



IBM Software Group

WebSphere® Enterprise Service Bus for z/OS® V6.1 WebSphere Process Server for z/OS V6.1

z/OS simple configuration



@business on demand.

© 2008 IBM Corporation
Updated February 12, 2008

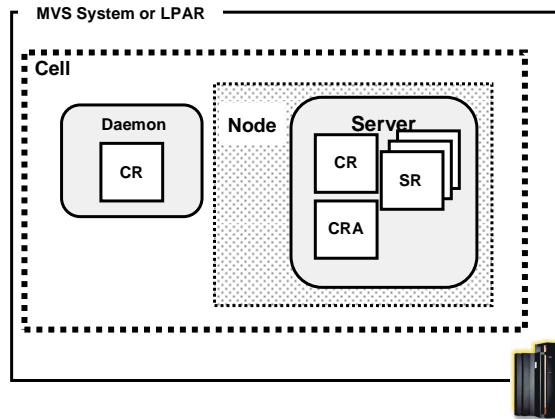
This presentation will look at the simple configuration of a stand-alone application server to enable WebSphere® Process Server for z/OS V6.1 or WebSphere Enterprise Service Bus for z/OS V6.1 function. You should look at the **z/OS installation and configuration overview** and the **z/OS DB2® configuration** presentation as prerequisites to this one.

Goals

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

The goal of this presentation is to take you through the necessary steps 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 simple configuration scenario. This configuration involves a stand-alone WebSphere environment.

Configure a stand-alone application server



- Start with a configured stand-alone application server
- Configure it with WebSphere Process Server or WebSphere Enterprise Service Bus

```

d:\cell
├─ AppServerNode1
├─ CL1BASE1.CL1BASE1.CL1DEM1
├─ CL1BASE1.CL1NODE1.CL1SR01
├─ CL1BASE1.CL1NODE1.CL1SR01.HOME
└─ Daemon
  
```

3

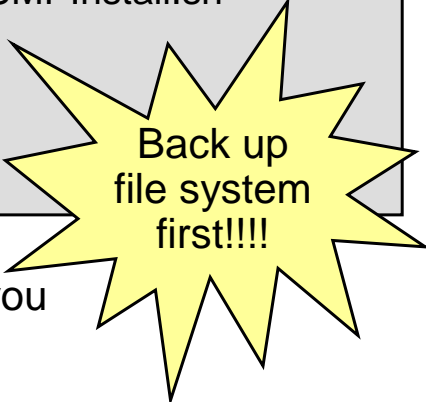
z/OS simple configuration

© 2008 IBM Corporation

This slide shows an already-configured stand-alone application server. This presentation begins with that server as the base and explains how to configure it to include the WebSphere Process Server or WebSphere Enterprise Service Bus function. Note that if you have the WebSphere Process Server for z/OS V6.1 product, you can configure the server with either WebSphere Process Server, which includes the WebSphere Enterprise Service Bus function or WebSphere Enterprise Service Bus only. If you have the WebSphere Enterprise Service Bus for z/OS V6.1 product, your only option is to configure the server with WebSphere Enterprise Service Bus function.

Configure stand-alone – ‘install’

```
<wps_smpe_root>/zos.config/bin/zSMPInstall.sh  
-smproot <wps_smpe_root>  
-runtime <app_server_root>  
-install
```



Back up
file system
first!!!!

- Will perform **-prereqonly** for you
- Run as *'WSADMIN'*
- 'install' creates links to product code, adds plug-ins to the administrative console for new functions

