



IBM Software Group

IBM® WebSphere® Application Server V6

Building Network Deployment Cells:

Add/Remove Nodes to the Cell



@business on demand.

© 2005 IBM Corporation
Updated May 13, 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

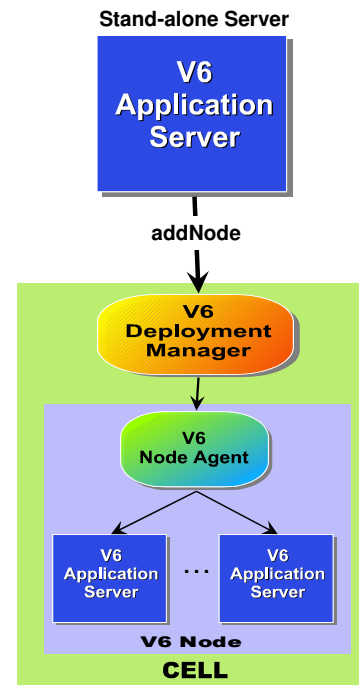
Adding (Federating) V6 Stand-alone Application Server to a Cell



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

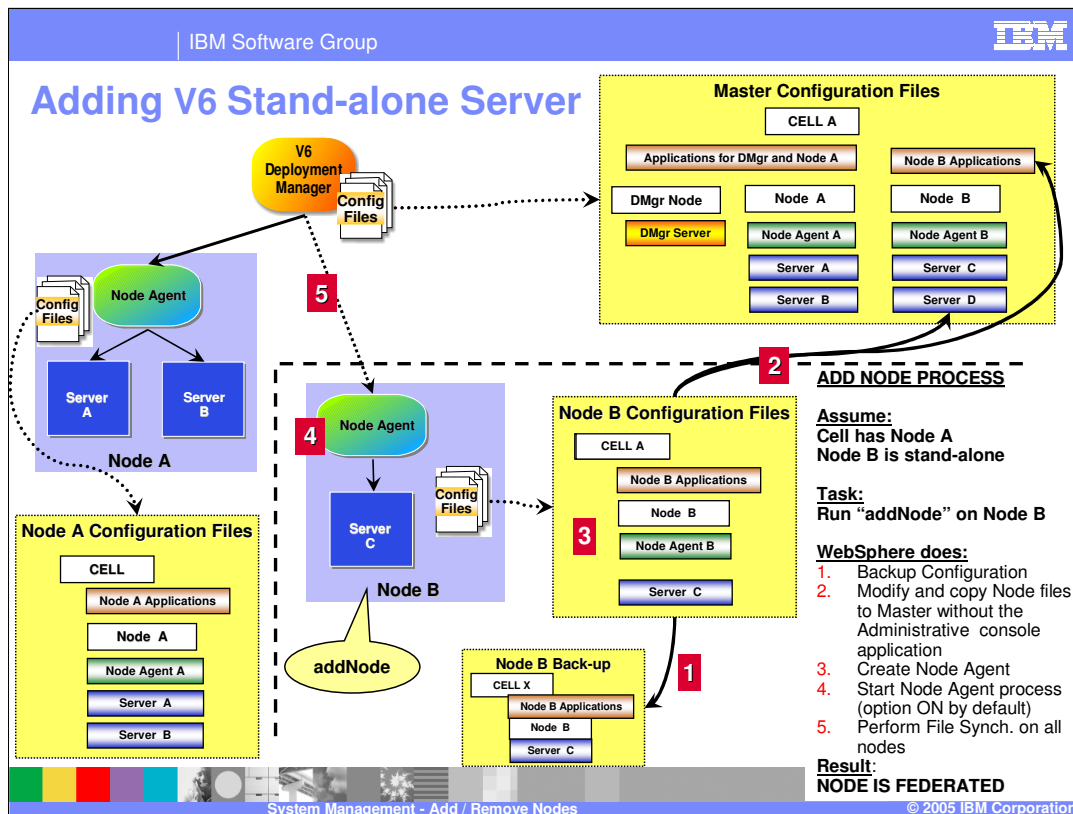
- Before you begin, you will need
 - ▶ V6 Deployment Manager profile
 - ▶ V6 Stand-alone Application Server profile
- The DMgr and Stand-alone server profiles can be:
 - From different WebSphere installs
 - On different physical Machines
 - In different distributed Operating Systems (like Linux, AIX, Windows and so on)



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 profile is the administrative focal point of the cell; you can federate any stand-alone node with that deployment manager. That node can be from a different installation, or on a different host system, even on a different operating system.

The process of adding a node can be initiated on the command line, or within the Administrative Console, or you can create and federate a node in a single step when creating a profile.

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 need to be running before you add it.



In this illustration, a base application server installation called Node A is federated into the cell. The addNode utility is a command line utility that is run from the Base Application Server's bin directory. The utility performs the following steps:

First, it backs up the current configuration files for the base 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 base 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 now, the node is now fully federated into the cell.

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

You can turn off the starting of the node agent in step 4, by using the "-noagent" option.

