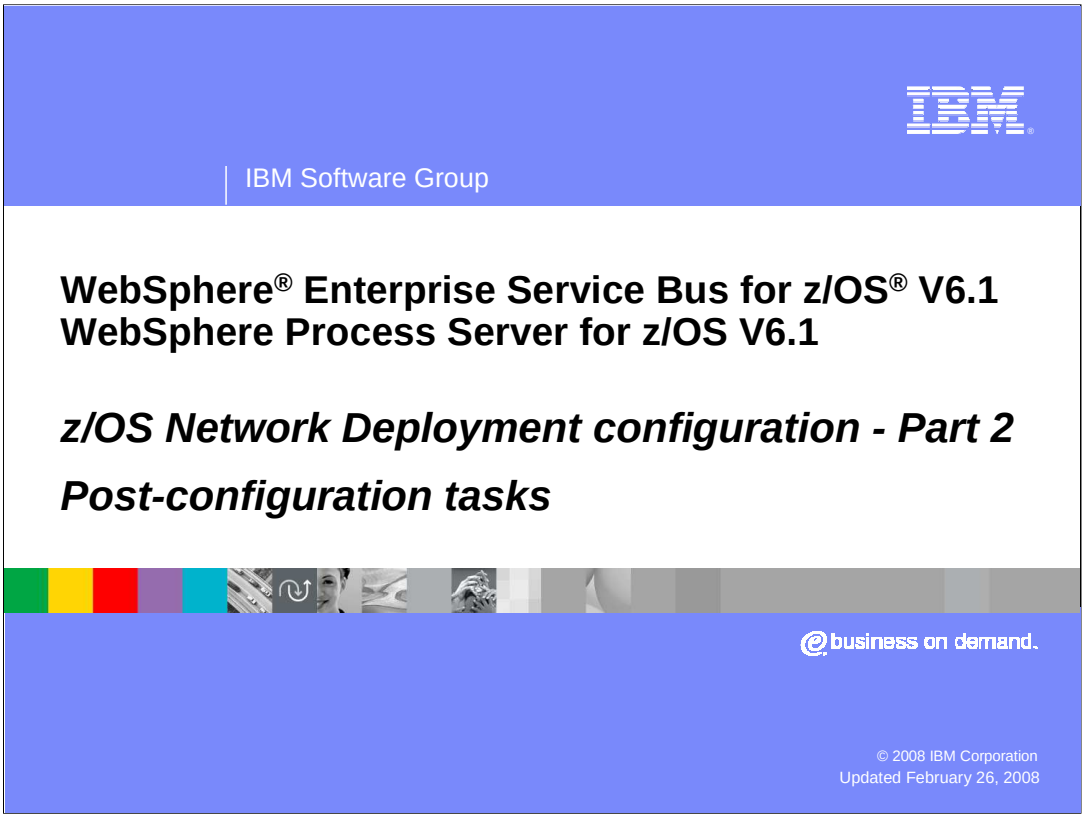

This presentation covers the post-configuration steps for setting up a Network Deployment cell to enable WebSphere Process Server for z/OS V6.1 or WebSphere Enterprise Service Bus for z/OS V6.1 function in its servers or clusters.

Goals

- Describe WebSphere Process Server for z/OS V6.1 and WebSphere Enterprise Service Bus for z/OS V6.1 configuration process using a Network Deployment configuration scenario

The goal of this presentation is to show you how to complete the configuration of WebSphere Process Server for z/OS V6.1 and WebSphere Enterprise Service Bus for z/OS V6.1 in a Network Deployment environment.

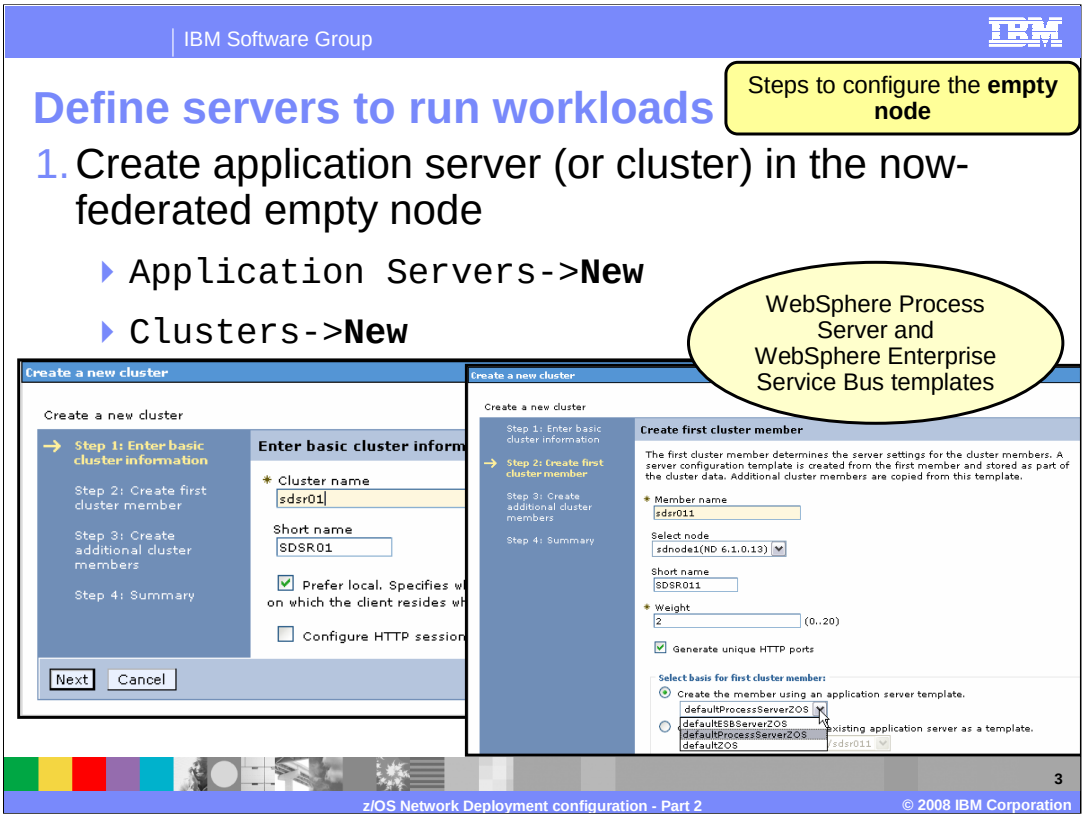

Part one of this topic suggests that you should configure an empty node rather than the stand-alone node in the network deployment scenario. This presentation covers what is needed for that scenario. The first thing needed is a server or cluster to run a WebSphere Process Server or WebSphere Enterprise Service Bus workload. You will notice new server templates are available for the WebSphere Process Server and the WebSphere Enterprise Service Bus functions as shown here, creating a new cluster.

Configure service component architecture

2. Add support for service component architecture components

```
Application servers > <serverName> >
   Service Component Architecture
Clusters > <clusterName> >
   Service Component Architecture
```

Found under Business Integration

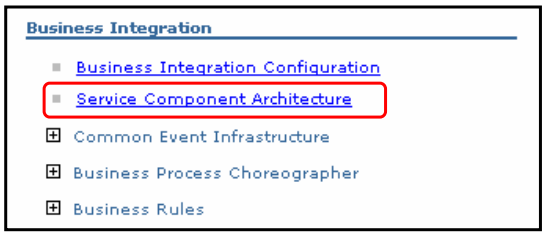

WebSphere Process Server applications or WebSphere Enterprise Service Bus applications need the service component architecture customized. You can configure it in the administrative console as shown. You need to fill in some information about the database you plan to use for the data stores that the messaging engines need. The messaging engines are created for you.

Configure service component architecture…

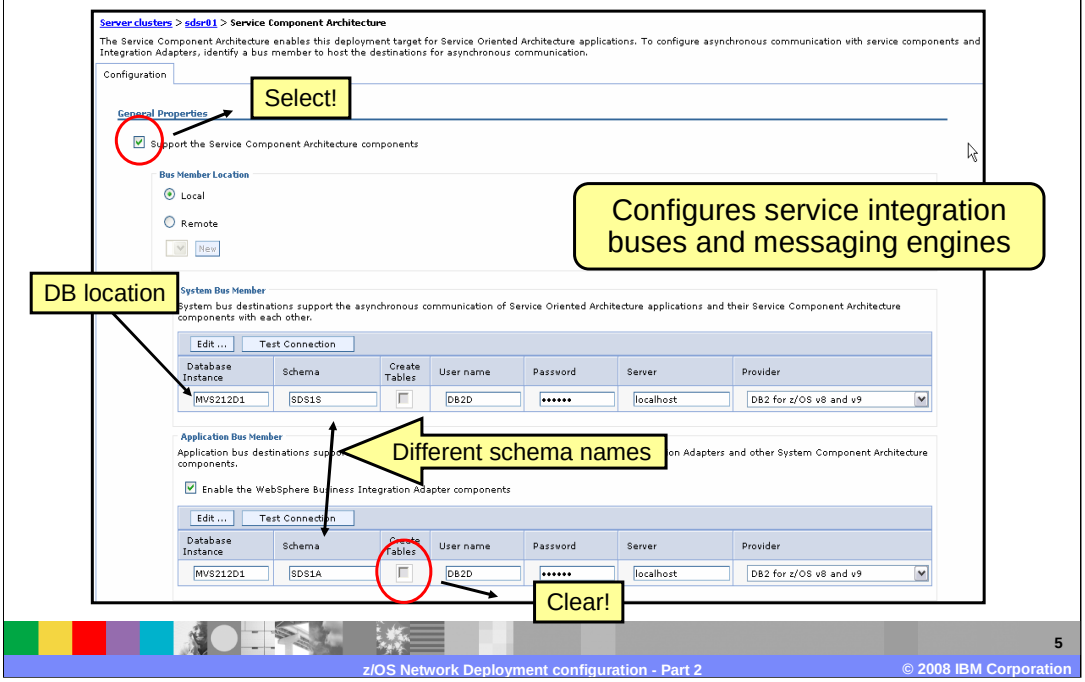

Two possible service integration buses are configured for you. These are the SCA system bus and the SCA application bus. The SCA system bus supports the asynchronous communication between service oriented architecture components and is required. The SCA application bus is optional and is used for asynchronous communication between WebSphere Business Integration adapter components and service oriented architecture applications. Each bus will have a messaging engine created for it. In order to configure the service component architecture, the first thing you need to do is select the check box near the top of the screen, indicating that you want SCA supported. You then need to fill in some information about the database that you plan to use for the messaging engines that are created for you. Keep in mind that the schema names must be different since the data stores have common table names for each of the buses.

Section

Configure business process and human task containers

The next several slides deal with configuring the business process and human task containers

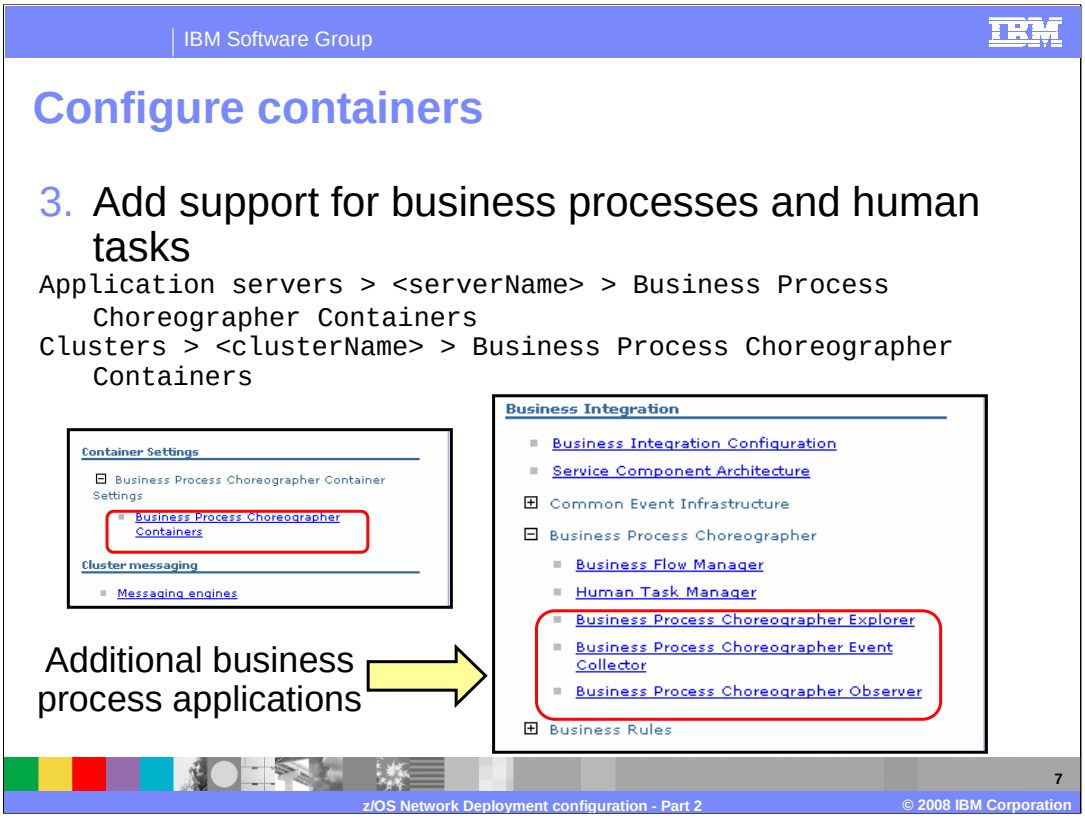

If you will be hosting applications that contain business processes or human tasks, you need to configure the business process choreographer containers. This is only available with the WebSphere Process Server product and is not be available with the WebSphere Enterprise Service Bus product. This slide shows where to find the installation wizard from the administrative console under the server or cluster that you configured. A later slide looks at the configuration using the bpeconfig.jacl script, which gives you more options. Also shown on the slide, under the Business Integration heading in the administrative console are some applications that you might want to install in order to administer or monitor your applications that contain business processes. These include the Business Process Choreographer Explorer, Event Collector and Observer.

The Business Process Choreographer Explorer is a Web application that implements a generic Web user interface for interacting with business processes and human tasks. There is information in the information center on customizing this to suit your needs.

The Business Process Choreographer Event Collector and Business Process Choreographer Observer are used together to monitor business processes. The event collector application reads event information from the common event infrastructure bus and stores it in the Business Process Choreographer Observer database, which must be configured. The observer application then uses this event data to produce reports on the business processes and human tasks.

Configure containers…

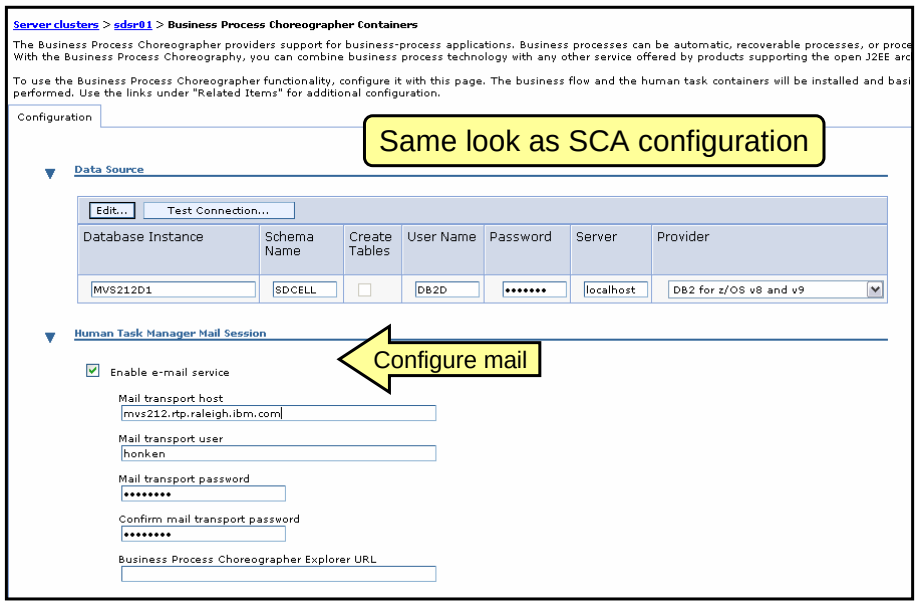

When you configure the Business Process Choreographer container in the administrative console, you are asked for information for the data source that holds the tables needed. The format is identical to the data source configuration for the SCA configuration. You also have the opportunity to configure the e-mail service for human tasks. This can be configured later as well so you can leave the defaults here if you choose.

Configure containers…

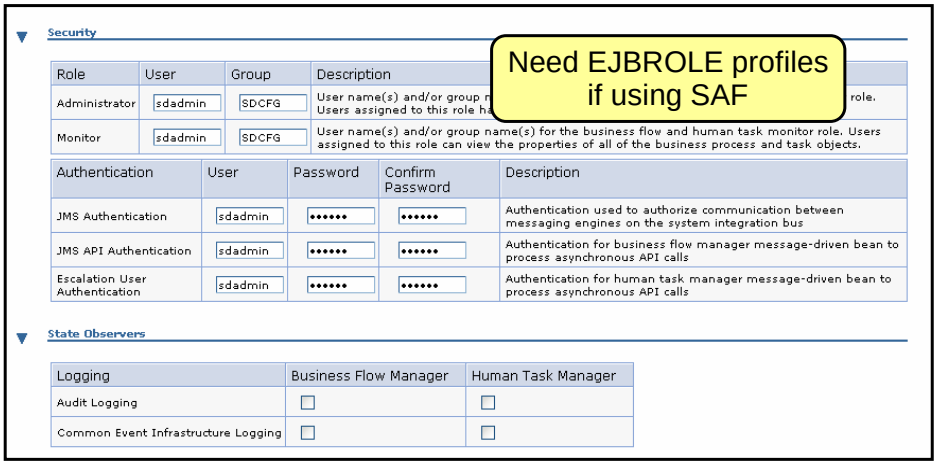

Next you are asked for some security information relating to the business process container. The user ID and groups that you specify need to be given access to the appropriate EJBROLE profiles if you are using SAF. See the **z/OS installation and configuration overview** presentation for more information on creating the EJBROLE profiles. You should give the user ID and group specified for the Administrator role access to the BPESystemAdministrator and TaskSystemAdministrator EJBROLE profiles. The user ID and group specified for the monitor role should be given access to the BPESystemMonitor and TaskSystemMonitor EJBROLE profiles. The JMS Authentication user ID specified is the user that is used to secure the BPC bus and also has the BPC_Auth_Alias, which is used to protect various other resources such as connection factories, created for it. The user ID specified for the JMS API Authentication becomes the user that should be specified for the RunAs role for the JMSAPIUser EJBROLE profile. Finally, the user ID specified for the Escalation User Authentication is the user that should be specified for the RunAs role for the EscalationUser EJBROLE profile.

You can turn on logging for the state observer from here but that should wait until it is used.

Configure containers…

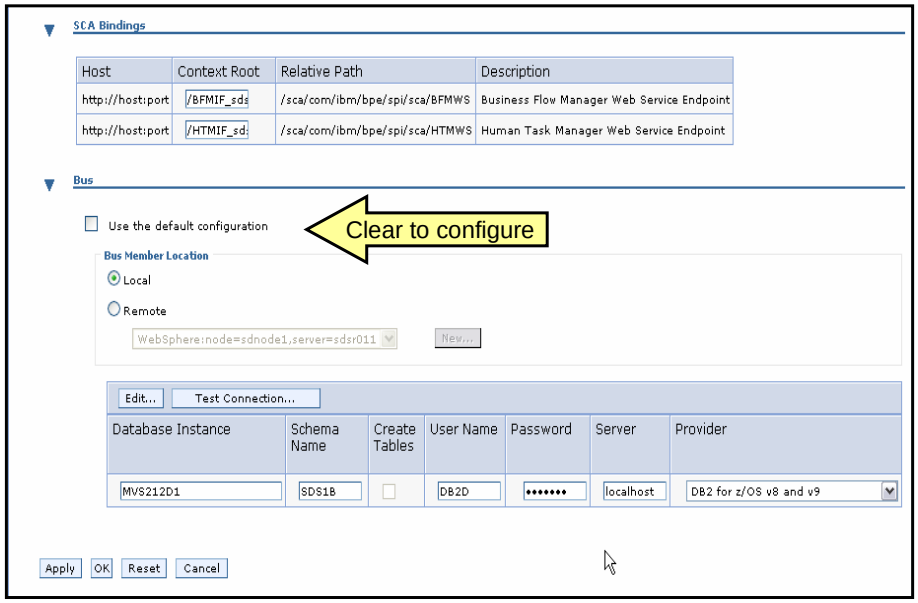

Finally, you are asked for some information about the SCA bindings and the bus that is created for the business process choreographer container. You have the opportunity to specify a different context root for the SCA binding but the default will work fine. In order to configure the bus for the business process choreographer container, you need to first clear the box to indicate to you do NOT want to use the default configuration. You then need to specify the schema name that is used on the tables for the data store on the messaging engine that is created for you. The other fields should be correct.

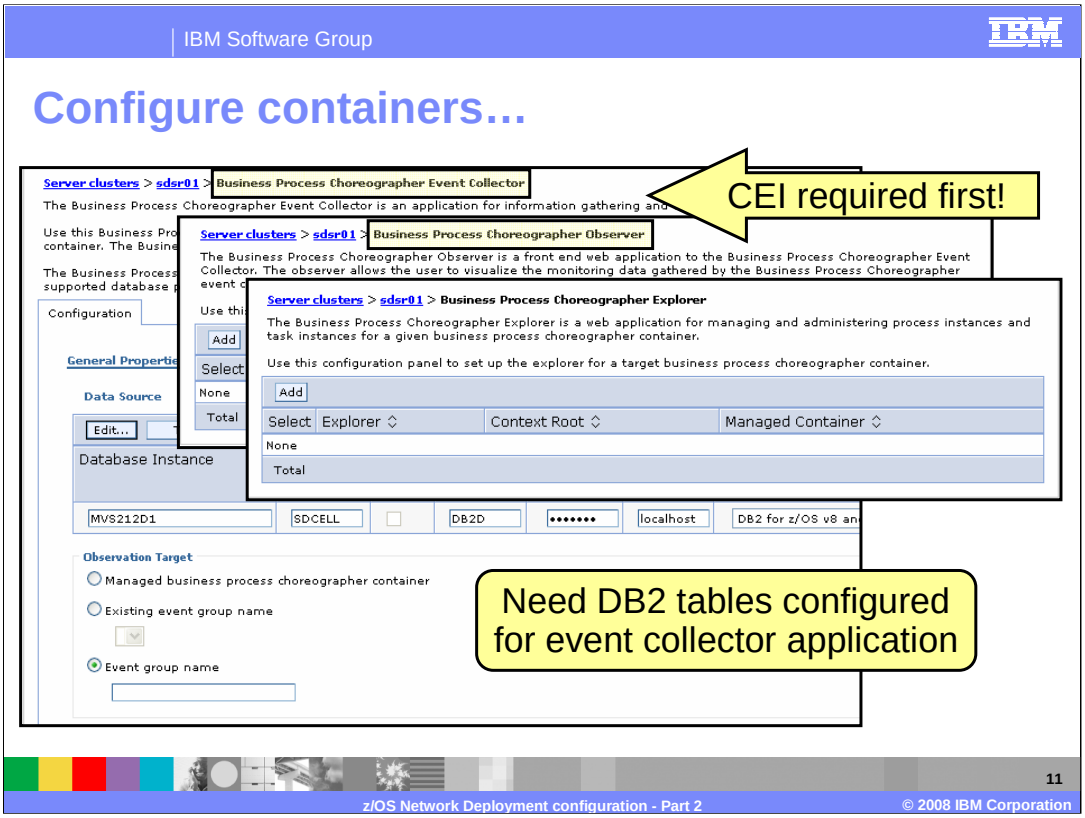

If you are planning to use the Business Process Choreographer Explorer or Collector and Observer, you need to configure those applications as well. Only the Collector has any real information needed where it needs to know about the database that is being used and the observation target. The other two applications have a simple 'add' box to indicate that you want the application deployed to the server or cluster. You are prompted for parameters such as the context root, but you can take the defaults. You must configure CEI before configuring the event collector and observer.

Configure containers using bpeconfig.jacl

```
<WAS_HOME>/ProcessChoreographer/config/bpeconfig.jacl
```

- ../../bin/wsadmin.sh -f bpeconfig.jacl *parameters*
  - Can be run non-interactively if all needed parameters are specified on the command line
  - See 'Using the bpeconfig.jacl script file to configure Business Process Choreographer' article in the Information Center for more detail

While the installation wizard in the administrative console gives you a nice sample BPC configuration, there is also a jacl script available to configure the business process and human task containers. It can be run interactively where you are prompted for all the values it needs or you can specify all values on the command line. To see all the possible parameters and for more information on running it, see the article in the information center.

Configure common event infrastructure

4. Add support for CEI

```
Application servers > <serverName> > Common Event Infrastructure Server
Clusters > <clusterName> > Common Event Infrastructure Server
```

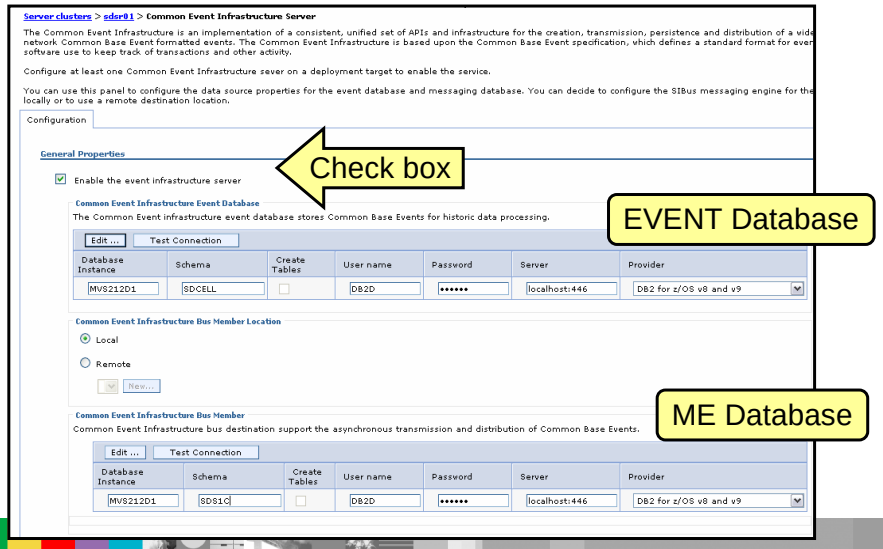

The last bit of function that you might want to add to the WebSphere Process Server or WebSphere Enterprise Service Bus cluster or server is support for the common event infrastructure. This is optional but can also be configured with an installation wizard. In order to configure the common event infrastructure function, you must select the check box indicating you want to enable it on the server. You will then be asked for the same data base information that you have seen before. The first box pertains to the EVENT database that is needed while the second one pertains to the data store that is to be used by the messaging engine that is created for the CEI bus.

Configure common event infrastructure… security

- Security is automatically enabled on the bus that is created
- A valid user or group needs to be added to the bus connector role and to the authentication alias

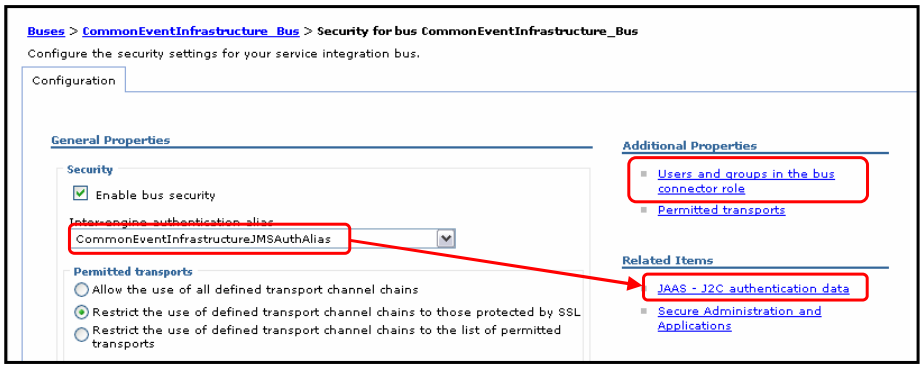

Security is automatically enabled on the bus that is created for the common event infrastructure. Therefore, in order to successfully use the common event infrastructure, you need to add a valid user to the bus connector role and for the specified authentication alias as shown on the slide.

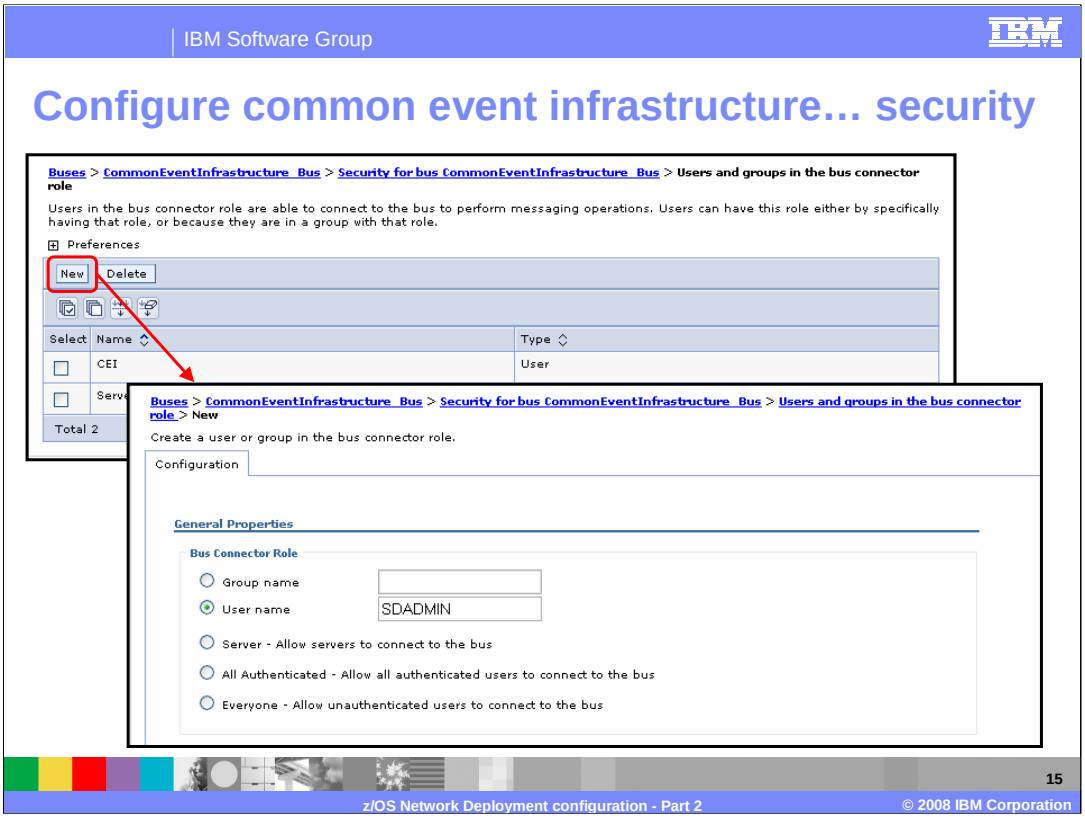

First, you should add a user or group to the bus connector role. This slide shows the addition of a user name. You can also add a group name or even allow everyone to connect to the bus.

Configure common event infrastructure… security

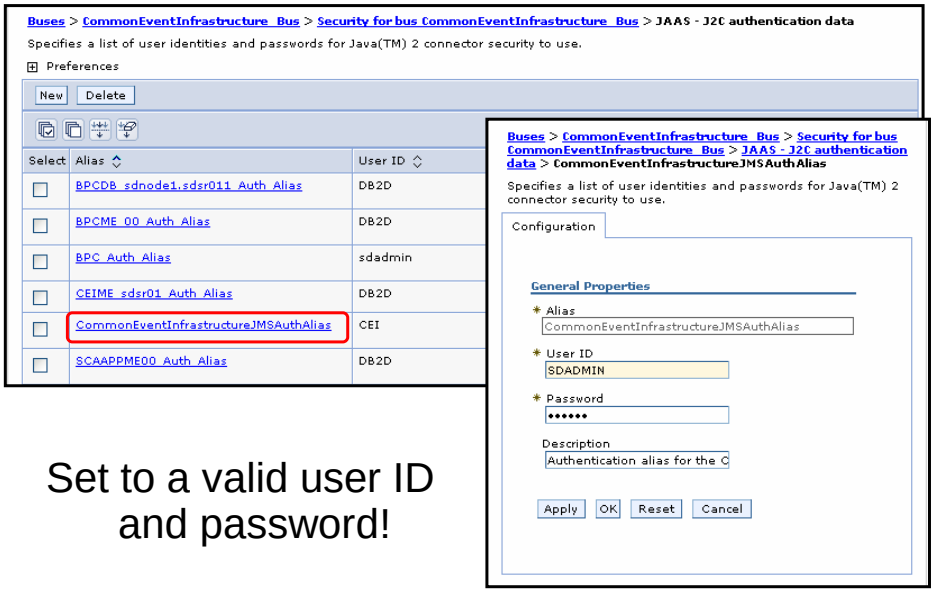

Set to a valid user ID and password!

Next, you should update the authentication alias that was created for you to contain a valid user ID. This example sets it to the administrator's user ID.

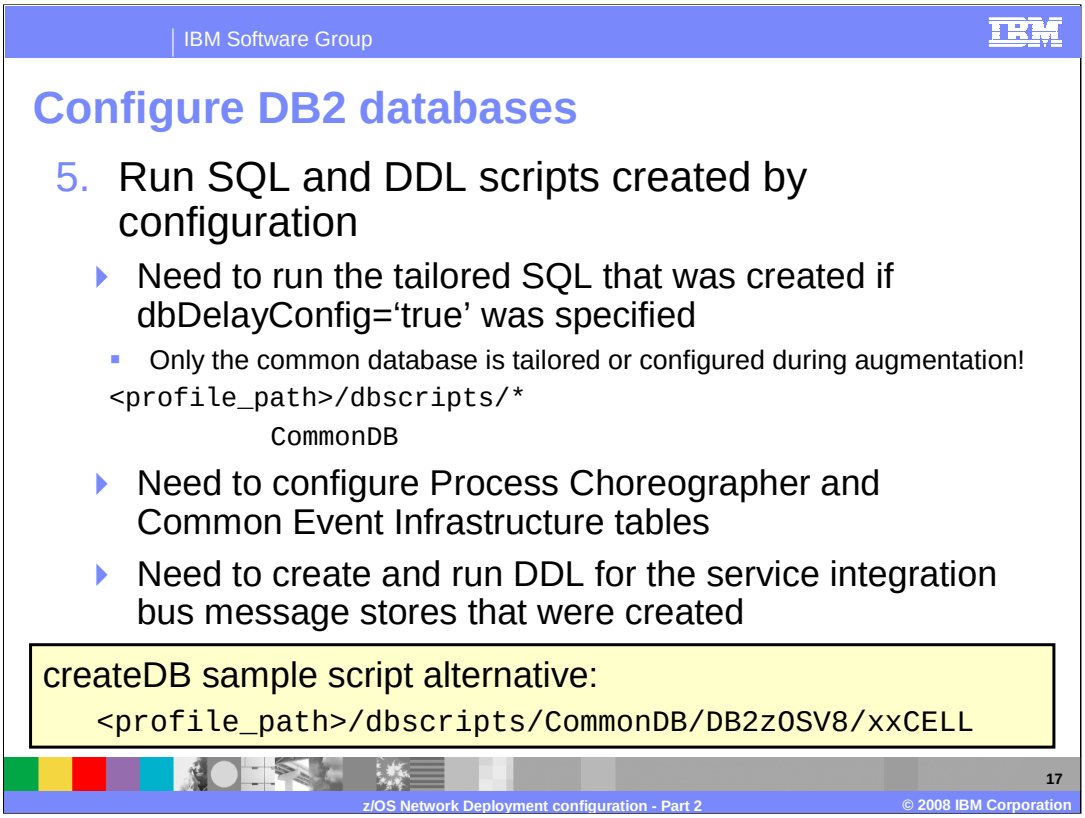

Now that the server or cluster is configured with all the functions you want, you need to configure the databases that the functions will use. Unlike the stand-alone case, only the common database is tailored or configured during augmentation. Partially tailored SQL is created during the business process choreographer configuration using the installation wizard. Fully tailored SQL is created during business process choreographer configuration using the bpeconfig.jacl script. For the common event infrastructure, you need to run a script in order to create the DDL necessary for that function. You also need to create and run the DDL for the service integration bus message stores that were created for you. A very nice alternative to all this is the createDB sample script that is discussed in the **z/OS DB2 configuration** presentation. If you decide not to use the createDB script, the next few slides will show you where to find the DDL or SQL that you need to run to complete the configuration.

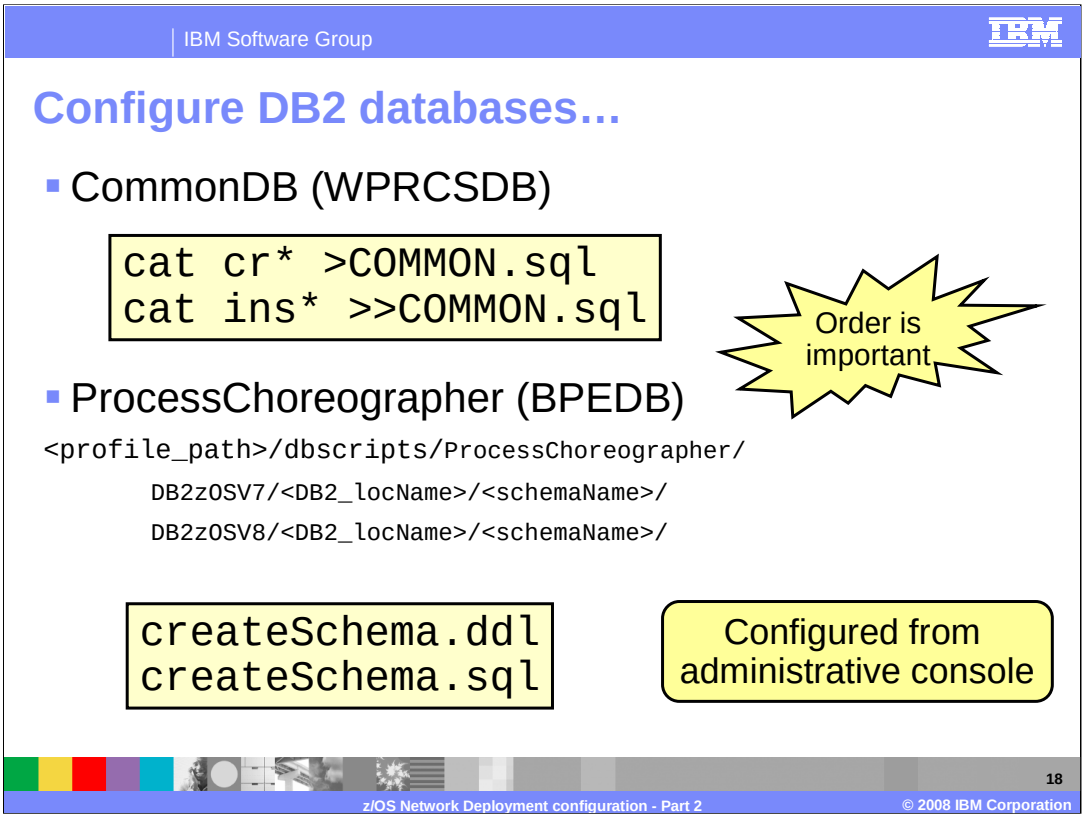

As seen in the z/OS DB2 configuration presentation, each component creates a directory under the dbscripts directory where you will find the SQL or DDL if it exists. It has been tailored (or partially tailored in this case) to your installation but consists of multiple files under each directory. You can combine the CommonDB files into one file, but an alternative is to use the createDB sample script as mentioned on the previous slide. To see the files created for the CommonDB component, see the z/OS DB2 configuration presentation.

When you use the administrative console to configure the business process choreographer containers, not all information to tailor the DDL or SQL is available. Therefore, both DB2 V7 and DB2 V8 SQL and DDL files are created and placed in the directories shown on this slide. The files need some additional tailoring such as substituting the storage group and the database name to match your naming conventions.

Keep in mind that all files ending in SQL are encoded in ASCII and all files ending in DDL are encoded in EBCDIC.

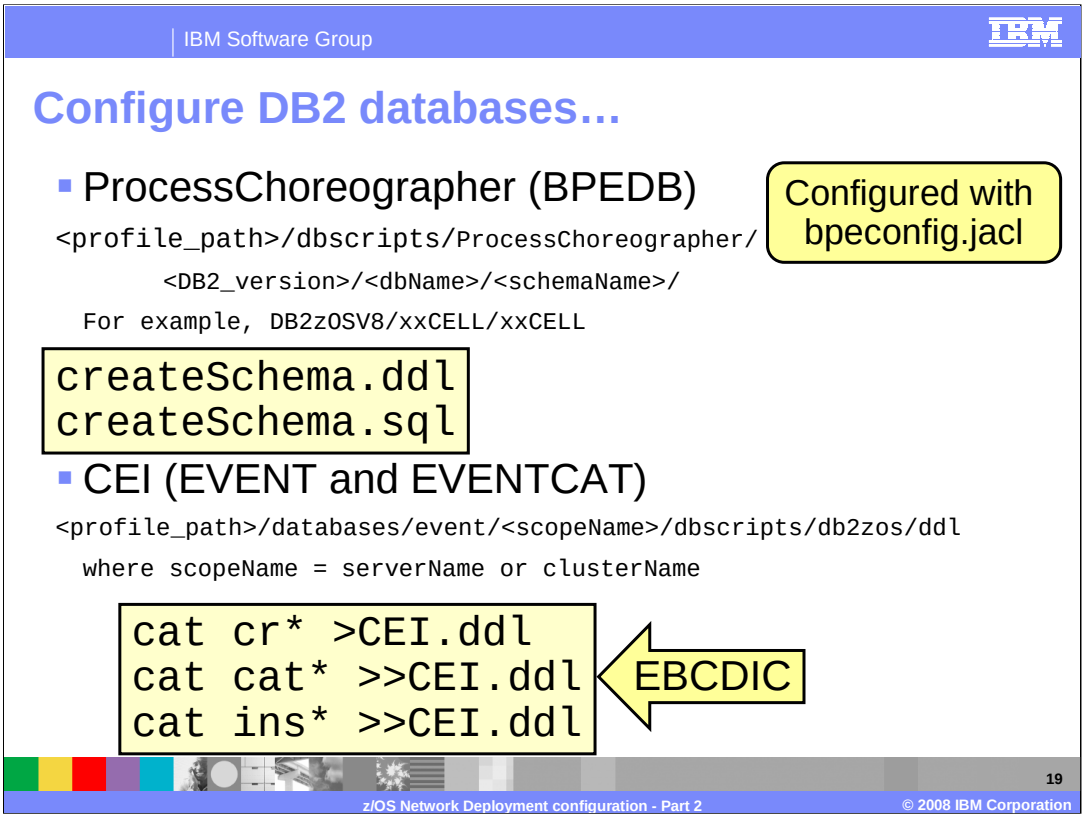

When you use the bpeconfig.jacl to configure the business process choreographer, more information is available so the SQL and DDL created are complete. These are found under the directory shown on the slide and one of the files must be run to configure the database. CEI is the last component that needs database configuration. Tailored DDL is found under the databases/event directory as shown on the slide. Again you can see here how you can combine the files into one. Be sure to examine the output, since there are some things that will need to be changed such as the database name and the storage group name.

Configure DB2 databases…

- Need to create DDL for the messaging engine data stores

```
./sibDDLGenerator.sh -system db2 -version 8.1
  -platform zos -schema xxyyS -user DB2D  -create
  -database xxCELL -storagegroup xxDBSTO -statementend
  ";" > /u/wsuser/wpswork/SIBSCA.ddl
```

- The schema name used here must match the administrative console!

createDB sample script alternative:

```
<profile_path>/dbscripts/CommonDB/DB2zOSV8/xxCELL
```


Regardless of the schema names you decide to use, you need to somehow create the DDL that can be used to configure the DB2 database. Shown here is the sibDDLGenerator command which is one way to generate the DDL for the service integration bus' message stores. One of the parameters specified on the command is 'schema' so you can customize it here to match your conventions . If you chose to use the createDB sample script instead, the SQL is found in the directory shown. The schema names in that case are shown in the **z/OS DB2 configuration** presentation.

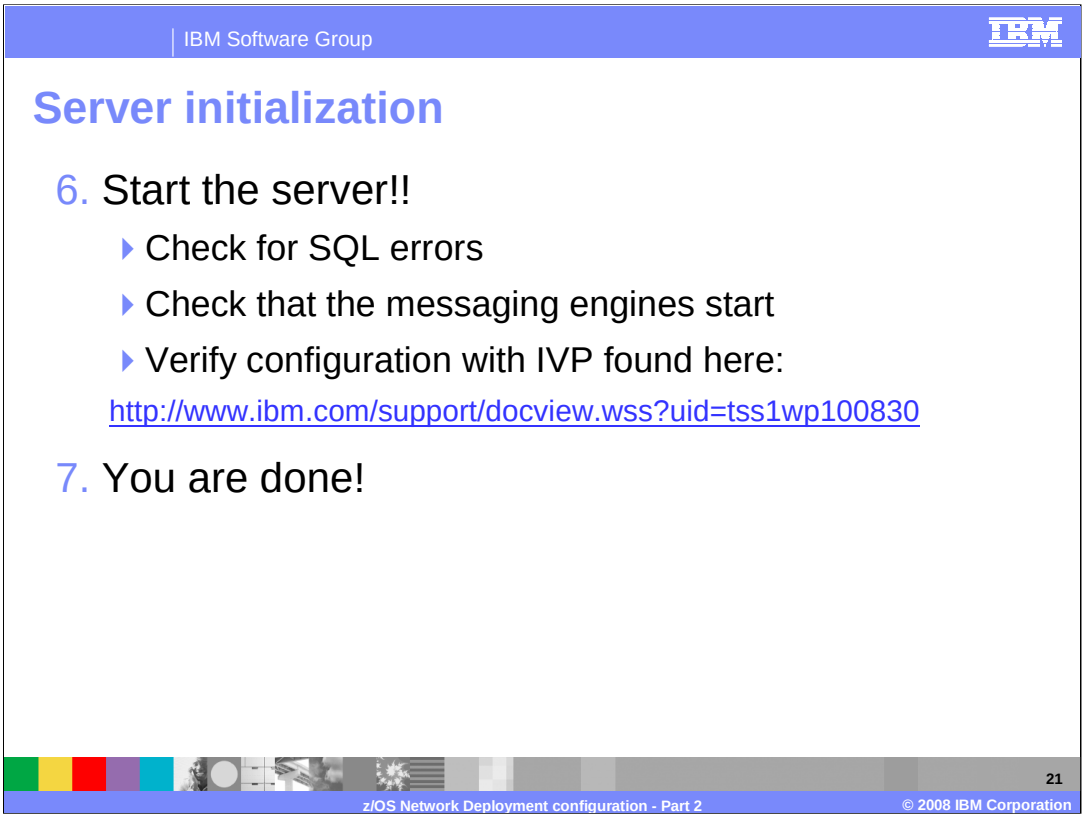

With that, you should have a fully configured WebSphere Process Server or WebSphere Enterprise Service Bus server and you can start it up. Check for any severe errors on startup and check that the messaging engines start. Errors are often seen with the database configuration so check for any SQL errors as well. SQL errors are very common. Once everything looks good, verify the configuration with the white paper found at the Web address shown on the slide. The white paper has you verify each of the various components in a very informative, methodical manner.

Summary

- WebSphere Process Server and WebSphere enterprise service bus configured in a Network Deployment environment has more post-configuration steps
  - Uses DB2 only

In summary, the configuration of a Network Deployment environment with WebSphere Process Server for z/OS or WebSphere Enterprise Service Bus for z/OS has some manual steps. The two parts of this presentation have looked at the steps involved to configure a fully-functional WebSphere Process Server or WebSphere Enterprise Service Bus server or cluster in a Network Deployment 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_WBIV61_zOS_ND_PostConfig.ppt

This module is also available in PDF format at: ../WBIV61_zOS_ND_PostConfig.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:

DB2 IBM 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.