4

z/OS simple configuration

© 2008 IBM Corporation

The first thing you need to do for the configuration is run the zSMPInstall.sh script in order to create symlinks in your WebSphere Application Server configuration to the WebSphere Process Server or WebSphere Enterprise Service Bus product code. This is really a task for the system administrator, since it is somewhat of an extension of the SMP/E install. You should use a WebSphere Administrator user ID to run the script. The zSMPInstall.sh script will also add plug-ins to the administrative console for new functions needed for the WebSphere Process Server or WebSphere Enterprise Service Bus. You need to specify the SMP/E root where you installed the WebSphere Process Server or WebSphere Enterprise Service Bus product code and the app_server_root for the WebSphere Application Server profile you are updating. The script will perform the prerequisite only function for you automatically to ensure that any prerequisites are met. Note that before running the script, you should backup your WebSphere Application Server configuration HFS first. The zSMPInstall script can be run from a telnet session or OMVS. You may find you need to increase some OMVS parameters such as MAXFILEPROC, MAXPROCUSER or MAXCPUPTIME when running the shell script this way. For example, you may need to specify: SETOMVS MAXCPUPTIME=86400.

Configure stand-alone – ‘install’ JCL

```

//INSTONLY JOB (ACCTNO,ROOM), 'HONKEN',CLASS=A,REGION=0M,
// NOTIFY=&SYSUID,TIME=NOLIMIT
//*
/*****
/* zSMPInstall.sh
/*****
//INSTO EXEC PGM=IKJEFT01,REGION=0M
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
BPXBATCH SH +
/etc/WAS60B/usr/lpp/zWPS/V6R1/zos.config+
/bin/zSMPInstall.sh +
-smproot /etc/WAS60B/usr/lpp/zWPS/V6R1 +
-runtime /etc/sscell/ssnode1/AppServerNode1 +
-install +
-trace '*=all=enabled' +
1> /tmp/installonly_84821.out +
2> /tmp/installonly_84821.err
/*
/*****
/* STEP Copy - Copy script output back to joblog
/*****
//MCFG EXEC PGM=IKJEFT01,REGION=0M
//SYSEXC DD DISP=SHR,DSN=WAS60A.SBB0EXEC
//SYSTSIN DD *
BBOHFSWR '/tmp/installonly_84821.err'
BBOHFSWR '/tmp/installonly_84821.out'
//SYSTSPRT DD SYSOUT=*
/*

```

- Can run the job from JCL as well

5

The zSMPInstall.sh script can also be run from JCL. This slide shows an example of running the shell script from JCL.

IBM Software Group IBM

<app_server_root> – after ‘install’

■ New directories with links and new links within existing directories


6

z/OS simple configuration © 2008 IBM Corporation

After running the zSMPIInstall.sh script, the WebSphere Application Server configuration is updated with new directories with links to the WebSphere Process Server or WebSphere Enterprise Service Bus product files, and with new links within existing directories.

Configure stand-alone – planning for DB2

- If using DB2:
 - ▶ Review the **z/OS DB2® configuration** presentation
 - ▶ Determine naming conventions
 - ▶ Talk to your DB2 administrator!
 - Optionally, run createDB sample script so SQL can be run while WebSphere Process Server or WebSphere Enterprise Service Bus is configured



Before going any further in the configuration of WebSphere Process Server or WebSphere Enterprise Service Bus, stop to do some planning. If you are using DB2 in your configuration, you'll need to know the names of the databases that will be used and schema names. Review the **z/OS DB2 configuration presentation** and talk to your DB2 administrator about the DB2 artifacts that are needed. If you decide to use the createDB sample script to create SQL for the databases, you can run that at this point and get your DB2 database configuration underway early.

Configure stand-alone - update response file

- Samples found in `<wps_smpe_root>/zos.config` directory: **standAloneProfile<DB2>.rsp**
 - Copy it to a place where it can be modified
 - Person running the `zWPSConfig.sh` or `zWESBConfig.sh` script needs to have 'R' access
- Global Properties (if DB2):

```
JMSUSER=ibmuser → ceiSampleJmsUser=$JMSUSER
JMSPASS=ibmuser
CONFIGSERVER=xxsr011
DBPRODUCT=DB2UDBOS390_V8_1
DBLOCATION=MVS215D1
...
```

8

z/OS simple configuration

© 2008 IBM Corporation

To prepare for running the configuration script, you will need to update one of the response file samples that are shipped with the product. The samples can be found in the `zos.config` directory in the SMP/E-installed HFS directory path. Since it is a read-only file system, you will need to copy it to a place where it can be modified. Make sure that the userid that you use to run the `zWPSConfig` script or `zWESBConfig` script has at least read-access to the file. There are two samples shipped for the Stand-alone configuration: `standAloneProfile.rsp` and `standAloneProfileDB2.rsp`. Select the correct one based on whether you plan to use DB2 or Derby for your databases.