Adding V6 Stand-alone Server: More details

- Different ways to add Node to a Cell
 - ▶ From Deployment Manager Administrative Console
 - ▶ Command line “addNode” utility
 - ▶ When creating a Custom profile (managed profile), you can federate to a DMgr as part of the profile creation
 - Available only with the Network Deployment package
- Add Node operation now contains a port conflict resolution algorithm
 - ▶ Looks at the ports used by all the profiles of the current installation only

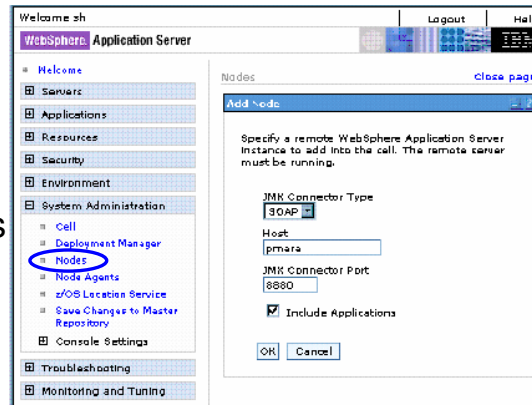


The next few slides will detail using the Administrative Console and the Command line utility to add a node to a cell. One of the new features of version 6 is the option to immediately add a node to a cell during the creation of a Custom profile. This was discussed in the Profile presentation.

New for version 6, the add node process has port conflict resolution algorithm, where it looks at the ports defined across all the profiles of the current installation. The algorithm does not look at other WebSphere installation, therefore this might not be sufficient. Another option is where you supply the ports through an external file, which is discussed later.

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 -> AddNode
- Enter Application Server's
 - ▶ hostname
 - ▶ JMX port
- Option: Include Applications



When adding a node from the Administrative Console, the application server you want 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 machine, 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 System out log file.

“addNode” Command line tool

- “addNode” command should be run from a stand-alone Application server
 - ▶ addNode.bat(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
- Windows™ services for Node Agent can be added
 - ▶ -registerservice option
- 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.

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 an application server is federated into a cell: Cell level configurations within the application server are overwritten by the cell level configuration of the cell. Applications are not included by default; you can use the dash include apps option to bring applications into the cell.

The administrative location changes; once an application server is added into a cell, the Administrative Console application is removed, and you use the Administrative Console of the cell to make changes to the application server that has become a node in the cell.

Finally, if there is an HTTP server on the machine where the application server resides, then the configuration should be changed to point to the plug-in file for the cell.

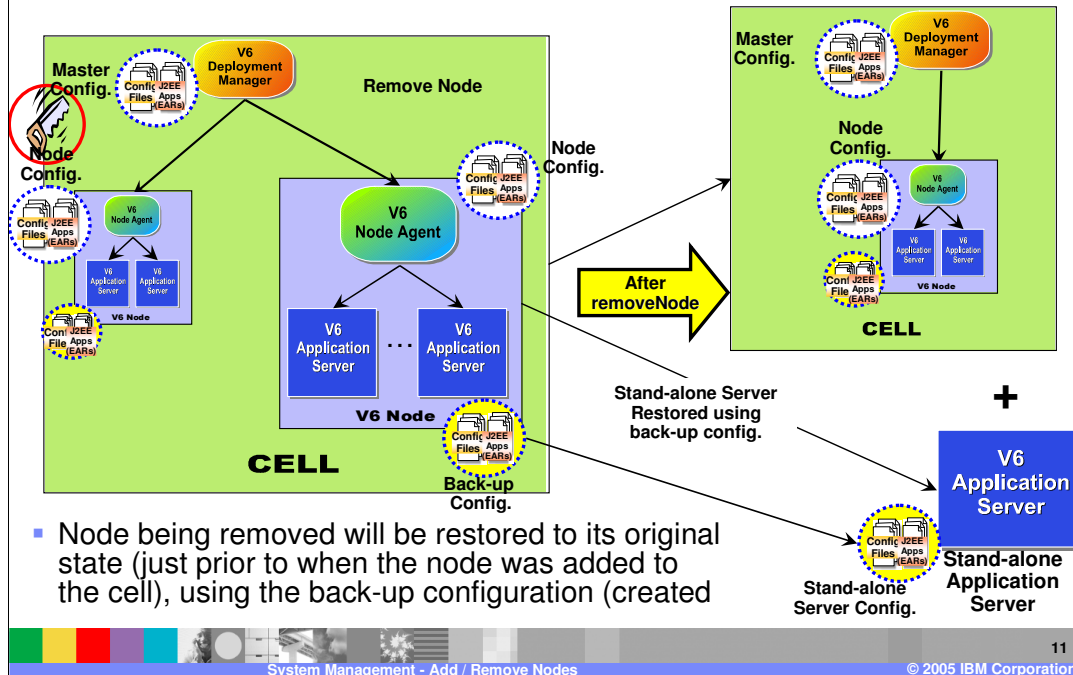
When adding a secured node to an unprotected cell, the node will now be unprotected. When adding an unsecured node to a secured cell, the node will now be secured. When adding a secured node to a secured cell, the node will continue to be secured, but with the security settings for the cell.

Section

Removing V6 Nodes from a Cell

This section covers removing nodes from the cell.

Remove V6 Managed Node from a Cell



This page shows the process of remove node. On the node that needs to be removed, a “removeNode” command is executed, as shown by the saw.

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 adding of the node, the add node utility had created a backup of the configuration. The remove node utility uses the backup configuration it to restore the configuration to its original state.

When doing an uninstall, this utility is called to remove the node from the cell. 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.bat (sh) command from the Node
 - ▶ Will need the administrator user ID and password, if security is ON
 - ▶ Node uninstall also makes use of the same command
- 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 2 ways to remove a node. One uses the remove node command and it is executed on the node being removed. The other way is through the Deployment Manager administration console in the nodes section.

If security is on, the remove command would need an administrator user id and password to be supplied.

When doing an uninstall of a federated node, the remove node utility is called to remove the node from the cell.

Since you go back 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 its 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 System out and error log files of DMgr and Node, along with the addNode or removeNode log files



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

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

Then 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 has 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.