


IBM WebSphere Application Server V8

AddNode-asExistingNode Option



© 2012 IBM Corporation

This presentation describes a new option for the addNode command that simplifies topology management tasks in IBM WebSphere Application Server V8.

Table of contents

- Overview
- Usage scenarios
- Summary
- References

This presentation describes a new option for the addNode command that can be used to simplify the management of Network Deployment cells.

You are provided with a brief overview of this new capability and then the presentation describes some scenarios where it can be applied.

Overview

This section will discuss an “AsExisting” option.

What is addNode -asExistingNode?

- A new capability in WebSphere Application Server V8 that can be used to:
 - quickly recover a damaged node
 - move a node to a different machine or a different operating system,
 - provision a cell

The addNode command was created to federate stand-alone WebSphere Application Server nodes to a Network Deployment cell. WebSphere Application Server Version 8.0 provides a new option, “AsExistingNode” for adding nodes to a cell. This option, subject to restrictions, can be used to quickly recover a damaged node, move a node to a different machine – potentially with a different operating system, or to provision a cell.

Usage scenarios

Following scenarios describe how this new command option can be used.

How is a node recovered with pre-Version 8 releases?

- Assuming a backup is not available, you must:
 1. Stop the deployment manager
 2. Make a backup of deployment manager configuration
 3. Use removeNode on the deployment manager to remove the node to be recovered
 4. Reinstall WebSphere Application Server
 5. Create a profile for a stand-alone application server node
 6. Use addNode to add the application server node the deployment manager
 7. Restore the backed up configuration to the deployment manager
 8. Run the syncNode command

First, if you are using a version of Network Deployment earlier than version 8, how do you recover a damaged node?

Of course, if the environment were backed up before the damage it can be restored using a backup/restore facility. If however, no backup is available, you must:

Stop the deployment manager.

Make a backup of the deployment manager configuration.

Remove the node to be recovered from the deployment manager using the removeNode command

Reinstall WebSphere Application Server

Create a profile for a stand-alone application server node

Use the addNode command to add the stand-alone application server node to the deployment manager

Restore the backed up configuration of the deployment manager

Run the syncNode command.

Recover a damaged node

- **Scenario:**
 - Recover a damaged node when no back up is available
- **Procedure:**
 1. Reinstall WebSphere Application Server at the same directory location as previously installed
 2. Create a profile with the same node name and path
 3. Recover the damaged node using `addNode -asExistingNode`

In version 8.0, recovery of a damaged node is greatly simplified with this new command option.

As before, assume that the configuration had not been backed up before damage to the node.

First reinstall WebSphere Application Server to the same directory location where it had been installed previously.

Then create a stand-alone application server profile using the same node and path as the damaged node.

Finally, run the `addNode -asExistingNode` to recover the damaged node.

Move node to new machine, at same path

- **Scenarios:**

- Move to a new machine with same operating system version as the current node
- Move to a new machine with a different operating system with a compatible path name. For example, move a node between different versions of Windows, or between UNIX variants

- **Restrictions:**

- Applications that use Scheduler only work with the same host name.
- Nodes cannot be moved between z/OS and non-z/OS systems

- **Procedure:**

1. Delete the original profile without using removeNode
2. On the deployment manager, change the host name of the node using the AdminTask.changeHostName(...) command
3. Reinstall WebSphere Application Server on the new machine at same directory location
4. Create a profile with the same node name and path
5. Run the addNode -asExistingNode command

With this new command option version 8.0 now supports the ability to move a node to a different machine. This can be done if the machine is running the same operating system as the current node's machine or if the new machine is running a different operating system but with a compatible path name.

In other words, with this procedure you can move a node between different Windows versions or operating system products, or between different UNIX operating systems. However you cannot move a node between Windows and UNIX or between z/OS and machines running a distributed operating system with this procedure.

One additional restriction is that the applications using the scheduler require that host name remain the same.

To move a node to a different machine but with the same path:

First: Do not use removeNode, but delete the original profile.

Then: On the deployment manager, use the AdminTask changeHostName command to change the node's host name to the new host name

Next: Install or reinstall WebSphere Application Server on the new machine at the same directory location as on the original host.

Next: Create a stand-alone application server with same node name and path as on the original host.

Finally: Run the addNode -asExistingNode command.

Move a node to a new operating system, or a different path

▪ Scenarios:

- Move to a new machine with the same operating system, but a different host name and path
- Move to a new machine with a different operating system but with compatible configuration files, e.g., between AIX and Windows or between two z/OS systems

▪ Restrictions:

- Applications that use Scheduler must have the same host name
- Nodes cannot be moved between Z/OS and non-Z/OS machines

▪ Procedure:

1. Delete the original profile without using `removeNode` (that is, delete the profile directory using operating system commands)
2. On the deployment manager, change the host name of the node: using the `AdminTask.changeHostName(...)` command
3. On the deployment manager, change the WebSphere Application Server install path and profile path in the variable maps
4. Reinstall WebSphere Application Server on the new machine, possibly at a different path
5. Create a profile with the same node name, possibly at a different path, with the new host name
6. Run the `addNode -asExistingNode` command

9

AddNode-AsExistingNode Option

© 2012 IBM Corporation

You can also move a node to a machine running a different operating system or to a different path on the destination machine.

The restrictions are similar to other scenarios. Applications using the scheduler require the same host name. Nodes cannot be moved between z/OS and other, non-z/OS machines. However, nodes can be moved between Windows and UNIX or Linux operating system machines as long as the configuration files are compatible.

The task is also similar to the more restrictive move operation:

First: without using the `removeNode` command, delete the profile of the node you are moving.

Then: On the deployment manager, use the `AdminTask changeHostName` command to change the node's host name to the new host name

Next, there is an additional step to handle the change in path between the original machine and the destination machine –

On the deployment manager, in the variable maps, change the WebSphere Application Server install path and profile path in the variable maps.

Then: Install or reinstall WebSphere Application Server on the new machine at the installation path specified on the deployment manager

Next: Create a stand-alone application server profile with the path specified on the deployment manager

Last: Run the `addNode -asExistingNode` command.

Provision environments

- **Scenario:**
 - Quickly create new cells from a template cell
- **Restrictions:**
 - The cell name cannot be changed
 - The Scheduler application does not work with multiple environments
 - Need to assess whether different resources, e.g., data sources, are required for each environment
- **Procedure:**
 1. Create and configure a cell
 2. Make a copy of the deployment manager profile
 3. For every new environment to be provisioned:
 1. Install WebSphere Application Server
 2. Create deployment manager and node profiles
 3. Restore the deployment manager profile
 4. Customize the deployment manager's configuration. For example, change:
 - host name, ports
 - WebSphere install and profile paths
 - security settings
 5. Run `addNode -asExistingNode` for each node. It can be run concurrently from each node

If your environment is to include multiple network deployment cells with similar configurations, this new `addNode` command can be used to simplify the provisioning of these cells.

This task is subject to these restrictions:

The provisioned cells must have the same cell name as the original, template cell.

The scheduler system is restricted to a single cell.

If different resources are required for different cells, these will need to manually configured.

To provision an environment with multiple cells:

First create and configure a cell to serve as a template.

Next make a copy of the deployment manager profile.

Then for each cell to be created:

Install WebSphere Application Server on each of the machines that are hosting WebSphere Application Server deployment manager and application server nodes.

Create deployment manager and node profiles on these machines

Restore the deployment manager profile on the deployment manager machine

Customize the deployment manager configuration as needed. For example change host name and ports, install paths, profile paths, security settings and so forth.

Finally, run `addNode - asExistingNode` on each node to be federated to the deployment manager. Note that it can be run concurrently from each node.

Summary

This section provides a summary of this presentation.

Summary

- The addNode command is enhanced with a new option, asExistingNode, that can be used to:
 - Recover a damaged node
 - Move nodes from one machine to another
 - Provision a cell with nodes
- These uses are subject to these restrictions:
 - Nodes can only be between machines running similar operating systems
 - Applications that rely on the Scheduler subsystem cannot be used on nodes that were moved to machines with different host names
 - Cell provisioning requires that the cell name is not changed, and is subject to restrictions on the use of the Scheduler and resource dependencies

In summary, the new `–asExistingNode` option for the `addNode` command simplifies the recovery of damaged nodes and the provisioning of cells. It also enables nodes to be moved from one machine to another. Some restrictions apply. Movement between some operating systems, use of the scheduler system by applications, and cell naming is limited in some scenarios.

References

AddNode command in the WebSphere Application Server Version 8.0 IBM Knowledge Center:

http://www-01.ibm.com/support/knowledgecenter/SSAW57_8.0.0/com.ibm.websphere.nd.doc/info/ae/ae/rxml_addnode.html

For additional information on this feature see the addNode command description in the Information Center.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, AIX, WebSphere, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Windows, and the Windows logo are registered trademarks of Microsoft Corporation 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.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.