The first set of parameters you need to set in the response file are ones that will be used later in variable substitutions as shown in the yellow box. This allows you to specify parameters once, making it easier to override on the command line and cutting down on typing errors. The user IDs and passwords that you are asked to supply will be used to create authentication aliases to protect various resources that are created for you. You are able to change their values in the administrative console at a later time if needed. Note that the `DBLOCATION` needs to be the location name, not the subsystem name.

Stand-alone response file – common properties

```
augment
profileName=default
profilePath=/etc/xxcell/AppServerNode1/profiles/default
templatePath=/etc/xxcell/AppServerNode1/profileTemplates/default.wbiserver
cellName=xxbase1
nodeName=xxnode1
serverName=xxsr01
```

Long names!!

- Leave templatePath set to 'default.wbiserver'

```
/WebSphere/V6R1/AppServer/profileTemplates/default.wbicare
/WebSphere/V6R1/AppServer/profileTemplates/default.bpc
/WebSphere/V6R1/AppServer/profileTemplates/default.wbiserver
/WebSphere/V6R1/AppServer/profileTemplates/default.*
```

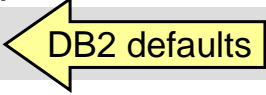
9

The next set of values you are asked to specify in both cases are some common properties such as cellName, nodeName and serverName. It is important to note that if you set up your cell such that you have different names for the long names and short names, you must specify the long name. On z/OS, the profileName will always be 'default' so that should not be changed. You need to change the path parameters to include your configuration HFS path, being careful to leave the non-highlighted part alone. On the templatePath parameter, this value determines which actions are performed during the profile augmentation portion of the install. It should be left 'as-is' with wbiserver specified unless you have a good reason to change it. One reason might be that the job is timing out so it does not finish. For that reason, you might specify one value at a time and run the job multiple times. In that case, you need to start with the wbicare template and move down the list shown. Each value specified will perform the augment actions for that template plus the augment actions for any pre-requisite templates. This means that the wbiserver value will perform the augment actions for everything. The asterisk value will also perform the augment actions for everything but the output log will get overwritten for each of the values and you are left with only the log from the last template augmentation. By leaving the default as wbiserver, the augmentation actions are performed for all three templates and the output is put in one log. The values shown here are valid for the WebSphere Process Server product. The values for the WebSphere Enterprise Service Bus product are different and can be found in the response file.

Stand-alone response file – database

Database configuration parameters

```
dbCreateNew=false
dbDelayConfig=true
```



DB2 defaults

- dbCreateNew – always false for DB2
 - Determines whether databases will be created during augmentation
- dbDelayConfig – defaults to true
 - Determines whether databases are configured during augmentation
 - If set to 'false', the tailored SQL will be run during augmentation
 - Must have authority to run the SQL
 - Databases must already exist!

10

z/OS simple configuration

© 2008 IBM Corporation

As you saw earlier, many databases and tables are required to run WebSphere Process Server or WebSphere Enterprise Service Bus. The two parameters shown here, dbCreateNew and dbDelayConfig, determine what configuration is done during the augmentation step. The dbCreateNew parameter determines whether the specified databases are created during augmentation. This is set to 'true', by default, when configuring with Derby but **MUST** be set to 'false' when configuring with DB2. If you are configuring your system to use DB2, the databases **MUST** be created manually. The augmentation will never create them for you. If you are configuring your system to use Derby, the default of 'true' should not be changed.

