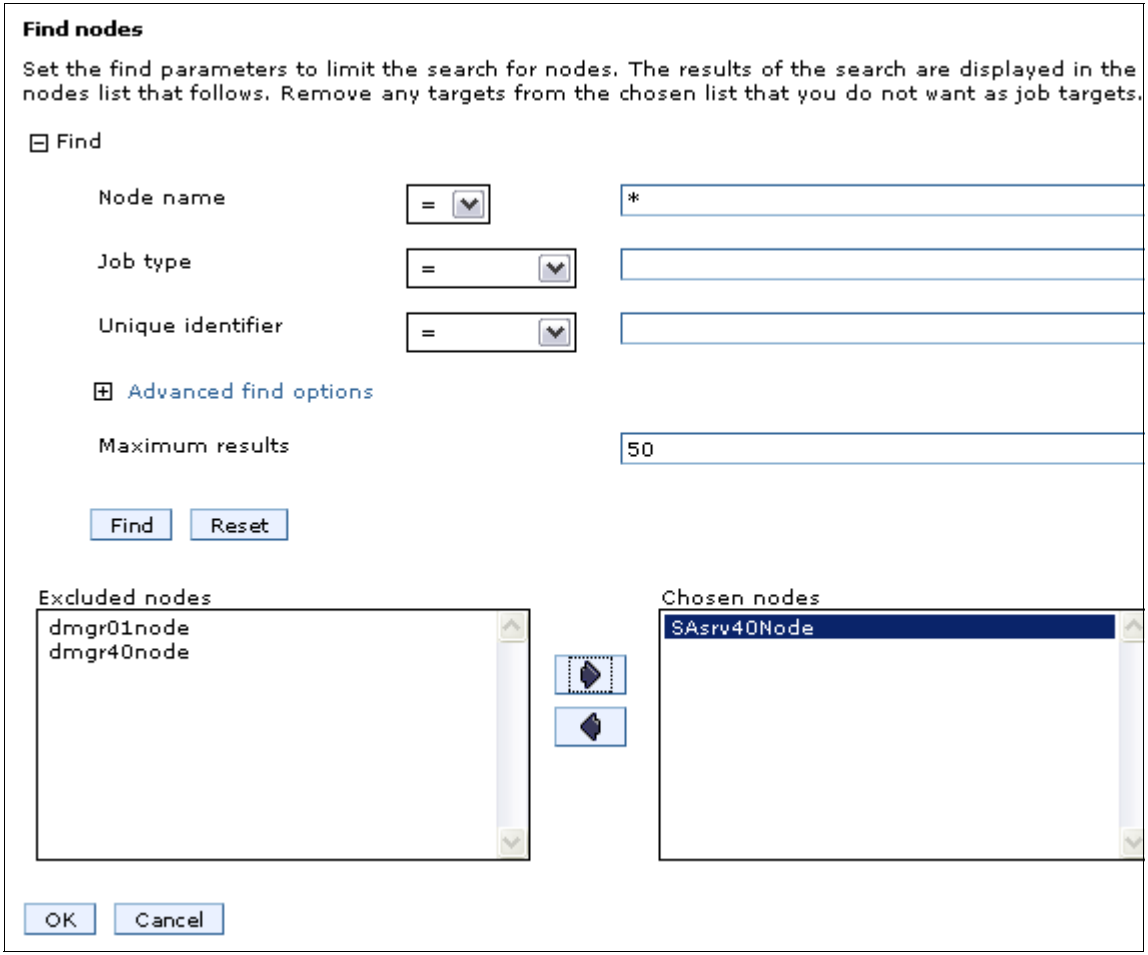

Figure 24 Search and select nodes

Click **OK**. This will return you to step 2 of the wizard with the node name filled in.

Click **Next** to continue the job submit process.

5. Specify the job parameters. These will vary widely depending on the type of job. The parameters provide the additional information the job will need to perform the task. For example, if you are running a job to start a server, you have selected the node in the previous step, but the server name must be entered as a parameter.

   Click **Next**.

6. The next step contains fields that specify how and when the job should run, and if a notification via e-mail should be sent (Figure 25 on page 41):

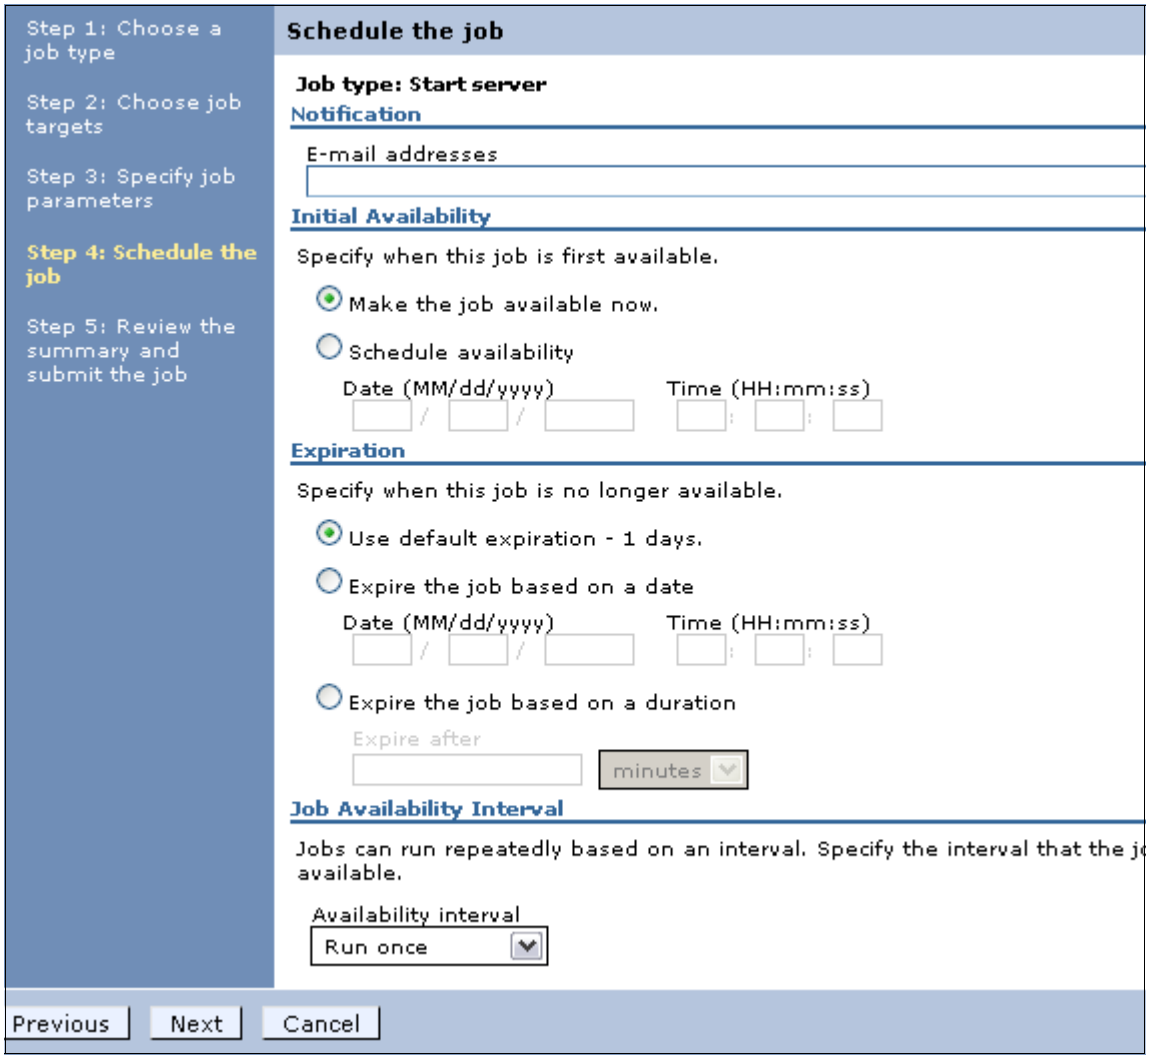

Figure 25 Specify job scheduling information

- Notification: The e-mail address specified will receive a notification when the job is finished. In order to use this field, you must configure a mail provider and mail session.
- Initial availability: You can make the job available now (it will run immediately after you have finished with the job submission process), or you can specify a date and time it will be available.
- Expiration: Specify an expiration date for the job.

– Job availability interval: This field allows you to repeat job submission at intervals. Depending on the selection, you will have an additional field displayed that allows you to choose the days, start and stop time, and so on (Figure 26).

Figure 26 Job interval options

If you select **Make this job available now** and **Run once**, the job runs right away and the Expiration settings have no meaning. The alternative is to set an Initial availability, Expiration date or duration, and select an interval that the job will run at.

7. Review the summary and submit the job.

   When a job is submitted from the job manager, the job details are saved in a database local to the job manager. The endpoint (deployment manager or administrative agent) pings the job manager at a predefined interval and fetches jobs that are to be executed. If the job submitted is a wsadmin job, the wsadmin script is executed. Otherwise, a corresponding job handler will execute the necessary admin code.

8. The Job status panel allows you to monitor the results. Use the Refresh icon in the Status summary column ( ) to update the status. The color in the Status summary field will indicate the success or failure of the job (Figure 27).

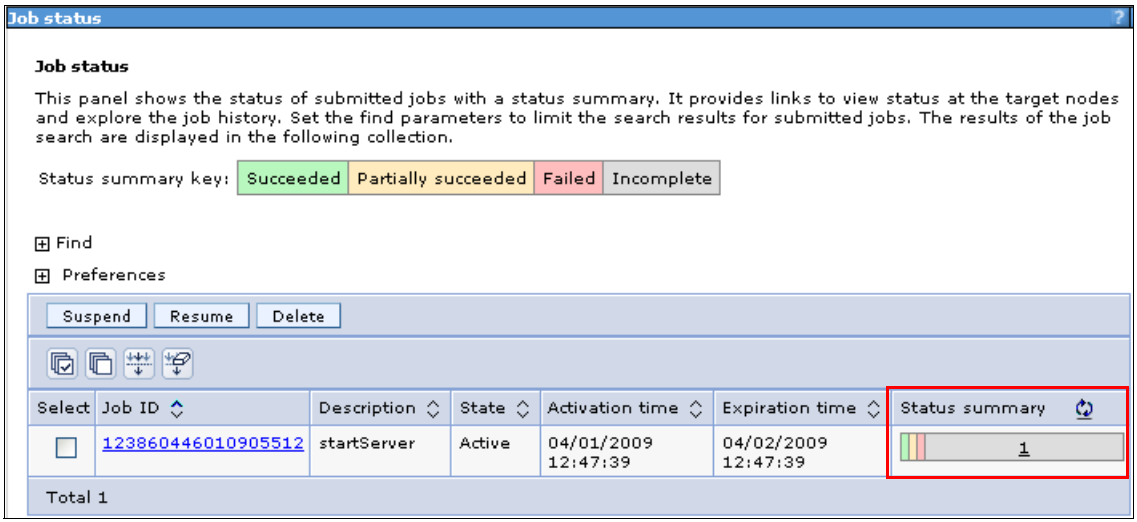

Figure 27 Status summary

9. Click the Job ID to see more information about the job (Figure 28).

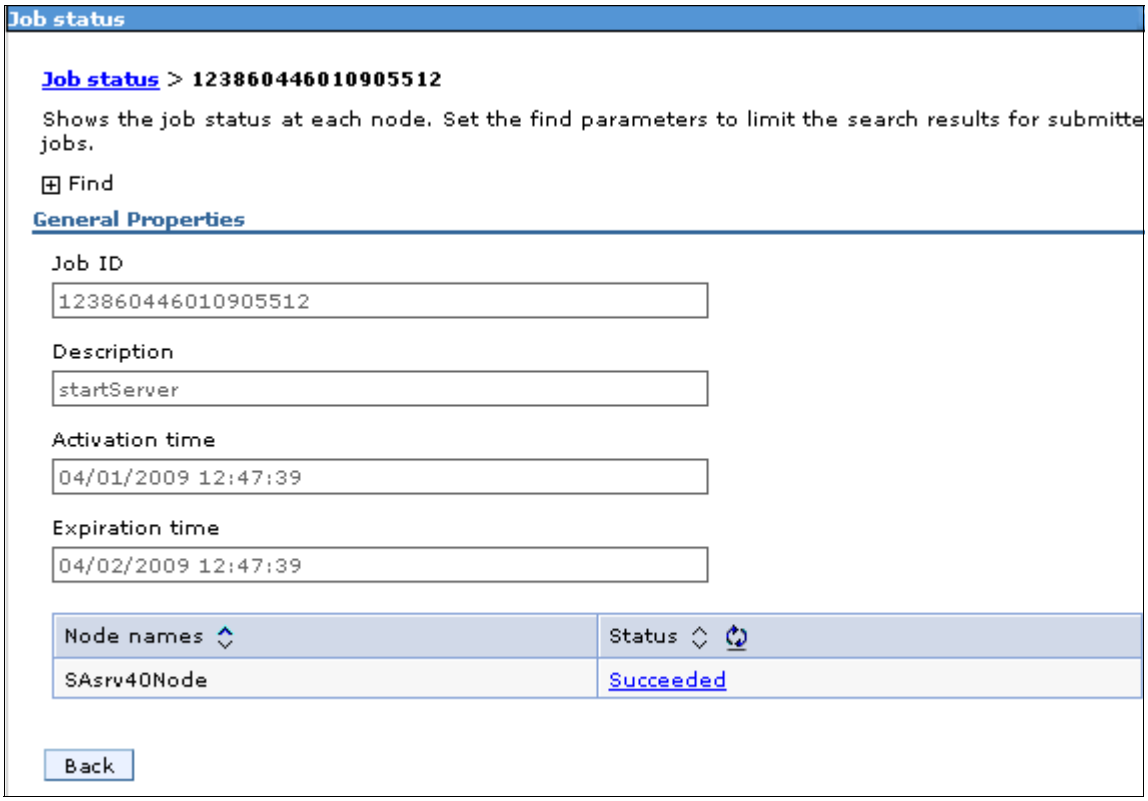

Figure 28 Job status information

Job status is always sent back to the job manager. Clicking the message in the Status column (Succeeded in this case) shows you additional information.

In the event of an error, you will see any messages produced by the job. Additional messages might be available in the logs for the server where the administrative action was to take place (Figure 29 on page 45).

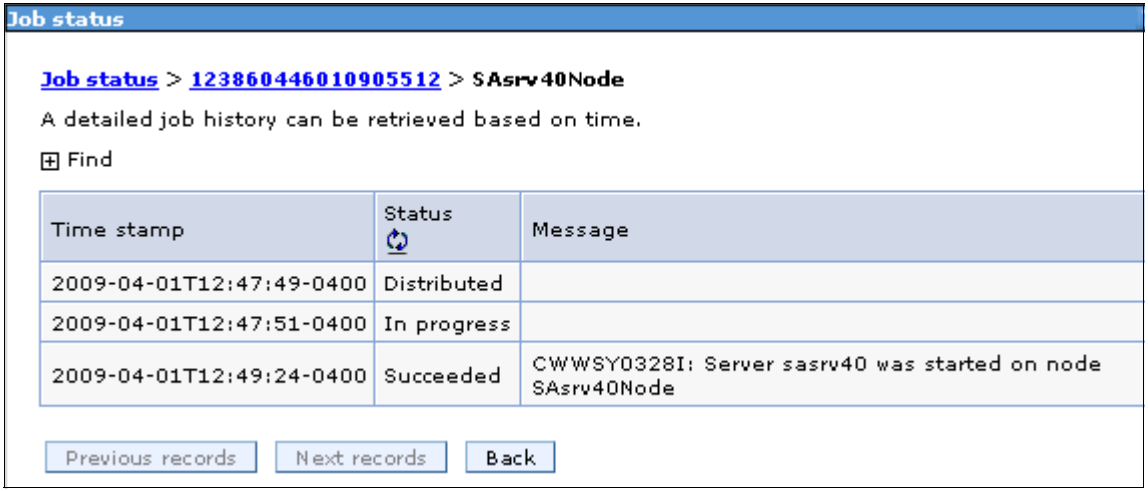

Figure 29 Job output

When you execute configuration type job (for example, create server) from the job manager to a deployment manager, the configuration will be saved if the job is successful. A job that submits a wsadmin script, however, will not save the configuration (the wsadmin script needs to do that).

Executing a job to a deployment manager does not cause node synchronization to occur. Synchronization will happen at the next automatic synchronization interval, or a wsadmin script can be submitted to synchronize.

Distributing files using the job manager

Some job types require that files be transferred to the node where the job will be run. The Distribute file job type can be used to transfer these files.

This is normally necessary in the following circumstances:

- When you want to run a wsadmin script on the node. The script must be distributed to the node before you can use the Run wsadmin script job.
- When you want to install or update an application. The EAR file must be distributed to the node before you can use the Install application or Update application jobs.

The following steps illustrate how to distribute a file to a node for use in later jobs. This example distributes a wsadmin script file to an admin agent:

1. The file to be distributed from the job manager must be in the /config/temp/JobManager directory of the job manager profile.

Create the ***jobmgr_profile_root*/config/temp/JobManager** directory and copy the file into it.

If you are developing a script or application in Rational Application Developer, you can export the file directly to the directory.

2. The distribute file job stores the file into the downloadedContent directory of the administrative agent or deployment manager profile. The destination parameter is relative to the downloadedContent directory. You must create this directory on the admin agent or deployment manager:
   - *adminagent_profile_home*/downloadedContent
   - *dmgr_profile_root*/downloadedContent
3. In the Job manager console, select the **Job** → **Submit** menu. This will launch the Job properties wizard.
   a. Select **Distribute file** as the job type and click **Next**.
   b. Enter the script file location on the job manager and the location to store the script file on the target node.

      In this example, the appInstall.py script was stored in the following location:

      jobmgr_profile_root/config/temp/JobManager/appInstall.py

      On the admin agent it will be stored as:

      adminagent_profile_home/downloadedContent/appInstall.py

      The arguments are entered as shown in Figure 30.

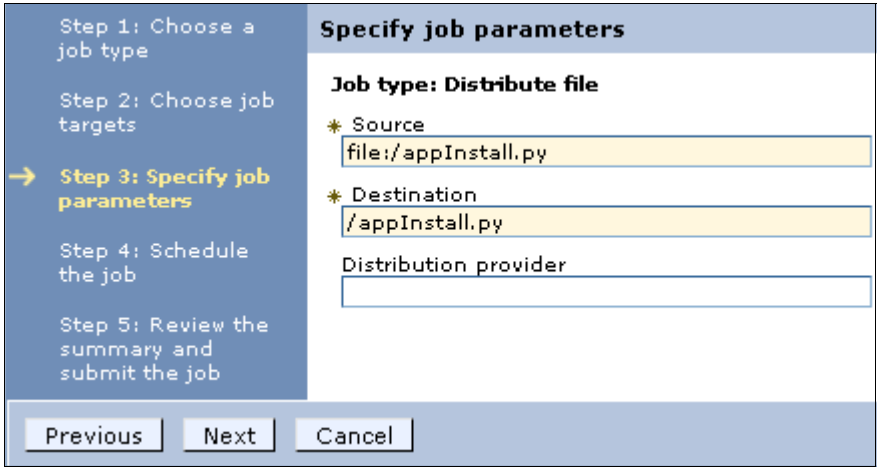

Figure 30 File distribution parameters

Click **Next**.

c. Take the defaults for the job schedule. The defaults will execute the distribute file job once. Click **Next**.

d. Click **Finish**. Monitor the status of the job and ensure it completes successfully.

Using command line tools

WebSphere Application Server provides commands that can be run from a command line. These commands can be used for many administrative tasks, for example, to start, view, or stop a WebSphere process. Many commands have an equivalent GUI interface, either created specifically for the command, through an administrative console, or through a First Steps console. However, it is often convenient to simply enter these commands manually from a command line.

Examples of commands are:

- startServer
- stopServer
- serverStatus
- manageprofiles
- addNode

Command location

Command line tools must be run on the system where the process you are entering the command for resides. For the most part, the commands exist in two places:

- *install_root*/bin

  Commands entered from this location will operate against the default profile unless you use the -profileName parameter to specify the profile.

- *profile_root*/bin

  Commands entered from this location will operate against the profile defined in *profile_root*.

Key usage parameters

The commands are consistent across platforms, though how you enter them, case sensitivity, and the extension will vary.

Note: Parameter values that specify a server name, a node name or a cell name are always case sensitive regardless of operating system.

There are several commonly used parameters that are valid for every command you should be aware of.

- **-profileName** specifies the profile the command is to run against
- **-username** specifies the user ID with the administrative privileges required to execute the command
- **-password** specifies the password for the user ID specified in -username
- **-help** will display the usage requirements and a list of parameters for the command

Entering commands

In this section we will show how commands can be entered on the various operating systems.

Commands in the paper: In the rest of this paper, the commands used were entered on Windows® operating systems in our test labs and will reflect that format. If you are using another platform, note that you will need to adjust the examples for your platform.

Windows operating systems

Commands in Windows operating systems have an extension of .bat. It is not necessary to use the extension. Commands are not case sensitive, though parameters and names are case sensitive.

To use a command:

1. Open a Command Prompt window.
2. Change to the directory where the command is.

   For example:

   ```
   C:\Program Files\IBM\WebSphere\AppServer\profiles\profile_name
   ```

3. Enter the command. For example:

```
serverStatus -all -username admin -password admin
```

Note: When running command line tools on the Microsoft® Windows Vista® operating system and Windows Server® 2008: On the Windows Vista operating system and Windows Server 2008, you can install WebSphere Application Server as either Administrator or non-Administrator. When it is installed as Administrator, certain operations (such as those involving Windows Services) require Administrator privileges.

In order to ensure that WebSphere Application Server command line tools have sufficient privileges, run them as Administrator. When you run these command line tools from a Command Prompt, run them from a Command Prompt window that is launched by performing the following actions:

1. Right-click a Command Prompt shortcut.
2. Click **Run As Administrator**.
3. When you open the Command Prompt window as Administrator, an operating-system dialog appears that asks you if you want to continue. Click **Continue** to proceed.

If you are using a Windows Server Core installation of Windows Server 2008, any WebSphere Application Server commands that require a graphical interface are not supported, because a Windows Server Core system does not have a graphical user interface. Therefore, commands such as pmt.bat or ifgui.bat are not supported on that type of Windows Server 2008 installation.

UNIX operating systems

Commands in UNIX® operating systems have an extension of .sh and are case sensitive.

To use a command:

1. Open a Command Prompt window.
2. Change to the directory where the command is.

   For example, for root users the directory would be:

   – (AIX®) `/usr/IBM/WebSphere/AppServer/profiles/`*`profile_name`*`/bin`

   – (HP, Linux®, Solaris™) `/opt/IBM/WebSphere/AppServer/profiles/`*`profile_name`*`/bin`

   For non-root users: *`user_home`*`/IBM/WebSphere/AppServer/profiles/bin`
3. Enter the command. For example:

   ```
   serverStatus.sh -all -username admin -password admin
   ```

i5/OS operating systems

For an i5/OS® system, proceed as follows:

1. From the i5/OS command line, start a Qshell session by issuing the **STRQSH CL** command.
2. Change to the directory where the command is.

   For example:

   ```
   cd /QIBM/UserData/WebSphere/AppServer/V7/ND/profiles/profilename/bin
   ```

3. Enter the command. For example:

   ```
   serverStatus -all -username admin -password admin
   ```

z/OS operating systems

You can manage application servers on an z/OS system from a UNIX System Services environment as follows:

1. Enter **uss** (to switch to the UNIX System Services environment)
2. Change to the directory where the command is.

   On z/OS, this will be always be *app_server_root*/profiles/default because only the profile name “default” is used in WebSphere Application Server for z/OS.

3. Enter the command. For example:

   ```
   startServer.sh server1 -user username - password password
   ```

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 give 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.

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.

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.

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 illustrates 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. 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 IBM's application programming interfaces.

This document REDP-4573-00 was created or updated on October 13, 2009.

Send us your comments in one of the following ways:
- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbook@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099, 2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks