



IBM Software Group

IBM® WebSphere® Application Server V6.01

Installation – z/OS® Platform



@business on demand.

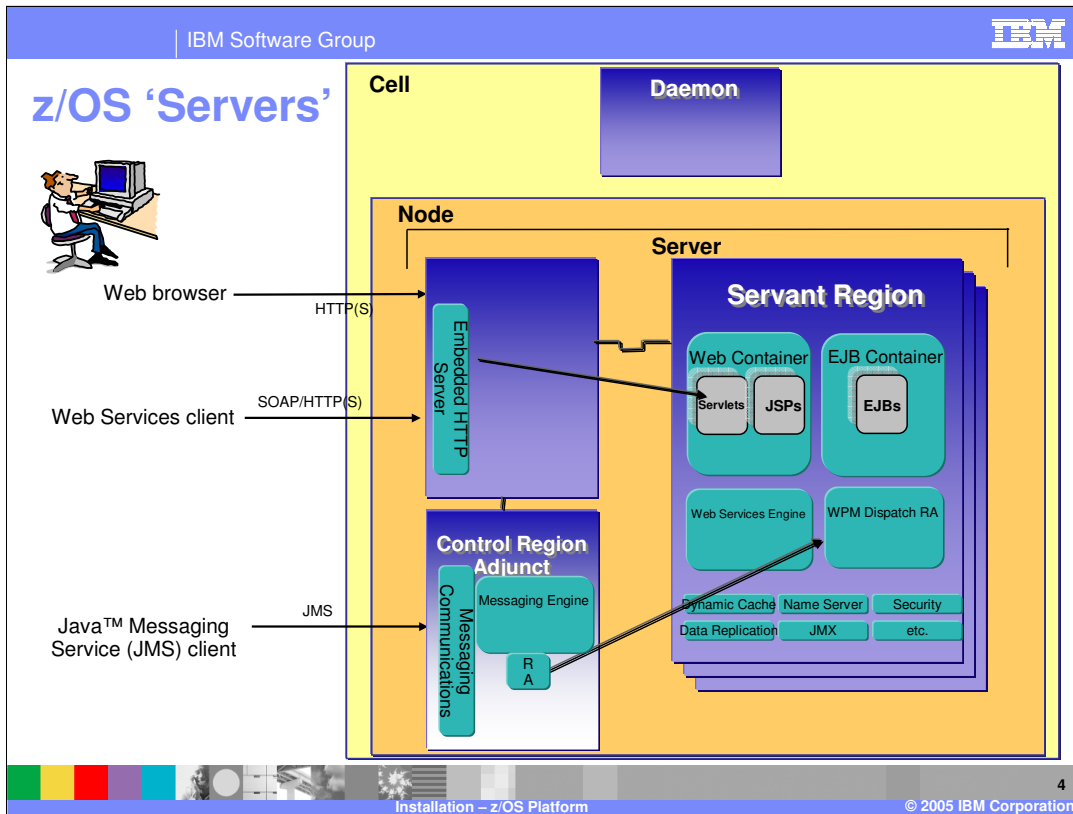
© 2005 IBM Corporation
Updated April 25, 2005

Agenda

- V6.01 System Structure on z/OS
- System Prerequisites
- Overview of Install Process
- Security Domain Configuration
- Stand-alone Application Server Configuration
 - ▶ Application Server
- Network Deployment Configuration
 - ▶ Deployment Manager
 - ▶ Empty Managed Nodes
 - ▶ Add Node (Federation)

Section

System Structure



WebSphere for z/OS V6.01 involves a few address spaces, including a possible new one that is used for communication to the new messaging engine. It works in conjunction with the current Control Region, working as the communications endpoint for messaging. The Control Region that is familiar from the V5 days is still there plugging away, taking care of the HTTP communications. The new Control Region Adjunct, or CRA, helps out with the messaging communications. Unlike the Control Region which is ALWAYS there, the Control Region Adjunct is only present when the a server becomes a bus member and has a messaging engine created. The messaging engine runs in the CRA, communicating with the Resource Adapter (RA) shown. The RA is really split across the two processes where the WPM Dispatch RA is where the MDBs eventually get dispatched.

Section

System Prerequisites

Pre-reqs

- Hardware requirements
 - ▶ Any hardware that supports z/OS 1.4 or higher or z/OSe 1.4 or higher
- Software requirements
 - ▶ z/OS 1.4 or higher or z/OSe 1.4 or higher with the following features installed, enabled and configured
 - z/OS Communications Server (TCP/IP) or equivalent
 - z/OS UNIX System Services and the hierarchical file system (HFS)
 - SecureWay Security Server (RACF) or equivalent security management product
 - System logger
 - System SSL security required when using SSL
 - Workload management in goal mode
 - Resource recovery services
 - ▶ SMP/E 3.2.0



Many of the requirements for V6 are the same as they were on V5. One noticeable change though is the requirement of z/OS 1.4 or higher. If you plan to exploit any zAAPs in your configuration, you'll need a z990 GA3 or z890 and will need to run z/OS 1.6.

You should also note the SMP/E level where you'll need to make sure you have SMP/E APAR **IR54653** installed. It is needed for the APPLY processing to complete. Also note that java now needs to be installed and configured on your driving system. This is a new requirement for V6.

Implementation Options (Optional)

- IBM HTTP Server for z/OS or equivalent
- LDAP
- DB2® Universal Database Server for z/OS, Version 7 or higher
- CICS® TS 1.3.
- WebSphere MQSeries® for z/OS V5.3.1 and higher
- WebSphere MQ Integrator Broker for z/OS V2.1 and higher
- Rational® Application Developer 6.0



This foil lists some possible optional products that you may use. Note that the tools for EJB, servlet and JSP development are now provided in the Rational Application Developer V6.0, the follow-on to WebSphere Application Developer.

Section

Overview of Install Process



This section looks at the install process.

Installation on z/OS


- SMP/E
 - ▶ Product Code now includes SDK
 - Unpaxed automatically during the APPLY process
 - ▶ adminconsole.ear, filetransfer.ear stays with product binaries even after configured
 - /usr/lpp/zWebSphere/V6R0/systemApps
- ISPF Customization Dialog
 - ▶ Same appearance as V5.x
 - ▶ A few new fields/panels to deal with V6 changes
 - ▶ Should be able to reuse some of V5 values
 - ▶ ex 'BBO.SBBOEXEC(bbowstrt)'

New in V6, IBM will now ship the SDK as part of the product code. It will be unpaxed automatically during the APPLY process into your product directory: /usr/lpp/zWebSphere/V6R0/java. The code is certified at particular java levels and so should not be changed unless instructed to do so.

Something else that is new in V6 is the idea of System Applications. These include applications such as adminconsole and filetransfer. Prior to V6, these apps were actually installed along with your other applications into your server and they became part of your configuration. With V6, they stay with the product binaries and a softlink points to them from a new systemApps directory in your configuration HFS. Some apps used for the PMEs will also reside in the systemApps directory including ManagementEJB.ear and SchedulerCalendars.ear.