The second parameter, dbDelayConfig, determines whether the tailored SQL is run during augmentation or whether it will be run manually after augmentation. Again, the default value for Derby is 'false' but the default value for DB2 is 'true'. Setting this to 'false' allows the configuration to be fully automated. If you are configuring your system to use Derby, you again should not change the default value of 'false' as you want to let the augmentation configure the databases for you automatically. If you are configuring your system to use DB2, you have the option of running the SQL during augmentation if you set this parameter to 'false'. Normally, this parameter is set to 'true' and the SQL will need to be run manually. Remember, you **MUST** have the required database (or databases) created manually when using DB2. If the databases are already created at the time of augmentation and you have the authority to create tables and indexes for them, then setting the dbDelayConfig value to 'false' will automatically run the SQL needed to configure them.

Stand-alone response file – WebSphere Business Integration core properties

```

CEI {
  ceiSampleJmsUser=$JMSUSER
  ceiSampleJmsPwd=$JMSPASS
  ceiSampleServerName=$CONFIGSERVER
  ceiDbName=$SQLDB
  ceiDbStorageGroup=$SQLSTO
  ceiDbSqlId=$SQLID
  ceiDbProduct=CEI_DB_DB2_ZOS
  ceiDbConnectionLocation=$DBLOCATION
}
SCA {
  configureScaSecurity=true
  scaSecurityUserId=$JMSUSER
  scaSecurityPassword=$JMSPASS
}
ESB {
  esbDbName=$SQLDB
  esbDbStorageGroup=$SQLSTO
  esbDbSqlId=$SQLID
  esbDbProduct=$DBPRODUCT
}

```

Need to know
DBNames and
STOGroups as these
will be used to tailor
SQL/DDL

11

z/OS simple configuration

© 2008 IBM Corporation

The next set of parameters deal with the configuration of the WebSphere Business Integration Core functions. At this point, you need to know the names you plan to use for the DB2 databases, schema names and storage groups as seen here in a portion of the DB2 version of the response file. These names are used to tailor SQL to be run either during the augmentation or manually after the augmentation. You can use the default database names for the cloudscape version of the response file. Not all CEI parameters are shown here but you should get a good idea of the types of things being asked. For the service component architecture function it will create an authentication alias to protect the system integration buses that are created if you set the `configureSCASecurity` parameter to 'true'. Finally, you are asked some information for the enterprise service bus setup having to do with the database that is required. Again, not all ESB parameters are shown. The names specified here must match what you decided on for your naming conventions.

Stand-alone response file – Business process choreographer properties

Business process choreographer configuration (optional)

```
configureBPC=true  
bpcDbName=$SQLDB  
bpcDbStorageGroup=$SQLSTO  
enableAdminSecurity=true  
adminUserName=$JMSUSER  
adminPassword=$JMSPASS
```

Other values
not shown

- If `configureBPC=true` specified, will create a sample business process choreographer configuration
- Administrative security **MUST** be turned on in order to configure BPC
 - `enableAdminSecurity=true` required for BPC configuration
 - Otherwise, will see error:
 - "Returning error message:cannotBeSetWithoutSecurity"

12

z/OS simple configuration

© 2008 IBM Corporation

This slide shows a subset of the business flow manager configuration parameters. The business flow manager is not part of the WebSphere Enterprise Service Bus configuration. It is available only when configuring WebSphere Process Server. The business flow manager includes both the business process choreographer and the human task manager. In order for these functions to be configured at this time, you must specify `configureBPC=true` as shown here on the slide. If `configureBPC=false` is specified, a sample configuration is not created at this time. It is possible to create a sample configuration later using `.jacl` scripts or the installation wizards in the administrative console. You can find a sample of running the `.jacl` script or using the administrative console install wizard in the `NetworkDeploymentConfiguration` presentation.

Here, 'optional' refers to whether you actually configure the Business Flow Manager at this time. The components needed will still be installed when you run the configuration but you will have to configure them at a later time. It's a good idea to configure the sample at this time, however. Later on, when you get an idea of how you will be using it, you can use the `.jacl` scripts to update the configuration more to your liking.

Stand-alone response file – Business process choreographer properties...


- