



IBM Software Group

WebSphere® Application Server for z/OS® V7

Migration overview



@business on demand.

© 2008 IBM Corporation
Updated January 13, 2009

This presentation covers migrating to WebSphere Application Server for z/OS Version 7.

Agenda

- Migration and coexistence
- Overview of migration process
- WebSphere customization tools
 - ▶ z/OS migration management tool
- zmmt shell script



First, an overview of the migration process for WebSphere Application Server for z/OS is provided. Levels of software needed for migration and coexistence is also covered, followed by a basic overview of the migration process. Finally, tools available to help you migrate your WebSphere Application Server for z/OS environment are discussed.

Migration and coexistence – software levels

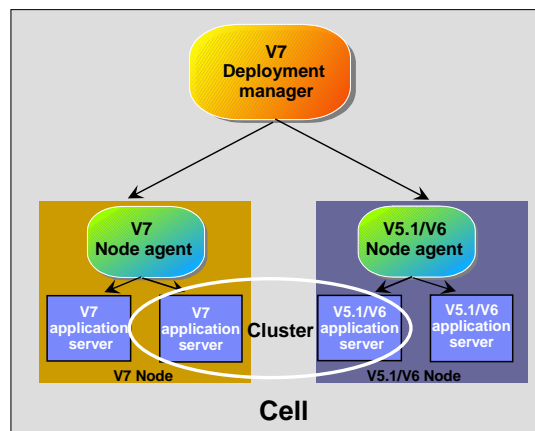
- WebSphere Application Server V5.1 and V7
 - ▶ Minimum level: 5.1
- WebSphere Application Server V6.0 and V7
 - ▶ Minimum level: 6.0.2.12, if security is enabled
- WebSphere Application Server V6.1 and V7
 - ▶ Minimum level: 6.1.0



There are some coexistence and prerequisite conditions that must be met before attempting a migration. The lowest release level of WebSphere Application Server that you can migrate to V7 is V5.1. If you are at a level lower than V5.1, you must first perform an intermediate migration to a level that is supported, such as Version 5.1.

V7 coexistence support: mixed version cell

- V7 mixed version cell support
 - ▶ Nodes at different versions are supported in the same cell
 - Supports V5.1, V6.0 and V6.1
 - ▶ Deployment manager must always be at V7 level
 - ▶ Must use the runtime migration tools to create the mixed version environment



5

Coexistence support refers to the ability to run multiple nodes of WebSphere Application Server on the same z/OS image or sysplex at the same time. This slide specifically shows the ability of V7 to coexist with multiple versions of WebSphere in the same cell. WebSphere Application Server Version 7 supports nodes configured at the V5.1, V6.0 and V6.1 levels in the same cell. However, the deployment manager must always be at V7. With mixed versions in a cell, you can minimize application downtime during migration because you can migrate one node at a time. If you have applications that run in a clustered environment, those applications can continue to run while the migration of one node takes place.

Migration on z/OS, two step process

- Install the product code
 - ▶ Using SMP/E
- Migrate configuration, one node at a time, by running customized jobs created:
 - ▶ Using WebSphere customization tools (z/OS migration management tool)
 - Workstation tool available for Windows® or Linux® Intel® operating systems
 - Available for download:
<http://www.ibm.com/support/docview.wss?rs=180&uid=swg24020368>
 - ▶ Using zmmt.sh command found in:
`/usr/lpp/zWebSphere/V7R0/bin`



In order to migrate your WebSphere environment, you must first install the Version 7 product code. This is done using SMP/E. Once the Version 7 code is installed, you can then proceed to migrate your cells, node-by-node. A stand-alone application server consists of one node while a network deployment configuration consists of at least two nodes, a deployment manager node and an application server node. In the network deployment configuration, you must always start with the deployment manager node. The migration process consists of running some customized jobs that are created for you with the help of two possible tools. The first tool is the z/OS migration management tool and it is part of the WebSphere Customization Tools. The WebSphere Customization Tools is a workstation tool that is available for the Windows or Linux Intel operating systems. It is available for download at the URL shown on the slide. The second option allows you to create the migration jobs completely on z/OS using a shell script, `zmmt.sh`, found in the product bin directory. Notice that the ISPF customization dialogs are no longer available starting with Version 7 of the product.

SMP/E

- Program directory available here:
<http://www.ibm.com/software/webservers/appserv/was/library/v70/was-zos/books.html>
- Requires SDK 1.4 or higher
- UNIX® must be in full function mode
- Userid to run some jobs needs authority to mount file systems and modify UNIX System Services files
 - ▶ Must also have read access to the facility class resources BPX.FILEATTR.PROGCTL, BPX.FILEATTR.APF, and BPX.FILEATTR.SHARELIB.




To install the product code, you will use SMP/E. The program directory for the product is found at the URL listed on the slide. Because WebSphere Application Server uses the SDK as part of the SMP/E processing, you must have SDK 1.4 or higher on the system where you plan to do the installation. UNIX must also be configured in full function mode in order to complete the product code installation. To run some of the SMP/E jobs, the user ID used must have the authority to mount file systems and modify UNIX System Services files and have access to the facility class resources listed.

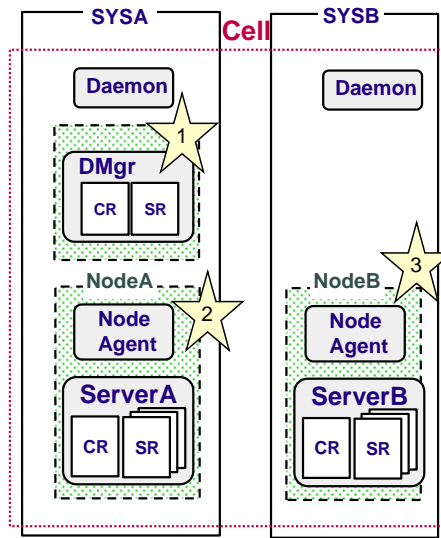
Overview of migration to Version 7

Migration is node-by-node

- ▶ From WebSphere Application Server V5.1, V6.0 or V6.1
- ▶ V5.1/V6.0/V6.1 configuration HFS remains intact for fall-back

Order of migration for network deployment cell:

- ▶ Migrate the deployment manager first 
- ▶ Update supporting software
- ▶ Migrate application server nodes  ↔ 



9

Migration overview

© 2008 IBM Corporation

Migration to Version 7 is done on a node by node basis, from a Version 5.1, Version 6 or Version 6.1 system. You cannot migrate from a Version 5 or lower system directly to Version 7. You must first migrate to a supported release. It is important to note that the previous release's HFS remains intact to allow for fall-back if necessary. On a network deployment cell, the deployment manager must be migrated first as denoted by the numbered stars. Next, you need to update any supporting software such as web servers, webserver plugins and databases. Finally, you migrate the application server nodes. Note that while the slide shows NodeA as the second node to be migrated, the order can actually be reversed and NodeB could be the second node migrated.

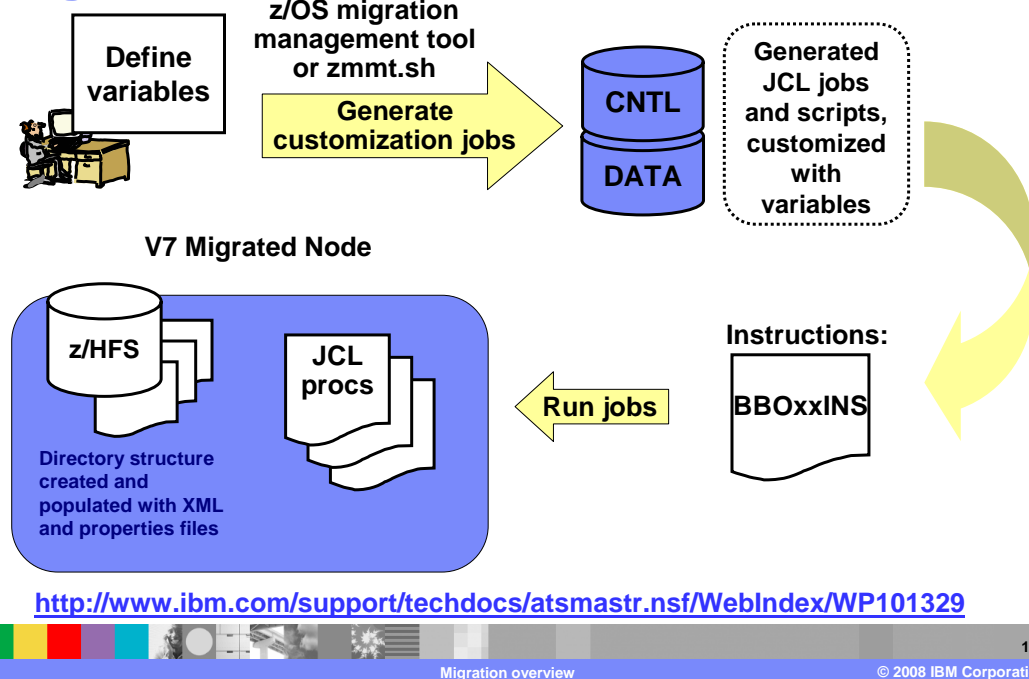
Overview of migration to Version 7

- Three types of nodes
 - ▶ Stand-alone application server node
 - ▶ Deployment manager node
 - ▶ Federated (managed) node



Migration is performed one node at a time and there are three types of nodes to consider. The stand-alone application server node consists of only one node so the migration is relatively quick and simple. The applications on that node are unavailable until the migration completes successfully and the server is restarted. A deployment manager configuration allows you to minimize application downtime during migration while you work on one node at a time. The deployment manager must be migrated first and then you can migrate each managed node one at a time. As long as your application servers are clustered, applications can remain running as you migrate one node at a time within the cluster.