For the most part, your configuration experience should be similar to V5.x if you are coming from there. V6 still uses an ISPF Customization Dialog to help you configure your WebSphere environment. You will notice a few new fields and you will notice some are missing but it has remained relatively unchanged. The command to run the ISPF Customization Dialogs is shown in the last bullet.

Version 6 Packaging



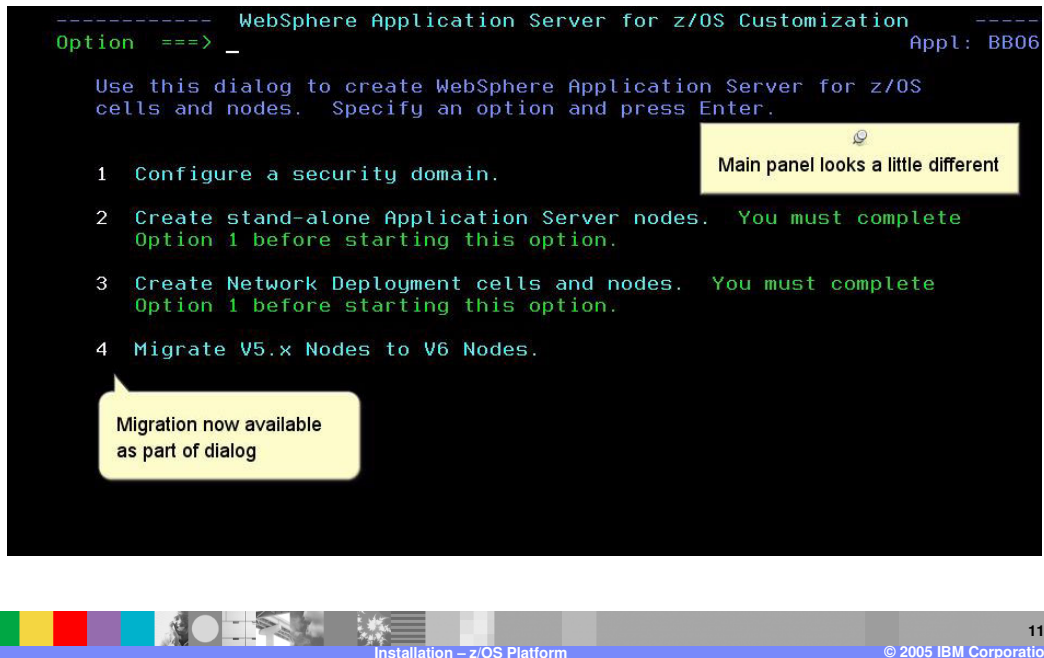
| | |
|-----------------|---|
| Component Name: | WebSphere Application Server for z/OS V6.0.1 |
| Program Number: | 5655-N01 |
| FMID: | H28W601 |
| COMPID: | 5655I3500 (same as WebSphere Application Server V.5) |
| PSP Bucket: | Upgrade=WASAS601, Subset=H28W601 |
| Documentation: | http://www.ibm.com/software/webservers/appserv/was/library/ |

This foil looks at the packaging of V6. One of the big changes in V6 is the addition of the PMEs to the base WebSphere product. Prior to V6, the PMEs were only available on the WebSphere Business Integration Server Foundation (WBISF) product.

Also new on V6 is a second very large HFS that contains CD images for the z/OS user to download. Included here are 26 CDs containing software not intended to run on z/OS. The CDs include Edge Components, the WebSphere Application Server Application Server Toolkit (ASTK), a trial

version of Rational Application Developer and Tivoli® Access Manager. See the appropriate articles in the WebSphere Application Server Information Center for installation and use of these CDs.

New Main Panel



The main panel of the ISPF Customization Dialog looks a little different. You'll notice the option to Migrate appears in the ISPF Dialog now. You can learn about that in the Migration Presentation. The first three options will be presented here.

You'll notice that there is now a 'stand-alone' application server instead of a 'base' application server as in V5. In V5, once you had your application server configured, you could go into the Administrative Console and add additional servers. This is no longer possible in V6, thus the new designation of 'stand-alone'.

Start by looking at option 1 and configure the security domain. This will be used for all the other options.

Preparation

- Before starting your customization
 - ▶ Plan, Plan, Plan!
- Many values are needed
 - ▶ Come up with a naming convention
- For an example, look here:

▶ <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100367>

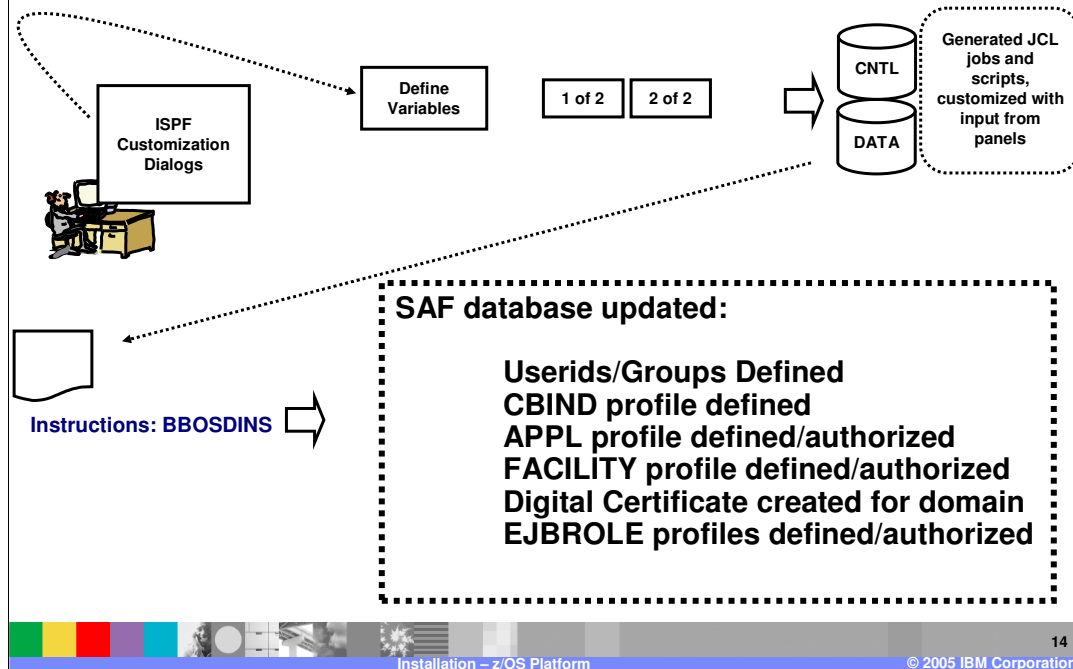


Lots of information is asked for in the ISPF customization dialogs to customize your environment. It's a good idea to sit down and plan out your configuration. In order to keep your sanity, you'll want to come up with a naming convention. The foil shows a techdoc where you can find a sample V5 configuration. There is also a spreadsheet available that helps you implement the suggested naming convention. Click on the Show Me for more information on that.

Section

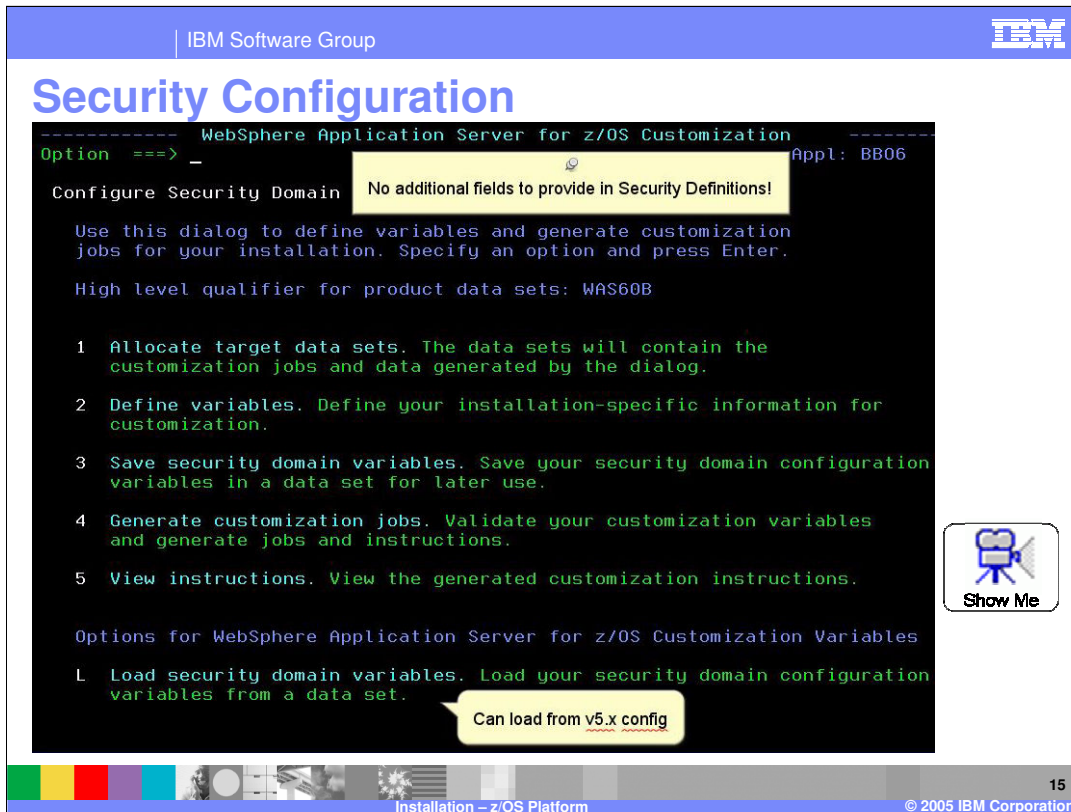
Security Domain Configuration

Creating the Security Domain



The first thing you want to do when configuring any type of WebSphere Application Server node is define the 'security domain' that you will use. The information you supply here will be loaded into the other configuration options as you will see a little later. There will be other security definitions that will be created as you configure the other options but the basic security domain definitions are set up in this first option. For instance, if you decide to use SAF authorization, it will create EJBROLE definitions for you and authorize the userid you select to those definitions so that you have an administrator available when you decide to turn on Global security.

Next look at how this is done.



This is the first screen you see when configuring the security domain. You can reuse the values from your V5.x security configuration if you wish. You'll see that there are no additional fields to provide in the Security Definitions.

This screen is the similar to how all the various configurations are done. If you have variables from an existing configuration, you can load those in as a base (option L). You will then want to allocate datasets (option 1) to hold the jobs that you will need to run (CNTL) and the various other execs/data that is needed to define your environment (DATA). Once that is done, you will actually define the variables (Option 2) necessary to configure the environment you are working on (the Security Domain in this case). Once everything is defined as you wish, you should save your variables using option 3 and generate the customization jobs using option 4. Once that is complete, you can view the instructions using option 5. This defines any manual steps that may be necessary as well as spelling out the jobs that you need to run (along with the userid it should be run under). Once the jobs have been run, the environment you have been configuring should be complete.

To access the a Demonstration of the security configuration, click on the Show Me icon.

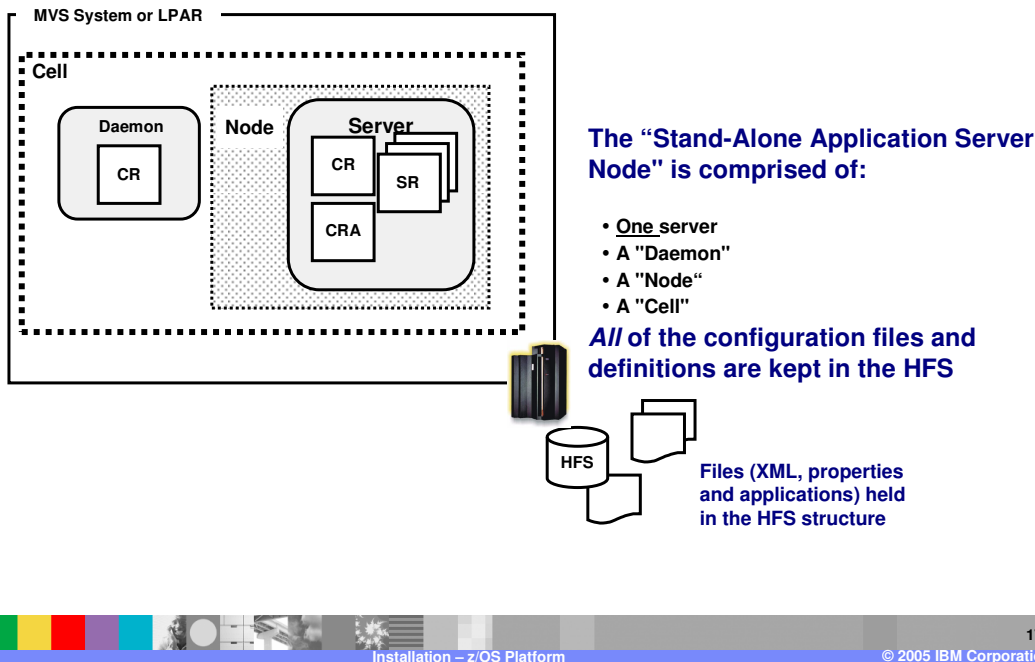
Section

Stand-Alone Application Server Configuration



The next configuration is the 'Stand-Alone Application Server Node'.

