



IBM Software Group

IBM® WebSphere® Application Server V6

Building Network Deployment Cell

Add/Remove Nodes to the Cell



@business on demand.

© 2005 IBM Corporation
Updated May 5, 2005

This presentation will cover building a WebSphere Network Deployment cell by adding and removing nodes.

Goals

- Describe how to build a WebSphere Application Server V6 Network Deployment Cell
- Pre-requisite
 - ▶ WebSphere Application Server V6 Network Deployment - System Management Architecture
 - ▶ WebSphere Profiles



The goal of this presentation is to describe the process of adding and removing nodes. In version 6, these processes are very similar to version 5. You should already be familiar with the architecture, including the concept of WebSphere Profiles.

Section

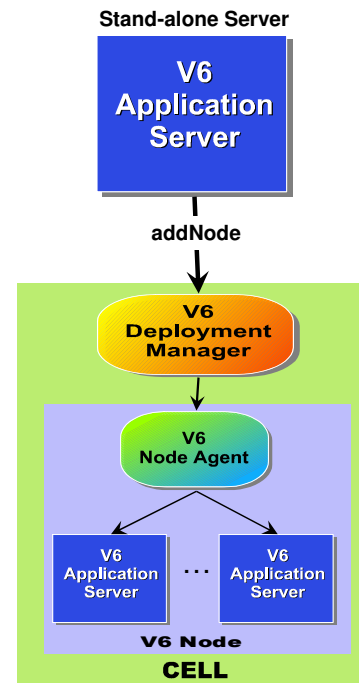
Adding (Federating) V6 Stand-alone Application Server to a Cell



The first section will discuss federating, or adding a Stand-alone Application Server (referred to as Stand-alone Node) to a cell. At this time, a new V5 Node cannot be federated to a V6 Network Deployment cell. However, a V5 node can exist in a mixed-node V6 Cell as part of a migration strategy. More details on the mixed node cell is provided in the Migration presentation in the Install and Migration section.

What Is Needed to Build a Cell

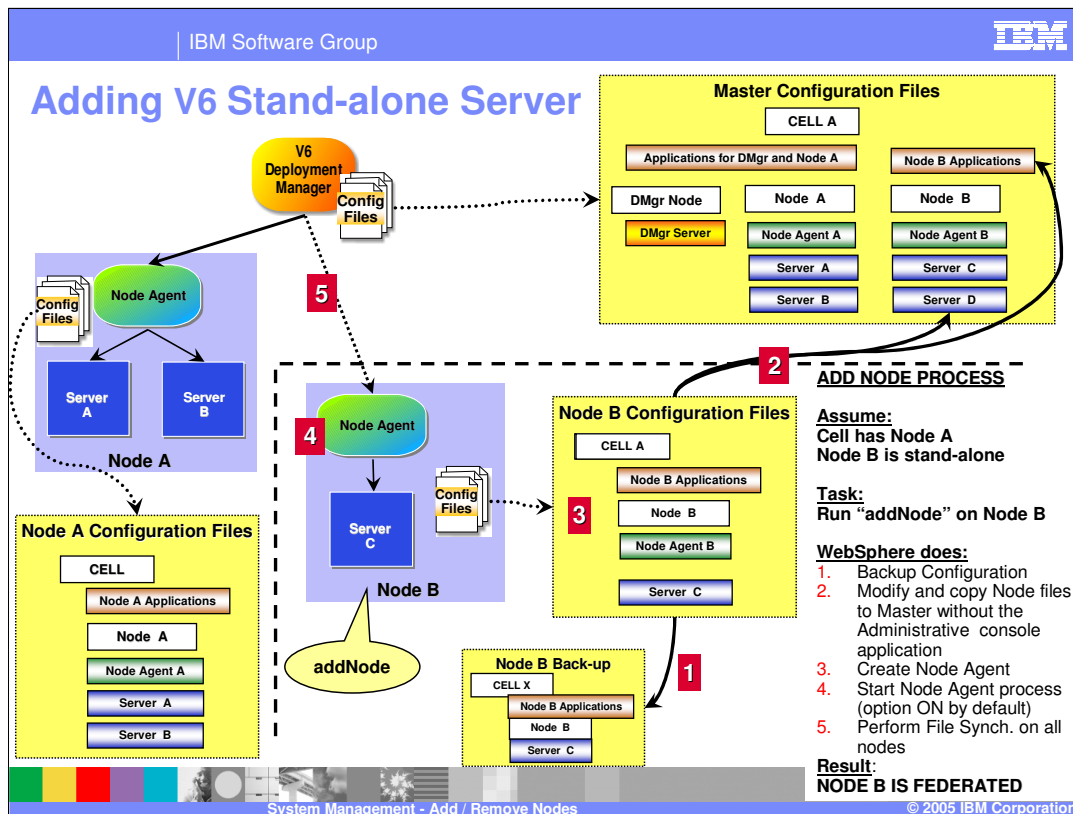
- Before you begin, you will need
 - ▶ V6 Deployment Manager configuration
 - ▶ V6 Stand-alone Application Server configuration
- The DMgr and Stand-alone server configuration can be:
 - From different WebSphere installs
 - On different physical Machines