Migrate node

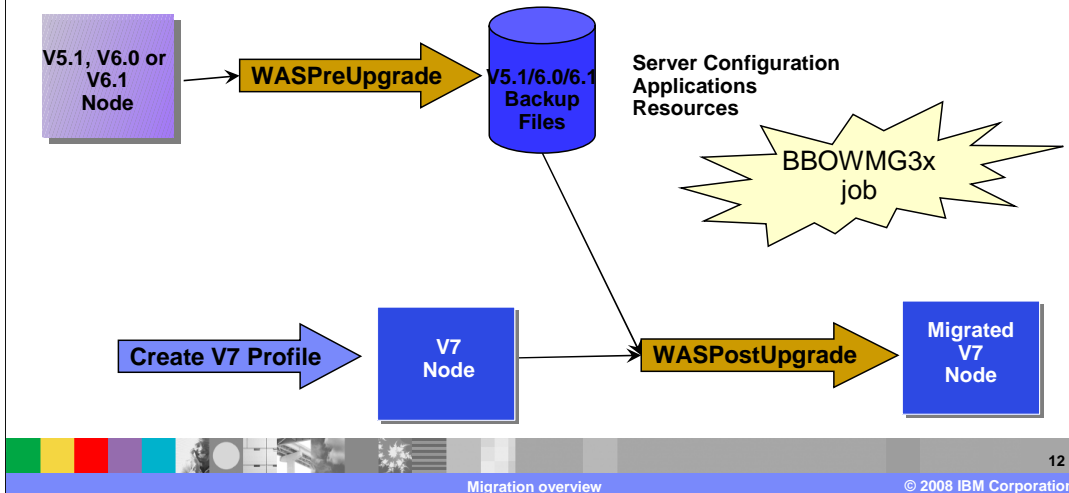


Once you have determined what type of node you are migrating, you will define variables that reflect your naming conventions and other choices you have made. Using either the z/OS migration management tool or the zmmt.sh script, you will generate customization jobs and scripts based on those variables. The jobs and scripts will end up in the CNTL and DATA PDSes on the mainframe where they can be run. In the CNTL PDS, you will find a BBOxxINS member, where the xx will depend on what type of environment you are migrating. The instructions to run the jobs are found there and they will tell you what jobs need to be run and in what order, along with the user ID that is needed to run the job. Once you have successfully run all the jobs, you will have a migrated V7 node with servers that are ready to start.

A migration techdoc is also available at the URL shown on the slide. It is a very detailed look at the entire migration process and is highly recommended to help you plan your migration.

Overview of migration to Version 7

- Run **WASPreUpgrade** command on the V5.1, V6.0 or V6.1 level that needs migration
- Run **WASPostUpgrade** on new V7 installation, pointing to the output produced by WASPreUpgrade – result is a migrated V7 environment



This slide shows the undercover processing that occurs on z/OS. The migration tools provide customized batch jobs that are run instead and the main migration job, BBOWMG3x, takes care of the steps shown here. The BBOWMG3x job will create the V7 default profile for you and populate the z/HFS as required. It also runs the WASPreUpgrade command to backup some files in the HFS to be used by the WASPostUpgrade step. The WASPostUpgrade step takes the backup configuration and converts it to the new V7 default profile, changing it as required for the new release. At the end of the job, you are left with a fully migrated V7 node.

Migration jobs

Create HFS or zFS	BBOMxZFS or BBOMxHFS
Copies tailored JCL to PROCLIB	BBOMxCP
Clear the transaction logs (for XA connectors only)	BBOWMG1x
Disable Peer Restart and Recovery (PRR) mode (for XA connectors only)	BBOWMG2x
Perform migration	BBOWMG3x



This slide shows the jobs that must be run with an explanation of what they do. The basic sequence of jobs are the same for each of the different types of nodes. The names, as denoted by the lowercase 'x', change slightly in each environment. The BBOWMG3x job is long running and can cause certain error conditions such as an ABEND 522. TIME=NOLIMIT on the JCL job card solves the problem. Also note that the BBOWMG1x and BBOWMG2x jobs are only needed if you have any XA connectors defined in your configuration. They do not apply to the deployment manager node migration.

BBOWMG3x job

- Multi-step, long-running job:
 - ▶ Create working directory (`/tmp/migrate/nnnnn`)
 - ▶ **WRCONFIG** – copy dialog generated variables to the HFS
 - ▶ **WRRESP** – create a profile creation response file from dialog generated variables
 - ▶ **MKCONFIG** – gather information from existing configuration (for instance, cell name, server name)
 - ▶ **VERIFY** – verify the variables generated from dialog
 - ▶ **CRHOME** – create a V7 WAS_HOME structure
 - ▶ **CRPROF** – create V7 default profile
 - ▶ **PREUPGRD** – backup some files in the HFS to be used by WASPostUpgrade
 - ▶ **UPGRADE** – run WASPostUpgrade to perform the migration (`serverindex.xml` renamed to `serverindex.xml__disabled`)
 - ▶ **FINISHUP** – run Config2Native, update file permissions and attributes

14

Migration overview

© 2008 IBM Corporation

To see a little more detail about the main migration job, BBOWMG3x, all the steps are listed here. A working directory in `/tmp` is used to do much of the processing. The `nnnnn` is a unique number that was generated during the creation of your migration jobs. For normal migration, the space used in `/tmp` is very small but if you turn on tracing, the space used can be quite large. Make sure you have the free space on `/tmp`. The rest of the jobs are explained on the slide. The **UPGRADE** step is where the actual migration occurs and will take the longest to complete. The **VERIFY** step before that attempts to check the information provided so that the migration does not fail because of bad input parameters.

BBOWMG3x job...




- May be run as three steps:

- ▶ Create working directory (/tmp/migrate/nnnnn) Common tasks
- ▶ **WRCONFIG** – copy dialog generated variables to the HFS
- ▶ **WRRESP** – create a profile creation response file from dialog generated variables
- ▶ **MKCONFIG** – gather information from existing configuration (for instance, cell name, server name)
- ▶ **VERIFY** – verify the variables generated from dialog
- ▶ **CRHOME** – create a V7 WAS_HOME structure 1
- ▶ **CRPROF** – create V7 default profile
- ▶ **PREUPGRD** – backup some files in the HFS to be used by WASPostUpgrade 2
- ▶ **UPGRADE** – run WASPostUpgrade to perform the migration (serverindex.xml renamed to serverindex.xml__disabled)
- ▶ **FINISHUP** – run Config2Native, update file permissions and attributes 3



The BBOWMG3x job may actually be run as three different jobs. They are supplied in the CNTL dataset for you. We will look at the names on the next slide but note that the first steps are repeated in each of the jobs.

BBOWMG3x job...

- ▶ **BBOWxPRO**
 - Common tasks +  (CRHOME, CRPROF)
- ▶ **BBOWxPRE**
 - Common tasks +  (PREUPGRD)
- ▶ **BBOWxPOS**
 - Common tasks +  (UPGRADE, FINISHUP)



Here you see the name of the three jobs corresponding to the numbers on the previous slide. The BBOWxPRO will create your configuration structure. The second job, BBOWxPRE, will do some processing to prepare for the actual upgrade process which occurs when the third job, BBOWxPOS, is run. Again, these jobs are included in the CNTL dataset for your use.

General recommendations

- **Keep the same procedure names**
 - ▶ Backup (save) your previous procedures
 - ▶ If you use different procedure names, RACF® profile updates required:
 - STARTED class profiles
 - Automation changes required
- **Use separate z/HFS for each V7 node** (even if you did not in prior versions)
 - ▶ This may require new procedure names if you had a shared HFS in the previous version
- **Review information in the 'Premigration considerations' article found in the information center:**

http://publib.boulder.ibm.com/infocenter/wasinfo/v7r0/topic/com.ibm.websphere.migration.zseries.doc/info/zseries/ae/cmig_pre.html



In order to simplify things, when planning for your migration, you should use the same procedure names. Before updating the procedures for Version 7 though, you should save your current procedures in case you need to fallback to the previous level. If you choose to use different procedure names, you will need to update the RACF STARTED class profiles. Sample RACF commands to accomplish this are found in the migration instructions provided. Keep in mind that automation changes may also be required when changing procedure names. You should also use a separate HFS for each Version 7 node. This might require new procedure names if you used a shared HFS in previous versions..

Migration troubleshooting

- Check the BBOWMG3x output for errors:
 - ▶ /tmp/migrate/*nnnnn*/BBOWMG3x.out and /tmp/migrate/*nnnnn*/BBOWMG3x.err written to JOBLOG
- Check the /tmp/migrate/*nnnnn*/logs directory for logs named WAS*Upgrade*<timestamp>.log
- Turn traces on for more information about failure
 - ▶ 'xxxx.DATA(BBOWMxEV)' can be updated to enable:
 - TraceState=enabled
 - profileTrace=enabled
 - preUpgradeTrace=enabled
 - postUpgradeTrace=enabled
- If job fails in the VERIFY step, correct error and rerun job
 - ▶ Most likely incorrect data was entered
- If job fails after the VERIFY step, correct error and rerun job
 - ▶ Must delete WAS_HOME directory that was created in CRHOME step first
 - ▶ Possibly need to rename serverindex.xml_disabled file back to serverindex.xml in original configuration

If the BBOWMG3x job fails, check the output for errors. If you need more information, you can turn traces on by editing the BBOWMDRF member in the DATA PDS. The trace states are disabled, by default, although you can also change them in the z/OS migration management tool when going through the configuration. If the job fails in the VERIFY step, it is most likely that you made an error when specifying information used to create the jobs. Correct the information and rerun the job. If the job fails after the VERIFY step, you need to delete the WAS_HOME directory that was created in the CRHOME step before re-running the job. Check the original configuration for the serverindex.xml file being renamed to serverindex.xml_disabled also. This is done to signal that the configuration has already been migrated so as to stop you from inadvertently migrating the node again. This is done by default but it is possible to change this behavior during the configuration phase. It is a check box in the z/OS migration management tool or you can set the keepDMGREnabled parameter to true in the response file.

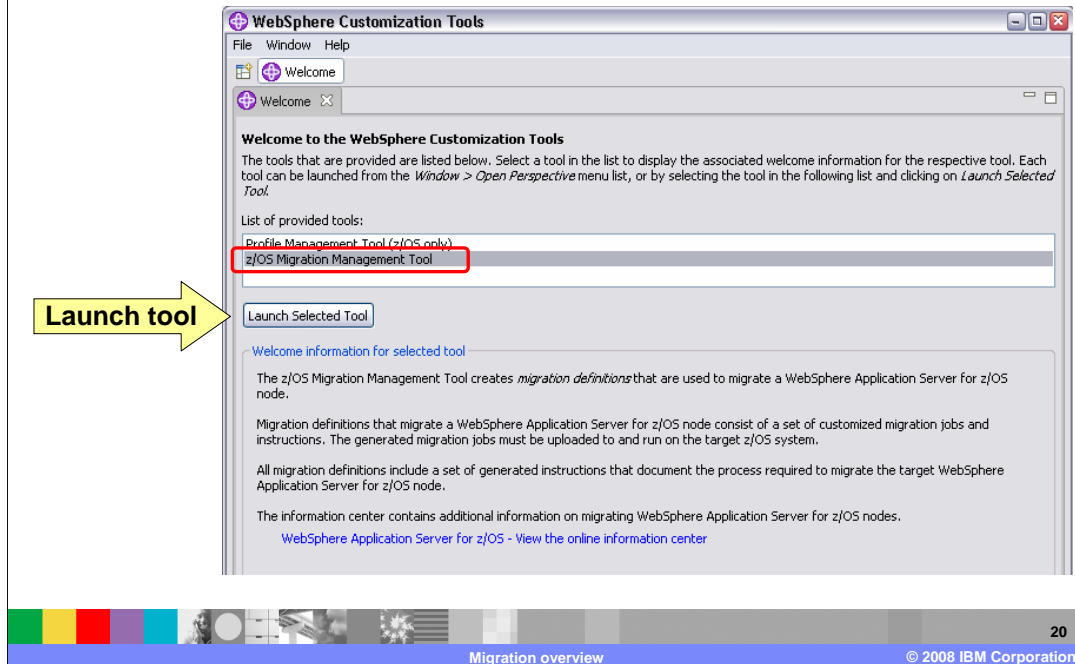
Section

z/OS migration management tool



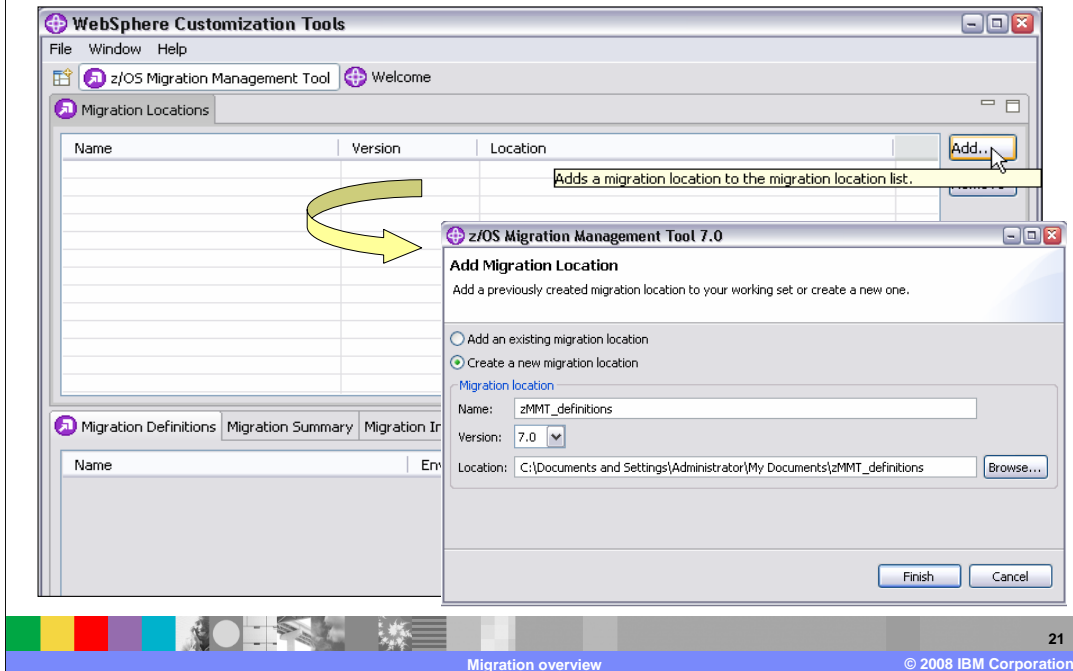
This section briefly covers the z/OS migration management tool.