V6.0.1 'Stand-Alone' Application Server Node

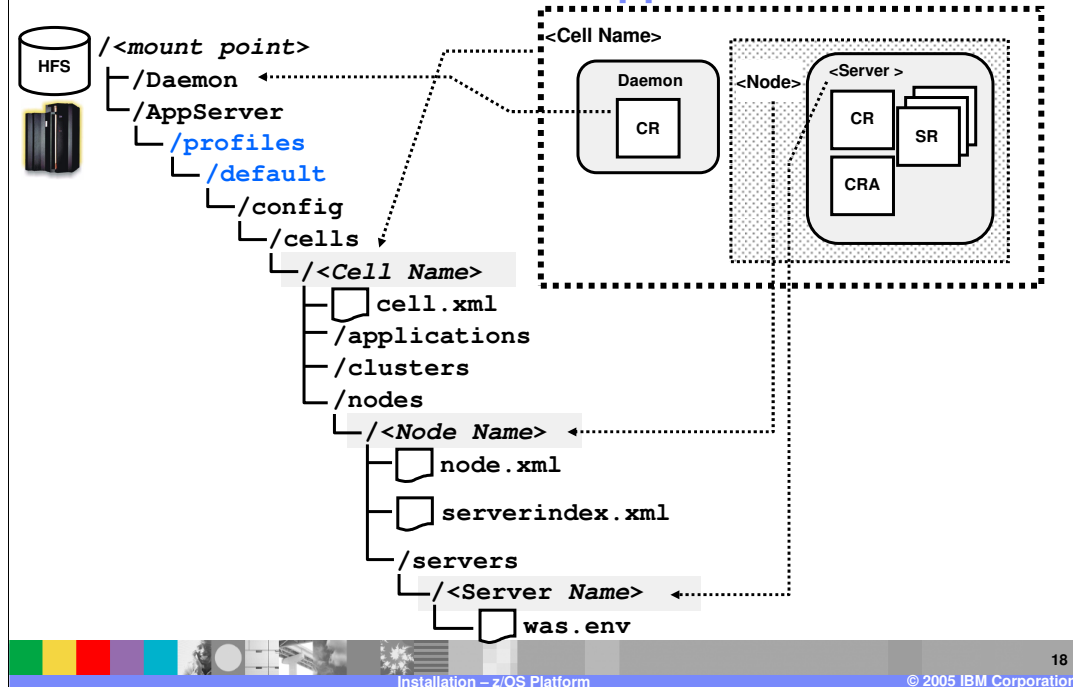


This foil has a picture of what will be configured once you have gone through the ISPF Configuration Dialog for this environment. You will end up with ONE **Server** which consists of a Control Region (CR), one or more Servant Regions (SR) and a Control Region Adjunct (CRA). WLM will start the Servant Regions as necessary to meet its goals and they will be given work from the Control Region working with WLM. The CRA structure is set up for you through the ISPF Dialogs but will not be started unless you add this server as a member of a service integration bus that you have created through the Administrative Console. Remember that in V6, the Administrative Console has no option to add additional servers...thus the 'stand-alone' designation.

The **Daemon** is a single CR server and one is required per cell per system or LPAR. It provides the location service in support of Remote Method Invocation and Internet Inter-ORB Protocol (RMI/IIOP).

The concept of "**Cells and Nodes**" does not reveal itself when looking at a Stand-Alone Application Server Node. It will make more sense later. For now, understand that a Stand-Alone Application Server Node has one node and one cell. The cell/node structure can be found in the HFS underlying the Stand-Alone Application Server which you will see on the next foil.

HFS Under the Stand-alone Application Server Node



All of the configuration information for a Stand-Alone Application Server Node is kept in the HFS.

The `<mount point>` is a directory that serves as the mount point for the HFS file system that will hold this stand-alone application server's information and applications. The name for `<mount point>` may be anything you wish.

Under the mount point you'll find two directories -- one for the Daemon server and one for the Application Server. We've expanded the AppServer directory to show you the details, but be aware that the Daemon directory has a similar structure as well.

Digging into the AppServer directory, if you are coming from V5, you'll notice a new **/profiles/default** directory highlighted in blue. This comes from the distributed idea of creating 'profiles'. Just remember that on z/OS there is just this one 'profile'. Under the `/profiles/default/config/cells` directory then, you'll see a directory that will have the same name as you provide for the "Cell Long Name" during the customization for this stand-alone application server

Further down you find a directory under the `/nodes` directory that bears the same name as you provide for the "Node Long Name."

Finally, down under the `/servers` directory you'll find a directory that has the same name as the name you provide for the "Server Long Name." It is in this directory that a very special file resides: the `was.env` file. The `was.env` file is the primary configuration file for the application server.

Basic Configuration Dialog

```

----- WebSphere Application Server for z/OS Customization -----
Option ==> 3_
Create a stand-alone application server node
Use this dialog to define variables and generate customization jobs
for your installation. Specify an option and press Enter.
High level qualifier for product data sets: WAS60B

1 Load security domain variables. Load your security domain variables.
2 Allocate target data sets. The data sets will contain the
  customization jobs and data generated by the dialog.
3 Define variables. Define your installation-specific information for
  customization.
4 Generate customization jobs. Validate your customization variables
  and generate jobs and instructions.
5 View instructions. View the generated customization instructions.
Options for WebSphere Application Server for z/OS Customization Variables
S Save customization variables. Save your customization variables
  in a data set for later use.
L Load customization variables. Load your customization variables
  from a data set.

```

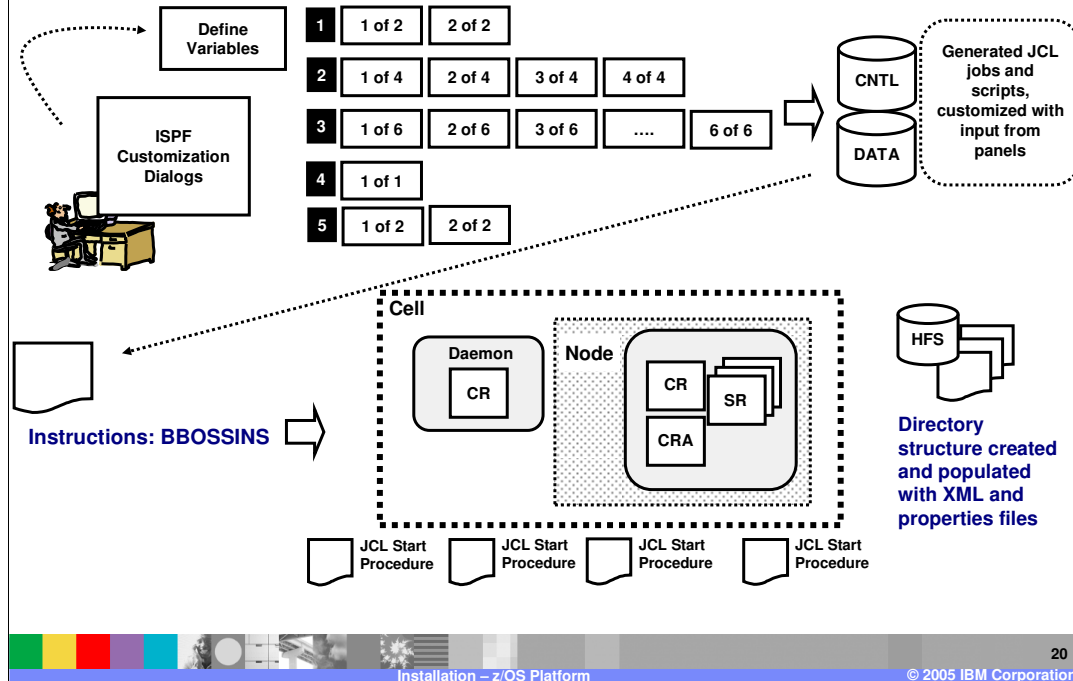
Each of the possible configurations start with the ISPF Customization Dialog. The screen shown here is what you see when you select that you want to configure a Stand-Alone Application Server node. Before you continue on with the actual Stand-Alone Application Server Node configuration, take a minute to look at this. All of the basic configurations start with a screen just like this so whether you are configuring a Stand-Alone node, a Network Deployment node, an empty managed node or federating a node you will start with this basic screen. If you are coming from V5, nothing has changed here. You may have a set of customization values that you want to start with. You can load them with option **L** if you have previously saved them with option **S**. For instance, if you are configuring a Network Deployment cell, you will want to Load your variables from your Stand-Alone Application node configuration first. If you are starting from scratch, you should first do some planning to come up with naming conventions. There are worksheets that are found in the Information Center or there is a spreadsheet that has been created that can help in your planning.

If you have loaded a previous configuration, as long as it has the security domain variables that you plan to use, you can skip Option **1** where you load the security domain variables. These should be included with your loaded customization variables. If you want or need to use a different set of values, you can load them in using option **1** to overwrite the ones currently defined. If you are starting from scratch, you will want to load the security domain variables that you previously defined as one of the first steps.

The second step you will take is to allocate the target data sets that will be used for your customized jobs and scripts. These are the CNTL and DATA PDSes you will see in the next foil.

The third step is where most of the configuration happens. This is where you will provide

Creating Stand-alone Application Server Node



Starting out at the ISPF Customization Dialogs, you'll go through pretty much the same steps as you did when you defined the security domain. You will define the variables needed for this configuration such as a cell name, node name and a server name. You'll describe where to find the SMP/E installed code. You'll name the JCL procs that will be used to run WebSphere and give you all the other information that you need to customize your environment. Once you are satisfied, you'll generate the jobs and scripts that you will need to run which will end up in a CNTL/DATA PDSes that you named. The CNTL dataset will also include a set of instructions for you to follow, giving you any manual steps that may be necessary and then instructing you to run the customized jobs that are created for you. Once you have followed the instructions, you will end up with the structure shown with a populated HFS that represents your WebSphere Application Server environment as well as the JCL Start Procedures needed to start the new Stand-Alone Application Server. JCL Procs are created for the Daemon, the Control Region (CR), the Servant Region(s) (SR) and the Control Region Adjunct (CRA).



Stand-alone Application Server Configuration

```
----- WebSphere Application Server for z/OS Customization -----
Option ==> 2_                               Appl: BB06

Use this dialog to create WebSphere Application Server for z/OS
cells and nodes. Specify an option and press Enter.

1  Configure a security domain.
2  Create stand-alone Application Server nodes. You must complete
   Option 1 before starting this option.
   Create Network Deployment cells and nodes. You must complete
   Option 1 before starting this option.
   Migrate V5.x Nodes to V6 Nodes.
```

'base' Application Server node has become 'stand-alone' Application Server node



Installation - z/OS Platform 21 © 2005 IBM Corporation

To see a demonstration of the Stand-Alone Application Server Node configuration, click on the Show Me icon.

Section

Network Deployment Configuration: Deployment Manager



This section will look at the Network Deployment configurations.

Network Deployment options

```
----- WebSphere Application Server for z/OS Customization -----
Option ==> 3
Appl: BB06

Use this dialog to create WebSphere Application Server for z/OS
cells and nodes. Specify an option and press Enter.

1 Configure a security domain.
2 Create stand-alone Application Server nodes. You must complete
Option 1 before starting this option.
3 Create Network Deployment cells and nodes. You must complete
Option 1 before starting this option.
Migrate V5.x Nodes to V6 Nodes.

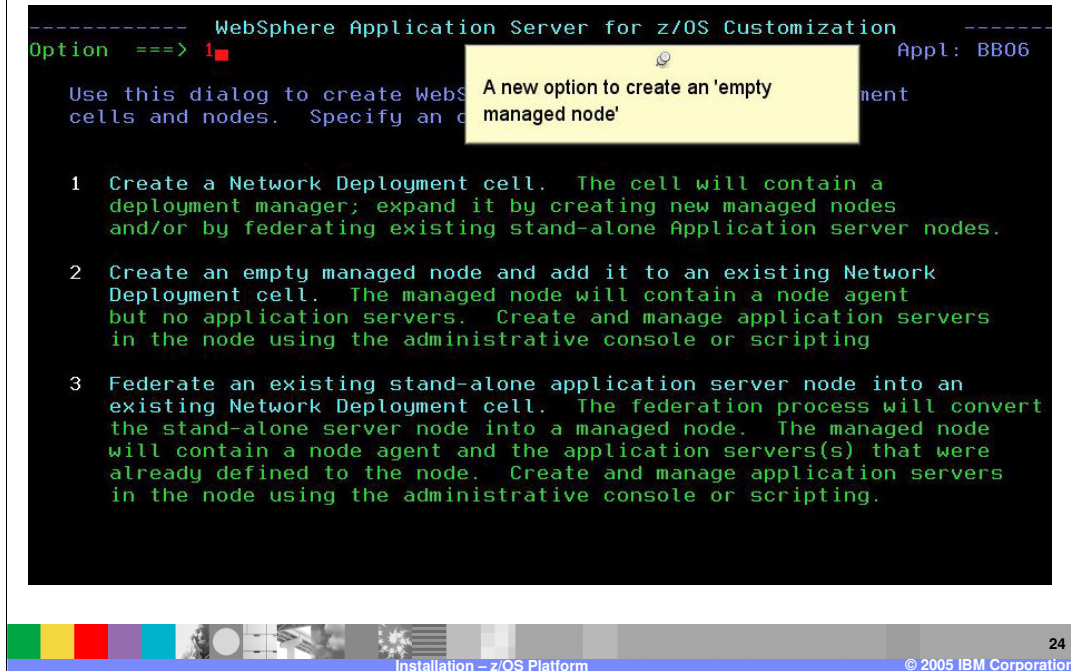
Note that Network Deployment 'options'
have been lumped together
```

Installation - z/OS Platform

© 2005 IBM Corporation 23

In V5, all options were on the main panel. The first 2 options were the same. The third option in V5, 'Configure integral JMS provider' is no longer necessary since embedded messaging is fully integrated in WebSphere in V6 and no additional configuration is needed. The last two options you had in V5, 'Configure deployment manager node' and 'Federate base Application Server node' are now lumped together under option 3. You can learn about option 4 in a separate presentation. The next few slides will explore option 3.

Create Network Deployment Cells and Nodes



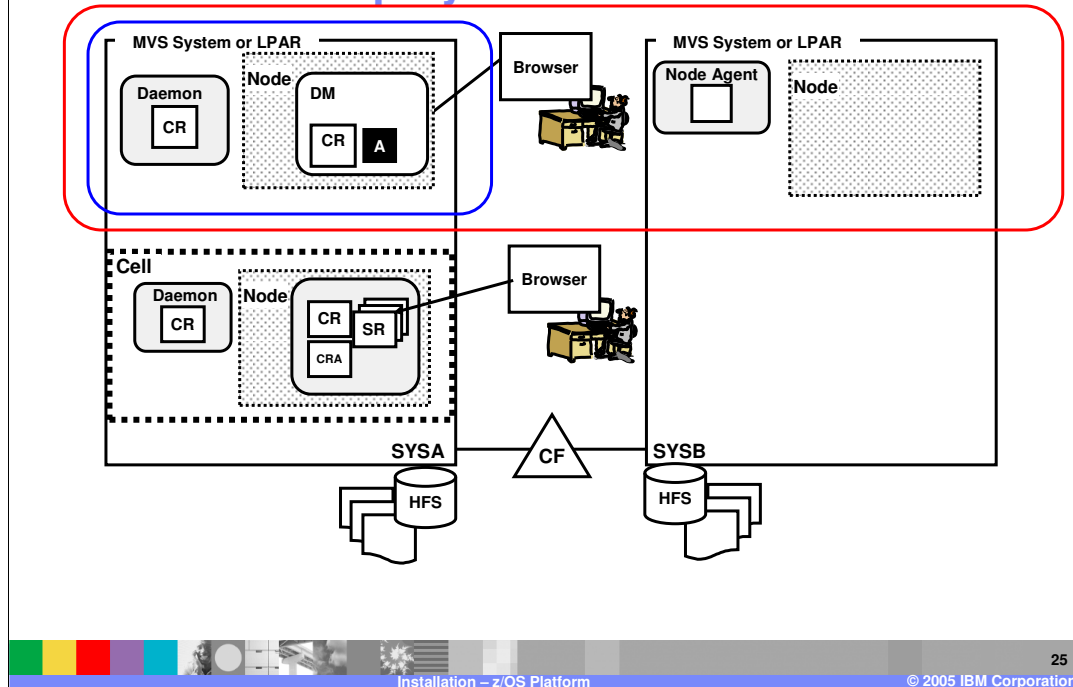
The first option on this screen allows you to create a Network Deployment Cell. This is an administrative node that allows you to add managed nodes to it later to be managed from this central point. To learn more about the structure of the Network Deployment Cell, you can review some of the architecture presentations.

Option 2 is a new option that has been introduced in V6. WebSphere now allows you to create an empty managed node so that you can create the structure of a node but are not forced to create a server in it. You can add application servers to the node via the Administrative Console since it will be federated into an existing Network Deployment cell. This is useful for server clustering where a node is being federated to run a clustered instance of a server (and the server has already been added to the cluster from another node). Often, the 'Stand-Alone application server' wasn't needed on the second node. This is also useful when you need to define a node where you want to manage a web server from and again do not need a Stand-Alone server defined.

Finally, the third option allows you to federate a previously configured Stand-Alone Application Server node into an existing Network Deployment cell (as defined in option 1).

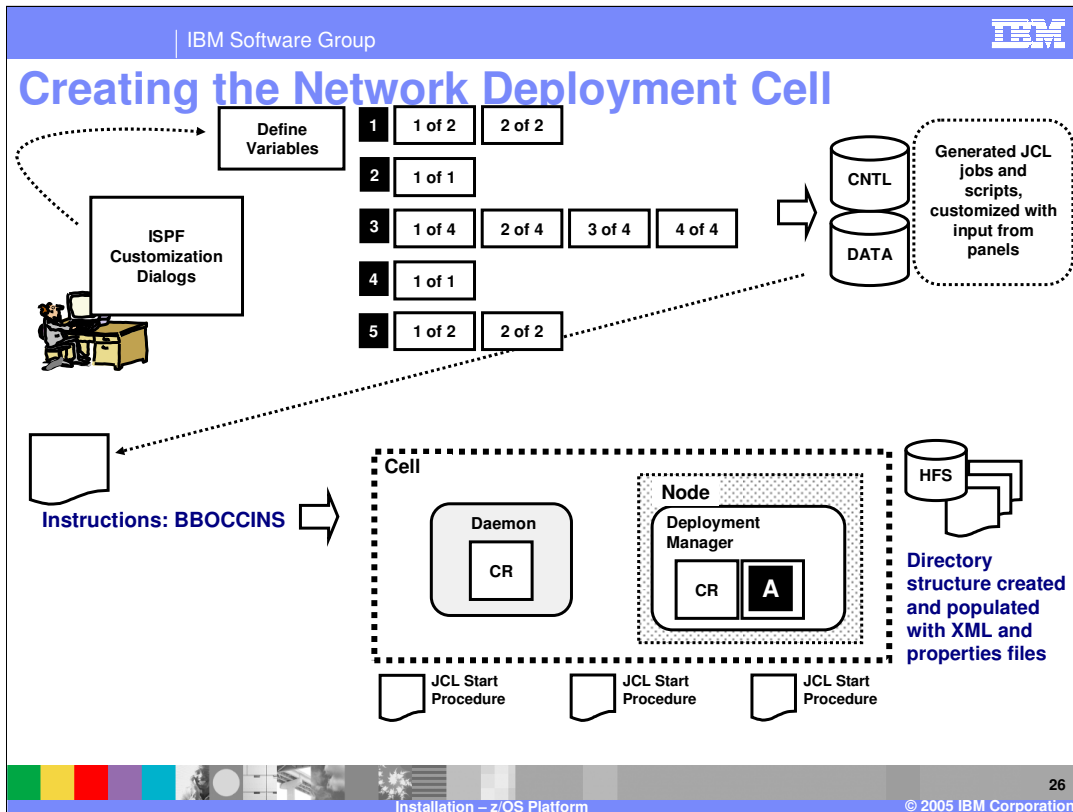
On the next slide, you will see a pictorial representation of the different options.

The Network Deployment Cell

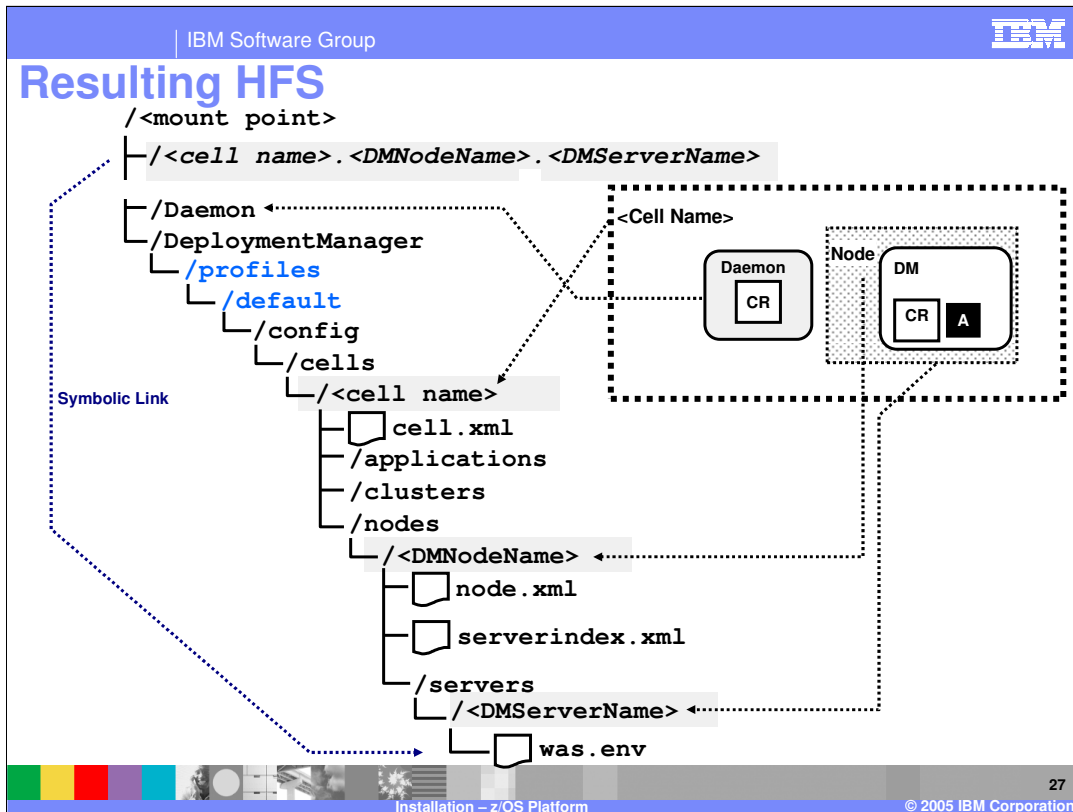


In the Network Deployment Cell Configuration, you will first create a Network Deployment Cell that will consist of a Daemon and Deployment Manager node and server where the Administrative Console will run. It is shown in the blue rectangular box. This picture also shows a 'stand-alone application server' on SYSA in the bottom of the box as well as an 'empty managed node' on SYSB. You'll notice that there are no Servers defined in the node on SYSB but it has been federated into the Deployment Manager cell, depicted by the red rectangular box. When you run option 3 to 'Federate an existing stand-alone application server', the stand-alone server in SYSA will become part of the Deployment Manager's cell.

We'll look at the option first where you will create a Network Deployment Cell (the blue box).



The configuration of a Deployment Manager cell is almost identical to the configuration for the 'Stand-Alone application server node', with a few less variables to supply. You should load the SAVECFG from your 'stand-alone application server' configuration as a starting point for your network deployment cell. You'll note that there are only three JCL procs created in this configuration...one for the Daemon, one for the Control Region (CR), and one for the Servant Region(s) (SR). No Control Region Adjunct (CRA) is needed in this environment.



The HFS for the Network Deployment cell configuration is similar to the Stand-Alone Application Server Node HFS structure. If you're coming from V5, you'll again notice the new `/profiles/default` directory highlighted in blue. This picture points out an important piece of the HFS structure that is used in the JCL to point to the important `was.env` file that was mentioned earlier. There is a symbolic link that is used to point to the `was.env` which is needed to start the various address spaces. The symbolic link name is built with the `<short cell name>.<short node name>.<short server name>`. This is the same for the Stand-Alone application server.

The `<mount point>` here again is selected by you. It can be shared with a stand-alone application server node or you can create a new one for the deployment manager cell. As you add nodes to this network deployment cell, the HFS here will become the place where the master configuration is kept.

Again, under the mount point you'll find two directories -- one for the Daemon server and one for the Deployment Manager. We've expanded the DeploymentManger directory to show you the details.

Create Network Deployment Cell

```
----- WebSphere Application Server for z/OS Customization -----
Option ==> !                               Appl: BB06

Use this dialog to create WebSphere for z/OS Network Deployment
cells and nodes. Specify an option and press Enter.

1 Create a Network Deployment cell. The cell will contain a
deployment manager; expand it by creating new managed nodes
and/or by federating existing stand-alone Application server nodes.

Create an empty managed node and add it to an existing Network
Deployment cell. The managed node will contain a node agent
but no application servers. Create and manage application servers
in the node using the administrative console or scripting

Federate an existing stand-alone application server node into an
existing Network Deployment cell. The federation process will convert
the stand-alone server node into a managed node. The managed node
will contain a node agent and the application servers(s) that were
already defined to the node. Create and manage application servers
in the node using the administrative console or scripting.
```



To see a demonstration of the Network Deployment cell configuration, click on the Show Me icon.

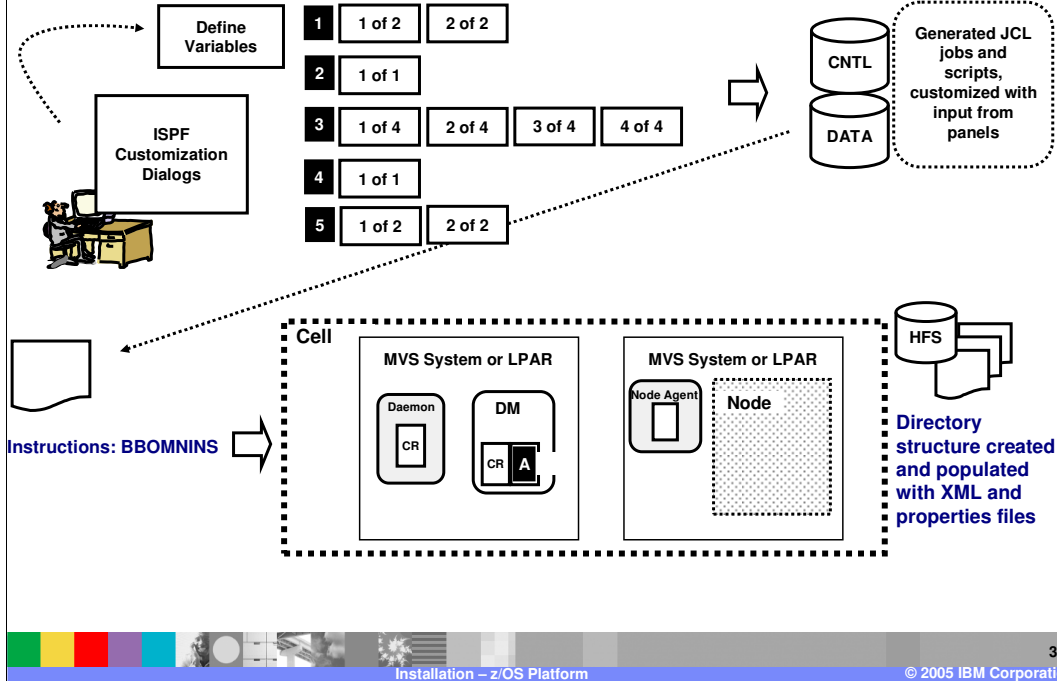
Section

Network Deployment Configuration: Empty Managed Nodes



You have looked at creating a Network Deployment cell. The next section will look at creating an empty managed node that will be federated to it.

Creating the Empty Managed Node and federate



For the Empty Managed Node, we'll again start with the ISPF Customization Dialog to define the variables that are needed to configure the empty managed node and federate it. Remember, the empty managed node is a node with no servers and the idea behind it is that it will be federated right away as shown on this slide. The empty managed node is nothing on its own. It is federated immediately where you can then create a cluster member in the empty node or a web server definition. The HFS will look just like a stand-alone server without a server directory.

Network Deployment: Empty Managed Node


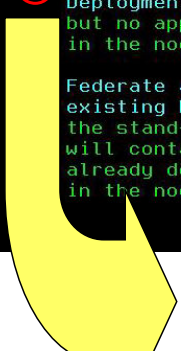
```
----- WebSphere Application Server for z/OS Customization -----
Option ==> 2_                               Appl: BB06

Use this dialog to create a Network Deployment cell. Specify the cell name,
cells and nodes. Specify the cell type as Network Deployment.
```

Let's look at 'empty managed node' option

- 1 Create a Network Deployment cell. The cell will contain a deployment manager; expand it by creating new managed nodes and/or by federating existing stand-alone Application server nodes.
- 2 Create an empty managed node and add it to an existing Network Deployment cell. The managed node will contain a node agent but no application servers. Create and manage application servers in the node using the administrative console or scripting

Federate an existing stand-alone application server node into an existing Network Deployment cell. The federation process will convert the stand-alone server node into a managed node. The managed node will contain a node agent and the application servers(s) that were already defined to the node. Create and manage application servers in the node using the administrative console or scripting.



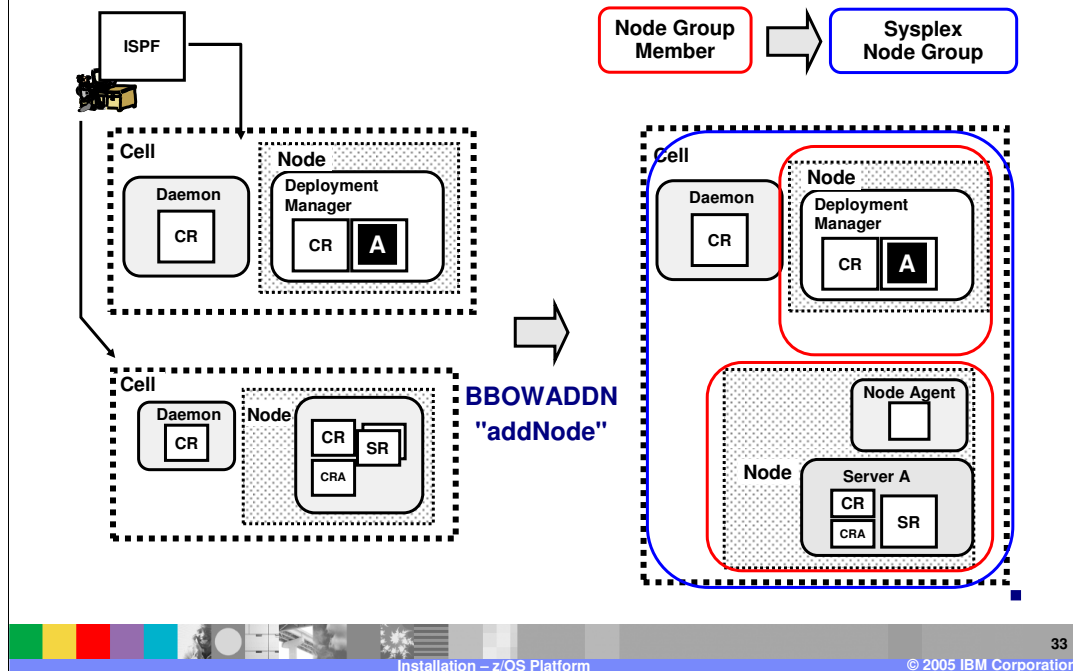
Show Me

To see a demonstration of the Empty Managed Node configuration, click on the Show Me icon.

Section

Network Deployment Configuration: Add Node (Federation)

“Federating” a Stand-alone Application Server Node



Installation – z/OS Platform

© 2005 IBM Corporation

33

The last configuration we'll look at is the Federation of an existing Stand-Alone Application Server node. This assumes that you have already configured both a Stand-Alone application server node and a Network Deployment cell. The purpose of the federation is to make the Stand-Alone Application Server node part of the Network Deployment cell. In the process, the Stand-Alone Application Server will take on both the Network Deployment cell's Daemon and its cell name. In the process of federation, a node agent will be created so that the Deployment Manager node can update the new node with any changes to its configurations. The master configuration for the federated node will be kept at the Deployment Manager after federation. The process of federation is accomplished with the 'addNode' command. Through the ISPF Dialogs, there will be a job created where the invocation is customized for your environment. The job is called BBOWADDN.

In V6, there are a couple of new parameters that you will see on the command:

- includebuses
- nodegroupname

These are parameters that will be specified in the ISPF Customization Dialogs. Note that the two nodes on the left are now part of the same cell on the right and ALSO become part of a sysplex node group. Each red rectangle signifies a node group member. Node groups are a new concept in V6. For more information about node groups see the node group presentation.

Federate an existing stand-alone node

```
----- WebSphere Application Server for z/OS Customization -----
Option ==> 3_                               Appl: BB06
Use this dialog for z/OS Network Deployment
cells and nodes                               Press Enter.

1 Create a Network Deployment cell. The cell will contain a
  deployment manager; expand it by creating new managed nodes
  and/or by federating existing stand-alone Application server nodes.

2 Create an empty managed node and add it to an existing Network
  Deployment cell. The managed node will contain a node agent
  but no application servers. Create and manage application servers
  in the node using the administrative console or scripting

3 Federate an existing stand-alone application server node into an
  existing Network Deployment cell. The federation process will convert
  the stand-alone server node into a managed node. The managed node
  will contain a node agent and the application servers(s) that were
  already defined to the node. Create and manage application servers
  in the node using the administrative console or scripting.
```

Business as usual, federate the 'stand-alone' node we created



To see a demonstration of the Federation configuration, click on the Show Me icon.

Section

Summary

Summary

- Nothing radically different in the WebSphere for z/OS V6.01 installation
- Same familiar ISPF Customization Dialog
- A few new concepts introduced:
 - ▶ Node Groups
 - ▶ SystemIntegration buses
 - ▶ Empty Managed Nodes



In summary, there is nothing radically different in WebSphere for z/OS V6.0.1 installation if you are coming from V5. You use the same familiar ISPF Customization Dialogs as have been used since v4. There are a few new concepts that are introduced including node groups, system integration buses and empty managed nodes. More information on the first two concepts can be found in separate presentations.

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 | iSeries | 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 multiprogramming 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.