To add a node to a cell, you need two things; A node to add, and a cell to add it to. The Deployment Manager is the administrative focal point of the cell. You can federate any stand-alone node with that deployment manager. The node can be from a different installation or on a different host system, but it must be a z/OS node.

The process of adding a node can be initiated on the command line, from the Administrative Console or by running the BBOWADDN job that is created when you run the Federate ISPF Dialog option. Of the three, the BBOWADDN method is the easiest. A job (BBOWMNAN) to federate your 'empty managed node' is also created for you when using the ISPF Dialog option to Create an Empty Managed Node. Both of these jobs are issuing the addNode command for you with the correct parameters that you supply.

The illustration shows that the application server is the active party in this process, reaching out to pull itself into the cell. Because the process is initiated from a batch or shell script, the application server does not have to be running before you add it.



In this illustration, a stand-alone application server installation called Node A is already federated into the cell. The slide shows the process of federating an additional node (Node B) into the cell. The addNode utility is a command line utility that is run from the bin directory on the Stand-alone Application Server. On z/OS, the BBOWADDN job is created to run this utility when using the ISPF Dialog Panels for federation. The utility performs the following steps:

First, it backs up the current configuration files for the stand-alone application server.

Second, all the configuration files are modified to fit the Cell model and then sent to the deployment manager by the addNode utility. The deployment manager sorts through the configuration files, adding most of them to the master repository. However, the cell level files are not integrated as the cell definition from the stand-alone application server no longer applies and the cell definition for the network deployment environment will now apply to Node A.

Third, it creates the configuration files needed to define a node agent for this node.

By default, the addNode utility launches the node agent in step 4 and the node is fully federated into the cell.

Finally, the deployment manager and the addNode utility perform a full file synchronization to update Node B.

You can disable the node agent startup in step 4 by using the "-noagent" option.

Adding V6 Stand-alone Server: More details

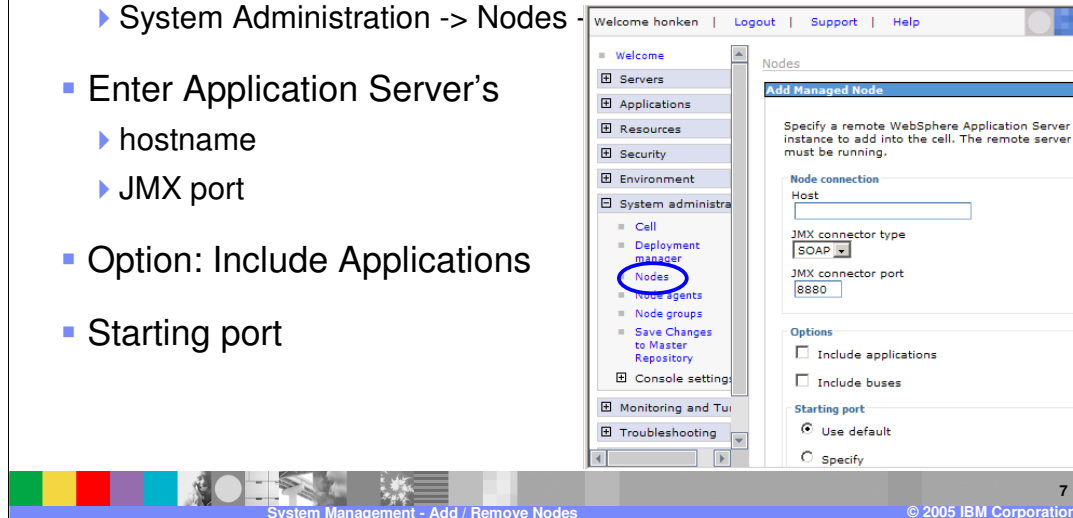
- Different ways to add Node to a Cell
 - ▶ From the Deployment Manager Administrative Console
 - ▶ By running the command line “addNode” utility
 - ▶ By running the BBOWADDN job that is created by the ISPF Dialog Federate option
 - ▶ By running the BBOWMNAN job that is created by the ISPF Dialog Empty Managed Node option



The next few slides will detail the use of the Administrative Console and the Command line utility to add a node to a cell. One of the new features of version 6 is the option to immediately add a node to a cell during the creation of a empty managed node. This was discussed in the Profile presentation.

Adding V6 Stand-alone Server from DMgr

- Application server must be running before it can be added
- In the Administrative Console navigate:
 - ▶ System Administration -> Nodes
- Enter Application Server's
 - ▶ hostname
 - ▶ JMX port
- Option: Include Applications
- Starting port



When adding a node from the Administrative Console, the application server you wish to add must be running. In the Administrative Console, navigate to the Nodes and click **Add Node**. Enter the host name of the application server system and confirm the default JMX Connector type and port. An option to include the applications from the Stand-alone server is provided. Any application that is already present in the cell will not be added. A warning message will appear in the log file.

You also have the option to specify what you would like your starting port to be. Adding the node through the Administrative Console will most likely result in you having to modify port values to avoid conflicts once the node becomes part of the cell since the ports will be selected for you.

“addNode” Command line tool

- “addNode” command should be run from a stand-alone Application server
 - ▶ `./addNode.sh <dmgr_host> [dmgr_port] [options]`
- By default applications are not added to the cell, unless “-includeApps” option is used
 - ▶ When adding an application, if an application with the same name already exists in the cell, a warning is displayed and the application is not installed
- Adding to a secure Cell requires appropriate Administrative user id and password
- New option “-portprops <file>” can list the specific ports to be used by the addNode utility



The addNode utility is run from the bin directory of the stand-alone application server installation and it connects to the host and port of the deployment manager of the cell it is joining. The host name is required, while the port is optional if the default SOAP port is used. The default SOAP port for a deployment manager is 8879.

There are several options that can be used with addNode, which can be seen by invoking addNode with the dash question mark option.

A new option allows you to specify the ports in an external file for the federated node to use. For example:

```
BOOTSTRAP_ADDRESS=9001
```

```
SOAP_CONNECTOR_ADDRESS=9002
```

```
ORB_LISTENER_ADDRESS=9003
```

Any ports not specified will be assigned default values, which will be unique across the current installation, although not guaranteed unique on the system.

“BOWADDN” job to federate node

```

----- WebSphere Application Server for z/OS Customization -----
Option ==> █

Define Variables for Federate stand-alone Application Server Node (1 of
Specify the following to customize your server, then press Enter
to continue.

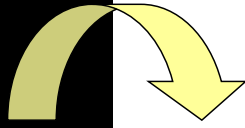
WebSphere Application Server home directory:
/etc/s6cellA
/AppServer
Deployment Manager Access
Node host name.....: MVS215
JMX SOAP port.....: 58879
Deployment manager security is enabled: N
User ID.....: V6ADMIN

Include Apps.....: Y
Application server's ORB port.: 59810
Node group name.....: DefaultNodeGroup

Node Agent Definitions
Server name (short)...: S6AGNT1
Server name (long)...: nodeagent
JMX SOAP connector port.....: 58878
Node Discovery port.....: 57272
Node Multicast Discovery port.: 55000
High availability manager communication port: 59355

ORB listener host name...: *
ORB port.....: 52810
ORB SSL port.....: 0
  
```

User specifies unknown values in the ISPF Dialog



```

----- WebSphere Application Server for z/OS Customization -----
Option ==> █

Define Variables for Federate stand-alone Application Server Node (2 of
Specify the following to customize your server, then press Enter
to continue.

Do you wish to federate service integration buses
that exist on this node?" (Y/N): Y
  
```

The BOWADDN job is created from the values you supply in the ISPF Dialog panel for the federate option. This gives you more control over port values. You are given the option of whether you want to include the applications (includeapps) and whether you want to include any service integration buses that have been created.