z/OS migration management tool

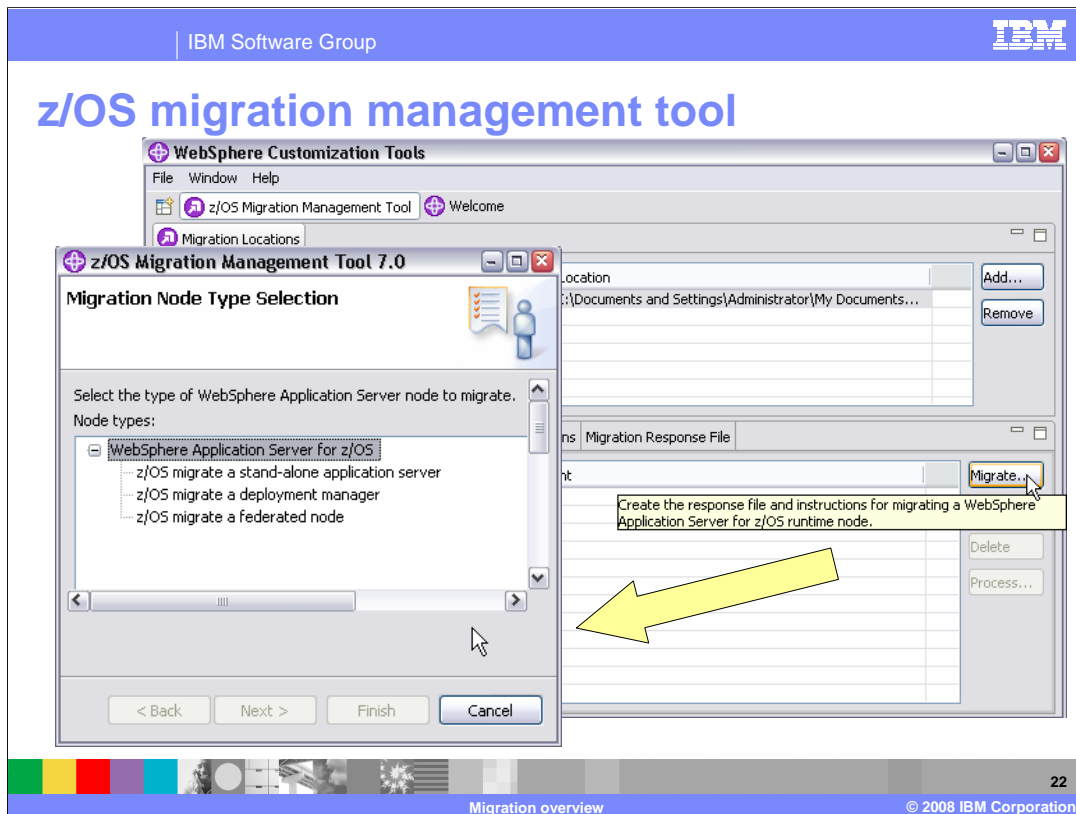


Now that you have seen an overview of the migration process, you will look at the tools available to do the migration. The first one is the z/OS migration management tool. As mentioned before, the z/OS migration management tool is part of the WebSphere customization tools that are available for the Windows or Linux Intel operating systems. The Welcome screen is shown on this slide and you must select the z/OS migration management tool and then launch it.

z/OS migration management tool



Once the tool has been launched, the first thing you must do is add a migration location. This is the location where the tool will keep the customization definitions you will define. This is shown in the smaller box on the slide.



Once you have a migration location defined, you can customize an environment of your choosing. Seen in the smaller box on the left are the various types of nodes you can migrate as noted earlier. You will need information from your down-level configuration to complete the panels before having the customized jobs created. Plan ahead!

z/OS migration management tool

The screenshot displays the 'WebSphere Customization Tools' environment with the 'z/OS Migration Management Tool' open. The tool has several tabs: 'Migration Definitions', 'Migration Summary', 'Migration Instructions', and 'Migration Response File'. The 'Migration Definitions' tab is active, showing a table with the following data:

Name	Environment
DeploymentManager migration	z/OS migrate a deployment manager

Buttons for 'Migrate...', 'Regen...', 'Delete', and 'Process...' are visible on the right. A red box highlights the 'Process...' button, with a tooltip that reads 'Process the selected migration definition.' A yellow arrow points from the 'Process...' button to a 'Select Process Type' dialog box. To the left of the dialog box, there are two blue cylinders representing 'CNTL' and 'DATA' files.

Select Process Type

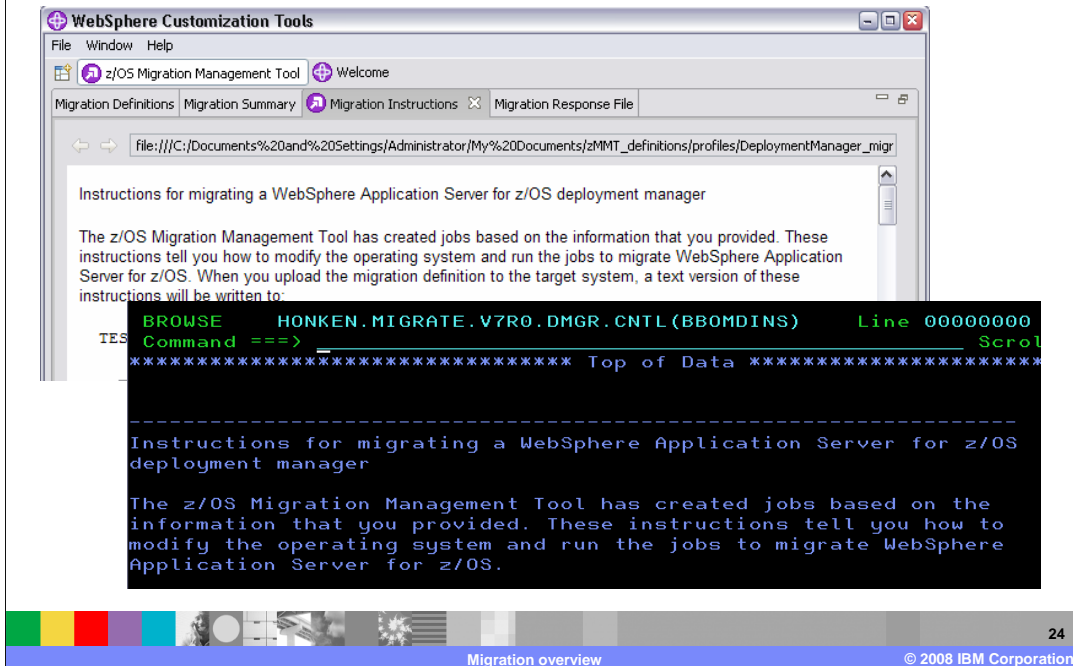
Select the type of processing to perform on the migration definition:

- Upload to target z/OS system
Create the migration jobs for the selected migration definition and upload them to a z/OS system. (This option requires an active ftp server on the target z/OS system.)
Note: If the migration data has been previously exported to the default directories, the migration jobs in these directories will be uploaded to the target z/OS system.
- Export to local file system
Create the migration jobs for the selected migration definition and export them to the local file system.

23
Migration overview © 2008 IBM Corporation

Once all the fields have been filled in and the migration definition created, you have the option of uploading the migration data which will create the CNTL and DATA PDS files up on the host. This is the process option highlighted on the slide.

z/OS migration management tool



Instructions on finishing the migration are found in both the z/OS migration management tool and on the host in the BBOxxINS member for the environment you are migrating. It will give you some possible manual instructions to complete and a list of jobs that need to be run. These manual instructions might include some additional security configurations that were introduced in V6.1 and some STARTED profiles that need to be defined if you changed your procedure names during the migration.

Section

zmmt.sh command

This section covers the zmmt.sh command alternative.

zmmt.sh command alternative

- Allows migration to fully happen on the mainframe
- Requires a fully-populated response file
 - ▶ z/OS migration management tool creates one
 - ▶ Examples found in the information center
- Will create the .CNTL and .DATA files needed to run the required jobs.



An alternative to running the z/OS migration management tool on the workstation is the zmmt.sh command. The zmmt.sh command runs on z/OS. There is no GUI interface to allow you to fill in your configuration variables. A response file that has been fully-populated is required. You can use the response file that is created when you use the z/OS migration management tool or you can find samples of the response files needed in the information center. When the shell script is run, it will create the CNTL and DATA files needed to run the required jobs.

zmmt.sh command alternative

- Found in the `/usr/lpp/zWebSphere/V7R0/bin` directory
- Parameters:
 - ▶ **-responseFile**
 - Specifies the path to your response file (ASCII or EBCDIC)
 - ▶ **-profilePath**
 - Fully qualified path name to an existing set of generated jobs This parameter cannot be used in combination with the `-responsefile` option.
 - ▶ **-workspace**
 - Specifies the Eclipse work space directory
 - ▶ **-transfer**
 - Copy generated jobs from a UNIX System Services (USS) file system to a pair of partitioned datasets. The **zmmt** command first writes the customization jobs to a USS file system.
 - ▶ **-allocate**
 - Attempts to allocate the target datasets

27

Migration overview

© 2008 IBM Corporation

The `zmmt.sh` command is found in the `bin` directory of the SMP/E install root. The `-responseFile` parameter is necessary to customize the resulting jobs and scripts to your installation's specifications. It can be coded in either ASCII or EBCDIC. Note that the sample shipped with the product is in ASCII. The resulting jobs and scripts are generated to the UNIX System Services file system. The `-transfer` command is necessary to actually move the generated jobs to CNTL and DATA PDS files. The `-profilePath` command can be used in conjunction with the `-transfer` command if you need to create the PDS files later. The `-workspace` parameter needs to specify a directory that you have read-write access to and is used as a work directory for Eclipse. Finally, the `-allocate` parameter determines whether it will attempt to allocate the CNTL and DATA PDS files first.

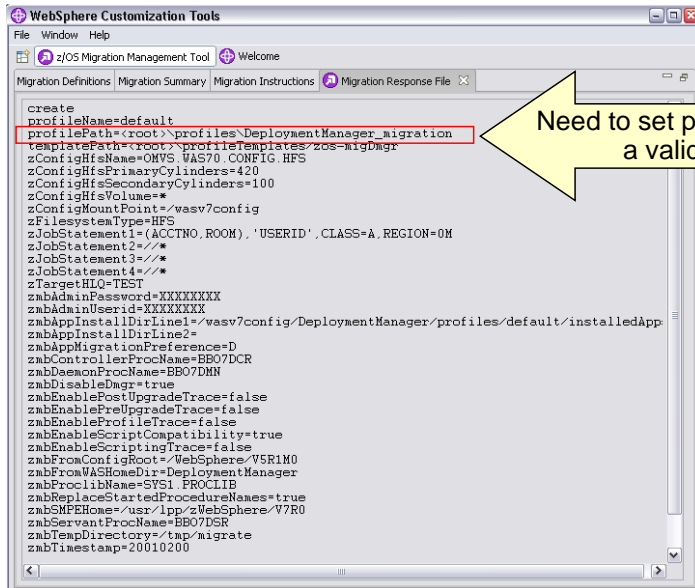
zmmt.sh command alternative example

- `zmmt.sh -workspace /xxx -transfer -allocate -responseFile /xxx/ZCellcmd.responseFile`
 - ▶ Generate the migration jobs to the location specified by `profilePath` in the response file
 - ▶ Allocate the target CNTL and DATA datasets, using the high level qualifier specified by `targetHLQ` in the response file
 - ▶ Transfer the jobs from the file system to the CNTL and DATA datasets



An example of the `zmmt.sh` command is shown here. Once the CNTL and DATA datasets are created, you can follow the directions found in the `BBOxxINS` member of the CNTL dataset to finish your configuration.

zgmt.sh command alternative response file



```

create
profileName=default
profilePath=<root>\profiles\DeploymentManager_migration
templatePath=<root>\profileTemplates\zos-migDmgr
zConfigHfsName=OMVS.WAS70.CONFIG.HFS
zConfigHfsPrimaryCylinders=420
zConfigHfsSecondaryCylinders=100
zConfigHfsVolume=*
zConfigMountPoint=/vasv7config
zFilesystemType=HFS
zJobStatement1=(ACCTNO,ROOM), 'USERID', CLASS=A, REGION=0M
zJobStatement2=/**
zJobStatement3=/**
zJobStatement4=/**
zTargetHLQ=TEST
zmbAdminPassword=XXXXXXXXX
zmbAdminUserid=XXXXXXXXX
zmbAppInstallDirLine1=/vasv7config/DeploymentManager/profiles/default/installedApp
zmbAppInstallDirLine2=
zmbAppMigrationPreference=D
zmbControllerProcName=BB07DCR
zmbDaemonProcName=BB07DMM
zmbDisableDegr=true
zmbEnablePostUpgradeTrace=false
zmbEnablePreUpgradeTrace=false
zmbEnableProfileTrace=false
zmbEnableScriptCompatibility=true
zmbEnableScriptingTrace=false
zmbFromVASHomeDir=/WebSphere/V5R1M0
zmbFromVASHomeDir=DeploymentManager
zmbFromLibName=SYS1.PROCLIB
zmbReplaceStagedProcedureNames=true
zmbSMPEHome=/usr/lpp/zWebSphere/V7R0
zmbServantProcName=BB07DSR
zmbTempDirectory=/tmp/migrate
zmbTimestamp=20010200
  
```

Need to set profilePath to a valid path



This slide shows the response file as generated by the z/OS migration management tool workstation tool. This can be uploaded and used to run the zgmt.sh command on z/OS. Note that you need to update the profilePath parameter here to point to a valid directory on your system.

zmmt.sh notes

- When run, you are put in the osgi command shell

- ▶ Will look like nothing is happening

```
/usr/lpp/zWebSphere/V7R0/bin#>rkspace -responseFile /tmp/zDMgr01.responseFile  
osgi>
```

- ▶ Will eventually come back with messages

```
osgi> Customization definition successfully written to /tmp/ZDMgr01  
Attempting to allocate dataset: BOSS.VICOM.BOSS0173.CNTL  
Allocation successful.  
Attempting to allocate dataset: BOSS.VICOM.BOSS0173.DATA  
Allocation successful.  
Copying CNTL files to BOSS.VICOM.BOSS0173.CNTL...  
Copy successful.  
Copying DATA files to BOSS.VICOM.BOSS0173.DATA...  
Copy successful.
```

When you run the zmmt command, you are thrown into the osgi command shell, where you will see the osgi command prompt for what seems like quite some time and you might think nothing is happening. Be patient and it will eventually come back with messages indicating either an error occurred or the definition was successfully created.

zmmt.sh notes

- When rerun, delete the profilePath directory
 - ▶ Otherwise, see message indicating the ProfilePath is not valid

```
osgi> The following validation errors were present with the command line arguments:  
profilePath: The profile path is not valid.
```



If you need to rerun the zmmt command, be sure to delete the directory you specified on the profilePath parameter. If you do not, you will see a validation error claiming the profilePath is not valid.

Summary

- Migration from V5.1, V6.0 and V6.1 supported
- Tools provided to help customize the process
 - ▶ z/OS migration management tool
 - ▶ zmmt.sh
- Performed on a node-by-node basis



This presentation covered the process of migrating an existing V5.1, V6.0 or V6.1 configuration to Version 7. An overview of the migration process and the tools provided to help in the process was also discussed. The migration must be done on a node-by-node basis, allowing for continuous availability of applications in a clustered environment.

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_WASv7zOS_MigrationOverview.ppt

This module is also available in PDF format at: ..\\WASv7zOS_MigrationOverview.pdf



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 RACF WebSphere z/OS

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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.

© Copyright International Business Machines Corporation 2008. 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.

