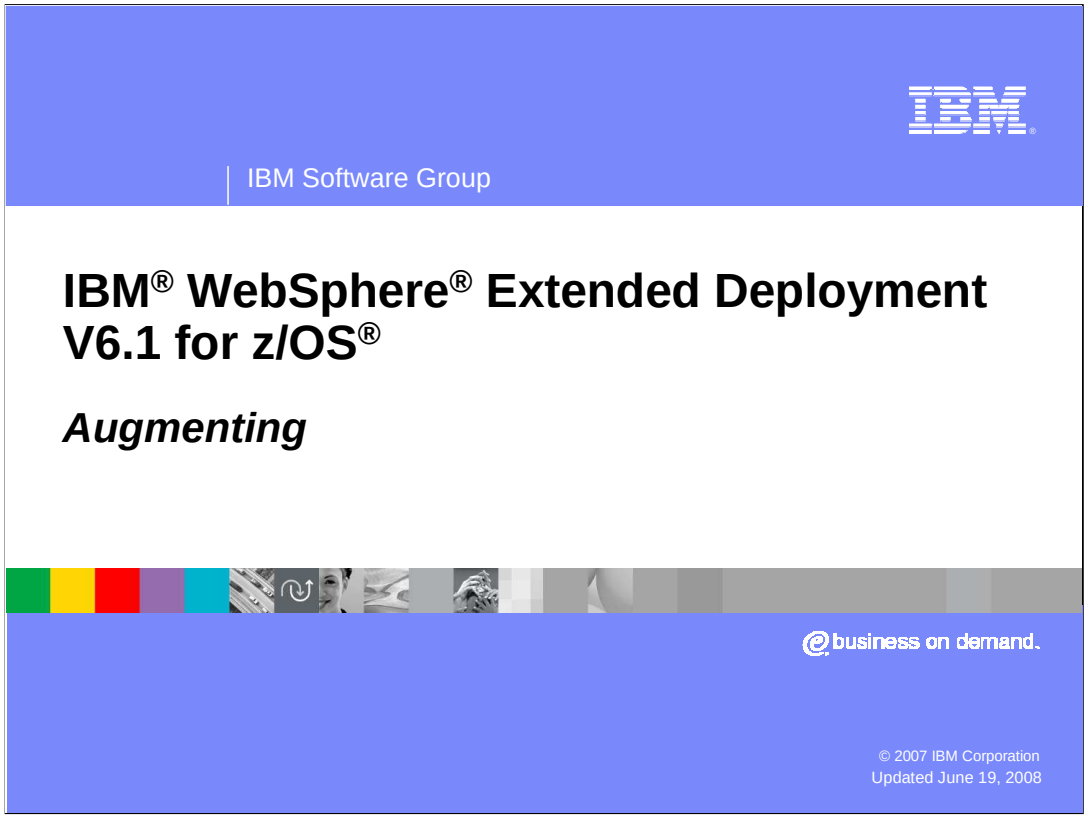

This presentation will discuss augmenting a WebSphere Network Deployment cell to a WebSphere Extended Deployment cell.

This module references

- WebSphere Extended Deployment Data Grid, which is now called WebSphere eXtreme Scale; and
- WebSphere Extended Deployment Operations Optimization, which is now called WebSphere Virtual Enterprise.

Though the module uses the previous names, the technical material covered is still accurate.

Agenda

- Augmentation overview
- Augmenting
- Troubleshooting
- Summary

The first topic is an overview of augmenting an Extended Deployment cell followed by augmenting details and then a slide on where to find information in case there is a problem.

Required new tool

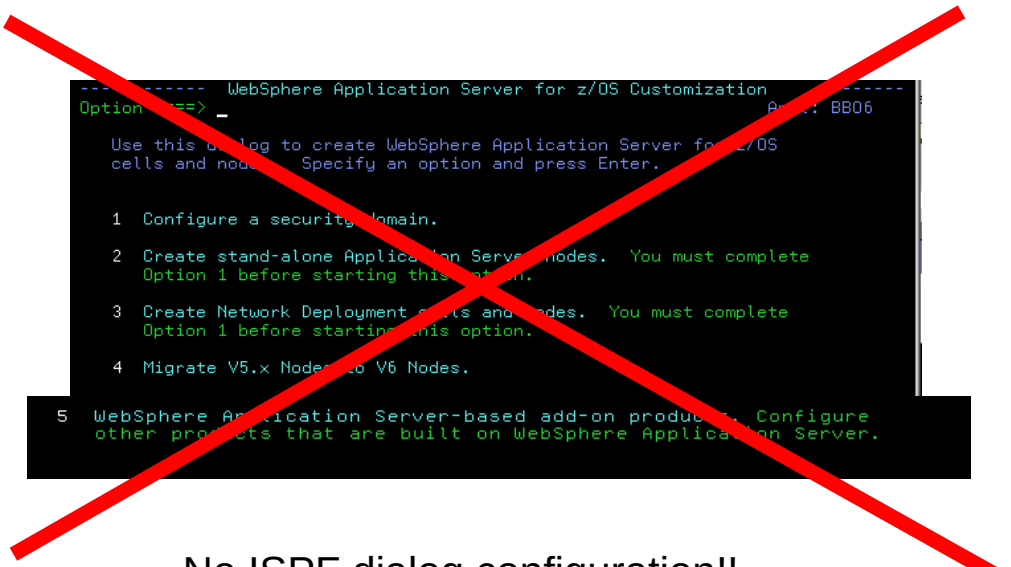

No ISPF dialog configuration!!

As you learned on a previous presentation, the WebSphere Extended Deployment for z/OS V6.1 libraries are fully contained in an HFS. You will notice that ISPF Dialogs are not used for configuration of the products because it would require the use of PDS. By using only shell scripts for configuration, the installation among the various platforms is more consistent and takes advantage of common code.

Creating profiles

- Extended deployment provides a plug-in to the profile creation and augmentation tool
  - Creates new WebSphere Extended Deployment augmentation profiles
  - Can create un-augmented profiles
    - These profiles have no WebSphere Extended Deployment capabilities unless they are augmented
  - Augments WebSphere Application Server profiles
  - Profiles not augmented during installation must be augmented to have Extended Deployment capabilities

The zPMT tool is delivered on CD as part of WebSphere Application Server V6.1 for z/OS. The first step is to install zPMT on your workstation. Install the Extended Deployment plug-in on the zPMT CD delivered as part of WebSphere Application Server V6.1, it should have a name similar to com.ibm.ws390.pmt.xd_6.1.1.v…..jar.

Building zPMT

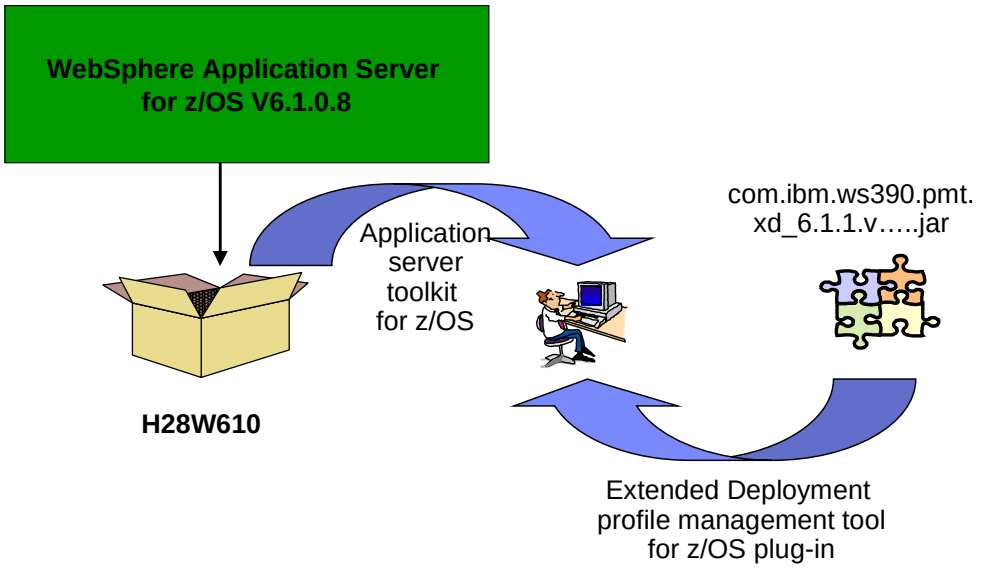

Here is a picture of the process just described.

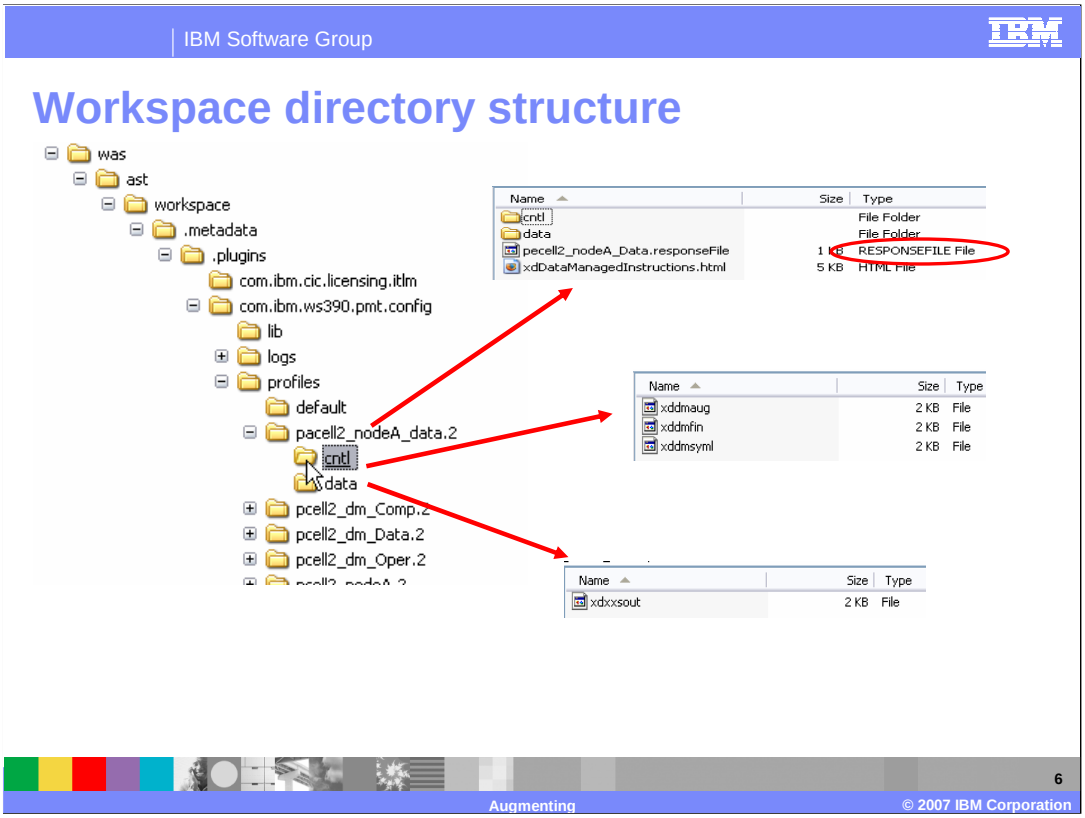

When working with the zPMT tool, your work is saved to a workspace you define. In this picture the workspace is “\was\ast\workspace”. This picture shows the directory structure inside the workspace. Each of your profile creates and augments will be kept inside of a directory under the directory “profiles”. In this picture, the profile pacell2_nodeA_data has been expanded. Here you see a response file. This file is very handy for creating multiple copies of similar profiles. A simple example is creating multiple empty managed nodes in a base WebSphere Network Deployment cell. Also under the profile directory, you see files that contain the various JCL that will be uploaded to your <HLQ>.CNTL and <HLQ>.DATA PDS’s on your target host system. The content of these files will be discussed later.

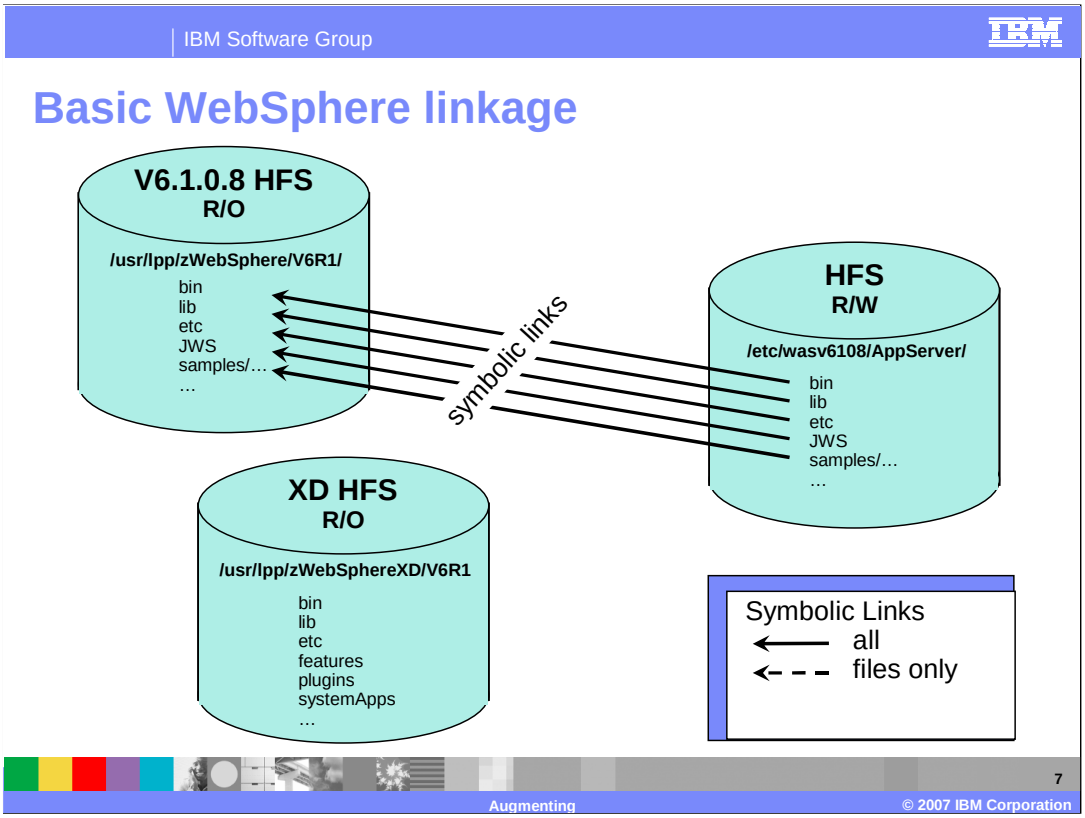

After configuring a base V6.0.2 Network Deployment cell and before converting the cell to an Extended Deployment cell, the picture appears as shown here. The SMP/E installation creates an HFS that contains all the WebSphere files that is referred to as the library file system or tree. The library HFS is shown in the upper left corner and is rooted at **/usr/lpp/zWebSphere/V6R1** as shown. When a WebSphere cell is configured, this library HFS is nominally mounted as read only. The configuration process creates one or more hfs's. One of these hfs's is shown on the right referred to as the configuration tree and is rooted at **/etc/wasv6108/AppServer** on a test system. The configuration tree has a complicated tree structure and only some of that structure is shown here. To save file system space, much of the configuration tree are soft links back to the library tree. In the basic case, directories are often linked back to directories. For example, if you were to look at the file **/etc/wasv61108/AppServer/etc** you would see that it is a soft link back to **/usr/lpp/zWebSphere/V6R1/etc**. In other cases, individual files are linked back to library files. For example, if you were to look at the file **/etc/wasv6108/AppServer/bin/wsadmin.sh** you would see it is a soft link back to **/usr/lpp/zWebSphere/V6R1/bin/wsadmin.sh**. The soft links provide sharing of library files in a natural way. This mixture of links is shown by the solid arrows in the picture above to indicate some links are for directories and some links are for files in a subdirectory tree. The first step to augmenting an Network Deployment configuration tree, is to add some soft links to the WebSphere HD library shown in the lower left of the picture above. But there is a problem as discussed on the next slide.

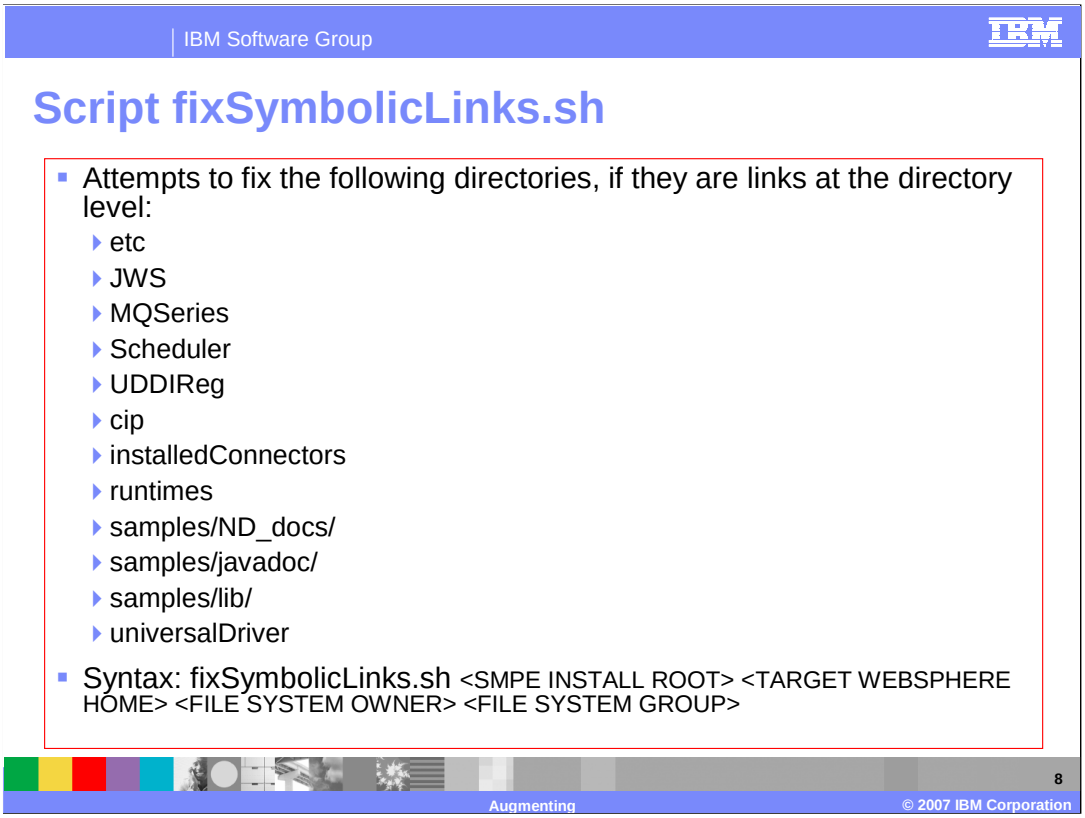

Because WebSphere Extended Deployment is a stack product, some care has to be taken about how the Extended Deployment configuration files are added onto a base WebSphere Network Deployment tree. In particular, if two or more stacked products are applied to a base Network Deployment configuration, neither can alter that tree in such a way as to interfere with another stack product. To see the problem, consider the directory <WAS_HOME>/etc. The WebSphere Extended Deployment install process needs to add some executables into this tree. Since etc is a soft link back to the WebSphere Network Deployment libraries, the quandary is where to put the Extended Deployment files. By rule, the Extended Deployment files can't go into the Network Deployment libraries and the augmentation can't break the soft link for etc. The solution is a standard way to alter soft links in the configuration tree. This is performed by a script in the WebSphere Network Deployment tree called fixSymbolicLinks.sh. The script breaks the soft links for the files listed above, mirrors the directory structure from the library tree to the configuration tree and places soft links back to the files in the library tree. Thus when you look in the configuration tree for etc, you will see a directory structure that looks exactly like the library tree except all the files have been replaced by soft links.

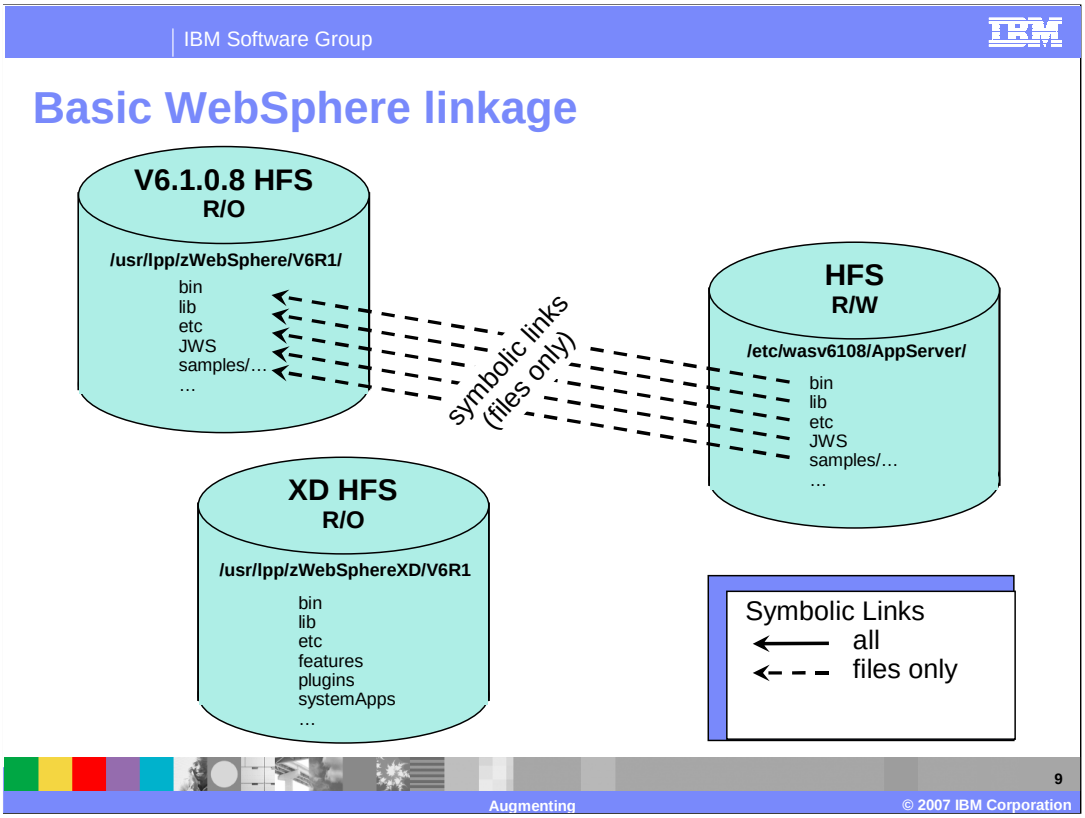

This is the same picture of the basic WebSphere Application Server HFS's you saw before except that all the links are now only to files. With this configuration, the augmentation jobs can add the Extended Deployment soft links to the configuration tree.

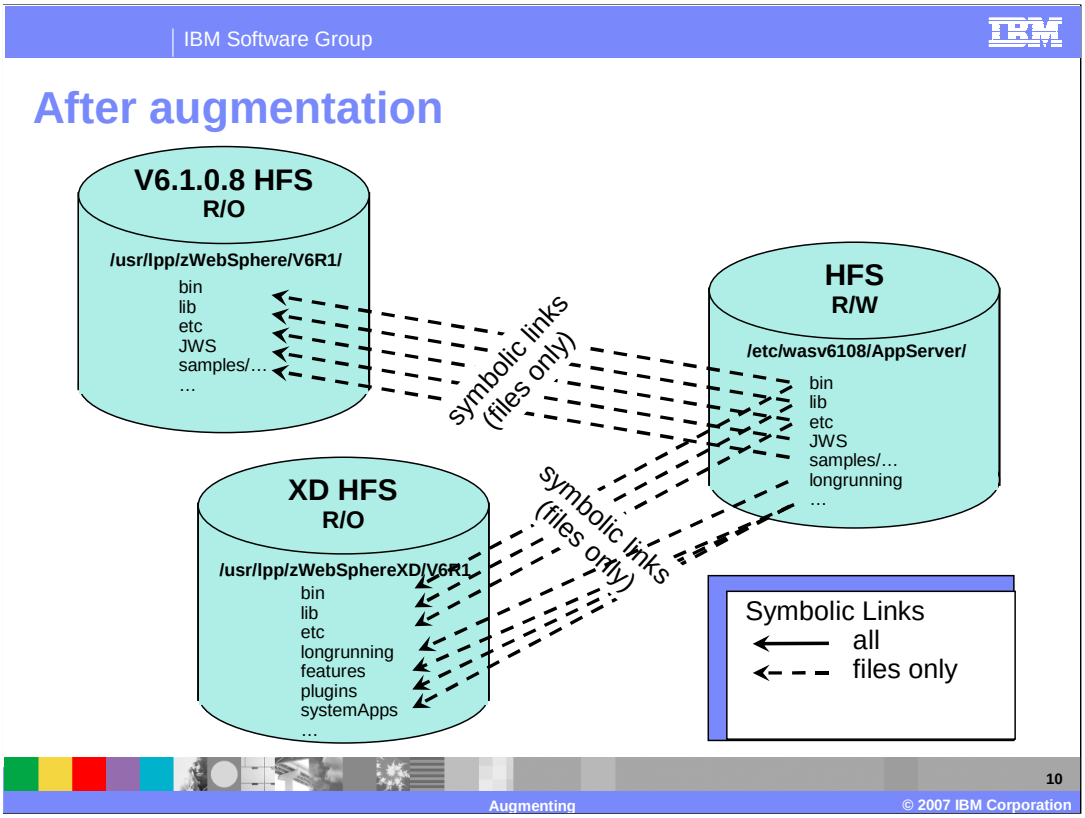

This slide shows the soft links after an augmentation. The augmentation has added soft links for WebSphere Extended Deployment specific files from the configuration tree to the WebSphere Extended Deployment library. A new directory called longrunning has been added to the configuration tree, and contains soft links back to WebSphere Extended Deployment specific files for you to use to perform functions like creating specific WebSphere Extended Deployment data bases. This is a sample of the types of files created and is not a complete list.

Agenda

- Augmentation overview
- Augmenting
- Troubleshooting
- Summary

The next section looks at the details of augmenting a cell.

Creating a new Network Deployment cell

- A deployment manager and empty managed nodes with either
  - ISPF panels
  - Or zPMT

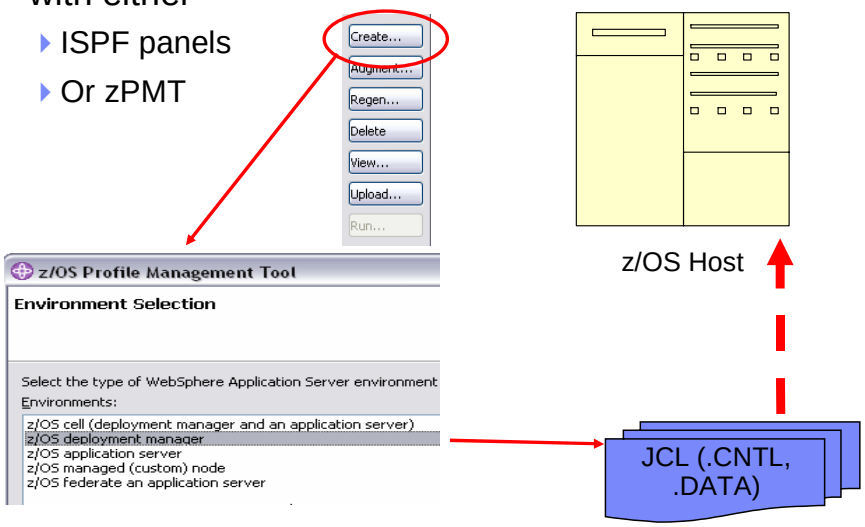

You can create a base Network Deployment cell with the ISPF panels or the zPMT tool as before. The picture shows the general process for creating a cell with the zPMT tool. Note that the windows shown during the profiles creation in the zPMT tool are very similar to the ISPF panels for performing the same task. Typically, a user will only create a deployment manager and one or more empty managed nodes – WebSphere Extended Deployment tends to create the required servers.

Augmenting

- Create augment jobs
  - Must use zPMT
  - One JCL package per sub-product
    - Compute grid
    - Data grid
    - Dynamic operations
  - One set of these per
    - Deployment manager
    - Node

Notice that you are going to create a lot of JCL to augment a cell. For example, if you have a deployment manager with three Nodes and a full Extended Deployment package, you will have to create 12 sets of JCL. Using the response files mentioned earlier can save some effort here.

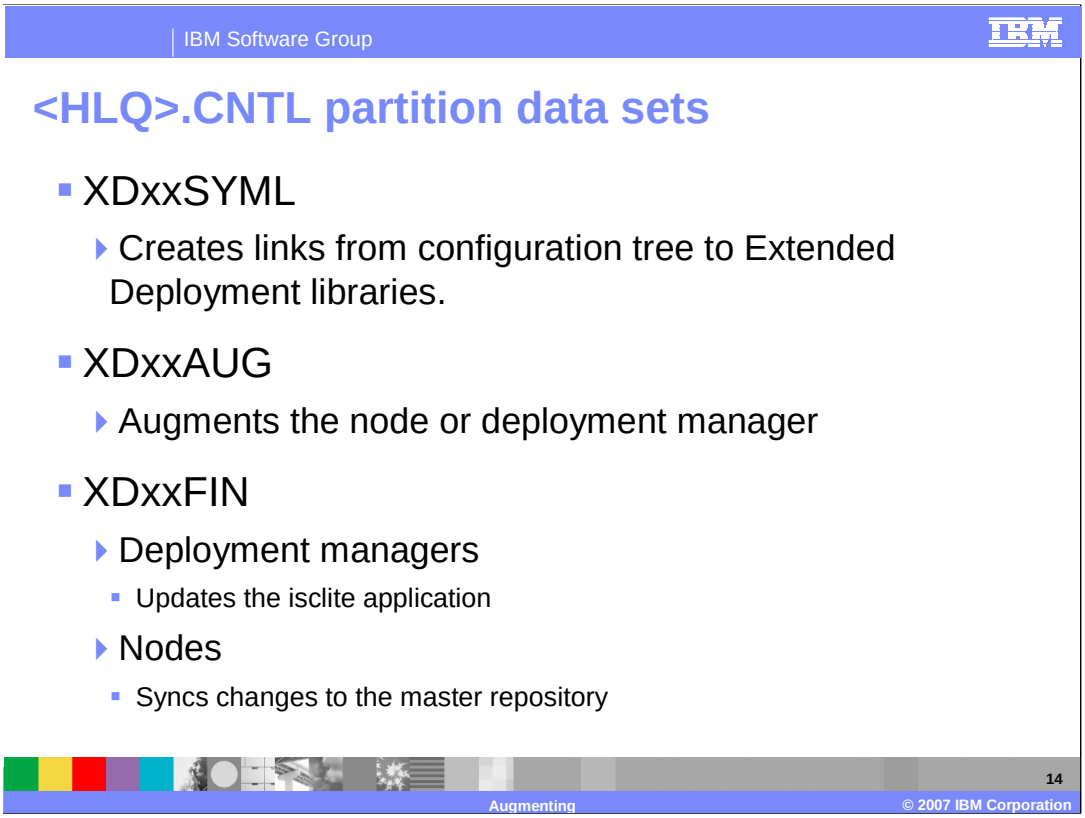

For all augmentations, there will be three files similar to these. The first two files perform the same operations for all augmentations and the node jobs do not affect the master repository. The third job performs different functions depending on whether a Deployment manager is being augmented or not. For a deployment manager, the job updates the administration console and the deployment manager can not be running during augmentation. It is probably best to augment all Extended Deployment components and then start the deployment manager. For Nodes the third job will sync changes to the configuration with the master repository and the deployment manager must be running. Note that you can run the XDxxSYML and XDxxAUG jobs for a single component on more than one node at the same time as an efficiency. However, the XDxxFIN jobs should be run sequentially to avoid a collision.

The <HLQ>.DATA contains a REXX script for copying an hfs file to the job log.

Viewing the instructions

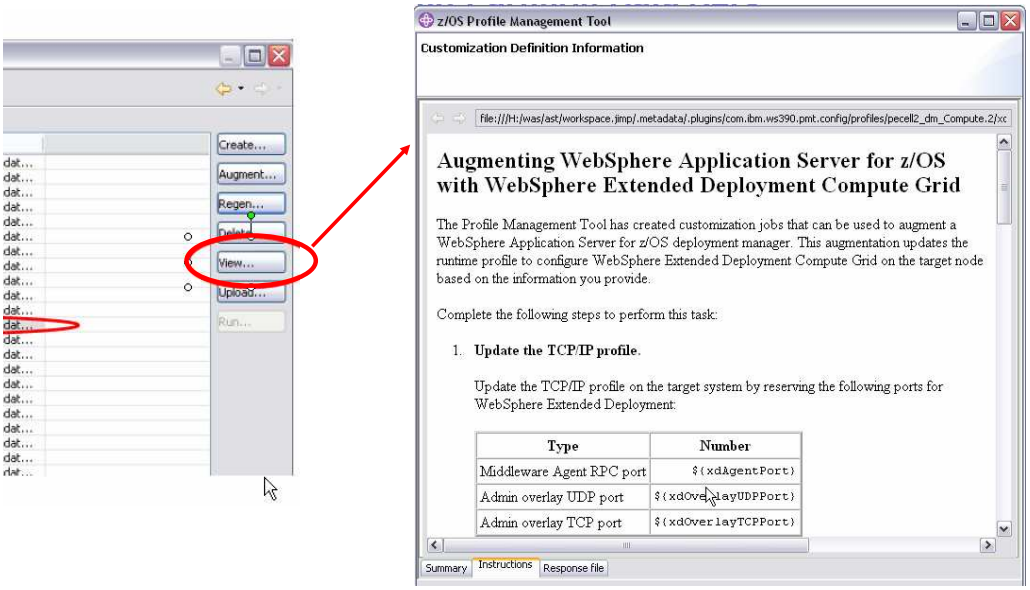

After creating and uploading the JCL, you can click on the “View” button and the “Instructions” tab to read the instructions. There is not a EEX*INS member so this is where you read augmentation instructions. Basically you will submit the three jobs described in these instructions per WebSphere Extended Deployment component for each deployment manager and node.

Agenda

- Augmentation overview
- Augmenting
- Troubleshooting
- Summary

This section presents where to find the augmentation logs.

Augmentation logs

- augmentProfile<*profilename*>.log
  - Logs for profile augmentation
  - Will contain reason for failure if the augment process fails

The installation logs are located in the logs subdirectory of the <WAS_HOME> directory that you selected for installation. The main installation log is named WAS.XD.install.log. If your installation fails, check this log for errors. Tasks launched by the installer, such as profile augmentation, also have their own logs.

Agenda

- Augmentation overview
- Augmenting
- Troubleshooting
- Summary

This section will provide a summary of this presentation.

Summary

- Extended Deployment V6.1 consists of three packages:
  - Operations optimization
  - Data grid
  - Compute grid
- Updated profile management tool used to create and augment Extended Deployment profiles.

Creating a WebSphere Extended Deployment cell requires that you first create a WebSphere Application Server V6.1.0.8 cell and then augment it, since WebSphere Extended Deployment is an add-on product, rather than a stand-alone product. Each of the three Extended Deployment packages must be augmented separately.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject= Feedback about XD61z_Augment.ppt

You can help improve the quality of IBM Education Assistant content by providing feedback.

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 MQSeries WebSphere z/OS

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 or changes in the products or programs 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 (for example, 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.

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.