“BBOWADDN” job...continued

- First thing the job does is copy the port values you specified into a file in your /tmp directory (see BBOWANPF in your configuration .DATA PDS for values being used)
- addNode.sh command is issued with the –portprops option specified:

```
//*****
//* Run addNode.sh
//*****
//BBOWAN EXEC PGM=IKJEFT01,REGION=0M
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
BPXBATCH SH +
  /etc/s6cellA+
  /AppServer+
  /bin/addNode.sh +
  MVS215 +
  58879 +
  -includeapps +
  -includebuses +
  -nodegroupname DefaultNodeGroup +
  -portprops /tmp/portFile_82180.props +
  -nodeagentshortname S6AGNT1 +
  1> /tmp/bbowaddn_82180.out +
  2> /tmp/bbowaddn_82180.err
```

10

The job that you run is created from the values you specify on the previous slide. The port values are copied over to an HFS file and specified using the –portprops option. This is a much cleaner way to federate than using the Administrative Console because you have the opportunity to specify the port values you want in advance.

Changes After Federating Node

- Cell level information overrides:
 - ▶ Virtual Host information, Security information, etc.
- Applications will not be migrated
 - ▶ unless the “-includeapps” option is used
- Administrative point changes to Deployment Manager
 - ▶ Administrative Console application is removed
 - ▶ wsadmin scripts should be changed to reflect the new cell name
- HTTP Plug-in xml file location for the local Web Server (if present) should be changed to point to the Plug-in file in the Deployment Manager master configuration
- The federated node inherits global security settings of the Deployment Manager Cell



Some things to remember when federating an application server into a cell:

Cell level configurations on the application server are overwritten by the cell level configuration from the deployment manager. Applications are not included by default. However, you can use the

-includeapps option to bring applications into the cell.

The administrative location changes; once an application server is added to a cell, the Administrative Console application is removed, and you use the cell Administrative Console on the deployment manager to make changes to the federated application server.

Finally, if there is an HTTP server on the same system with the application server, the configuration should be changed to reference to the plug-in file for the cell.

When adding a secured node to an unsecured cell, the node will become unsecured as well. When adding an unsecured node to a secured cell, the node will become secured as well. When adding a secured node to a secured cell, the node will continue to be secured, but with the security settings for the cell.

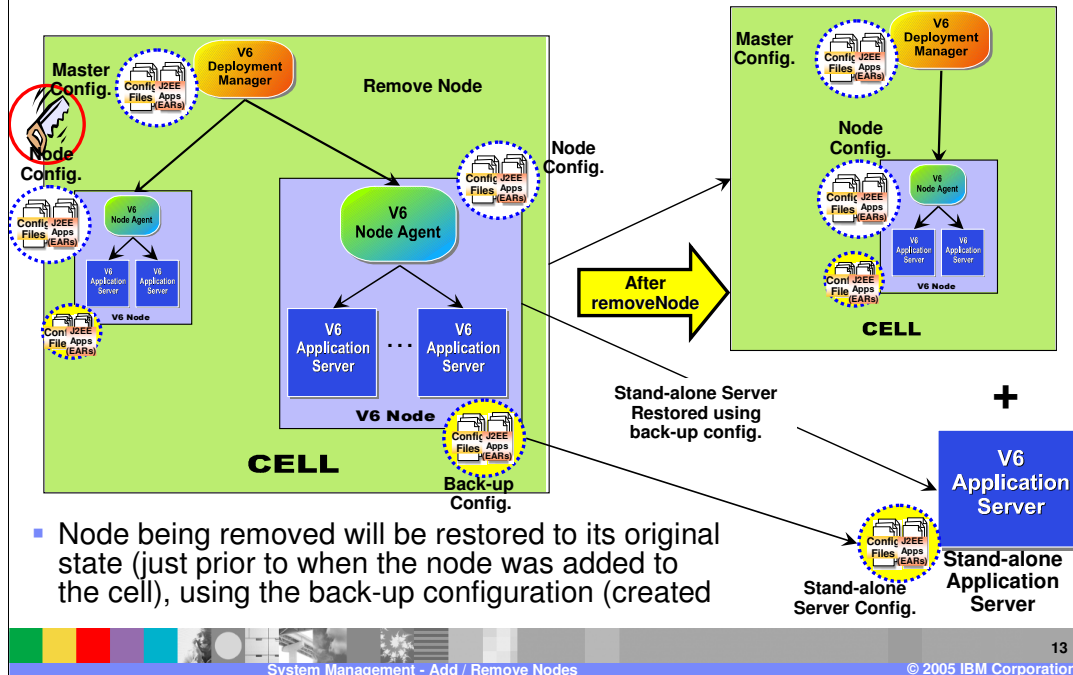
Section

Removing V6 Nodes from a Cell



This section covers removing nodes from the cell.

Remove V6 Managed Node from a Cell



This page illustrates the process of removing a node. The **removeNode** command is executed on the node being removed, as shown in the diagram.

The **removeNode** command line utility removes a node from the cell configuration and restores it to the state just prior to when the node was added to the cell. During the addition of the node, the **add node** utility created a backup of the configuration. The **remove node** utility uses that backup configuration to restore the configuration to its original state.

Remove Node can be called directly from the command line to allow you to restore the original application server environment. The configuration that was backed up during **addNode** will be restored.

Since you go back to the original application server configuration, you lose any configuration changes that were made while the node was federated.

Remove Federated V6 Node from a Cell

- Run `./removeNode.sh` command from the Node
 - ▶ Will need the administrator user ID and password, if security is ON
- From the Deployment Manager Administrative Console
 - ▶ System Management -> nodes -> select a Node and click remove Node
 - ▶ Remote Node Agent must be running
- Application Servers in the Node are stopped during the remove node process



There are two ways to remove a node from a cell. One uses the remove node command and is executed on the node being removed. The other way is through the Deployment Manager administrative console in the nodes section.

If security is enabled, the remove node command requires an administrator user id and password.

Since you revert to the original application server configuration, you lose any configuration changes that were made while the node was federated.

During the removal of the nodes, any running Application Servers will be stopped, since they will be restored to their original state.

Section

Problem Determination

The next section will address problem determination.

Problem Determination and Log files

- Deployment Manager must be running to:
 - ▶ Add or Remove a node
- Node Agent must also be running to remove a node using the Administrative Console
- Log files will be in the Node profile's log directory
 - ▶ <WAS_PROFILE>/logs/addNode.log
 - ▶ <WAS_PROFILE>/logs/removeNode.log
- Problems with Adding/removing a node
 - ▶ Requires participation from the Node and DMgr ,
 - ▶ Look at the job log of the DMgr and Node Agent, along with the addNode or removeNode log files



When troubleshooting a failed addNode or removeNode operation, first check that the deployment manager is up and running. If removing a node through the administrative console, the node agent on the node being removed must be running as well.

Next, verify that the systems can resolve host names to IP addresses in both directions. You can test this by pinging each system by host name from the other system.

Check the log files, located in the logs directory of the node being operated on.

Problem Determination

- Node names in the cell must be unique
- Cell name of the node being federated must be different from the name of the cell it is being added



Node names in a cell must be unique. Keep in mind that the deployment manager also has a node name, so a node with the same name cannot be added to the cell.

A stand-alone application server has a node name and a cell name, but they are transparent in single-server operation. However, when being added to a cell, the cell name and the node name matter and cannot conflict with existing names.

Summary

- Discussed the several ways to add a Stand-alone Application Server to a Cell
- Discussed what happens to the federated node
- Discussed how to remove Node from the cell

This presentation covered the process of adding and removing nodes.

Trademarks, Copyrights, and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
eIogo business	DB2	Series	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel 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.

Linux is a registered trademark of Linus Torvalds.

Other company, product and service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprocessing in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2004. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.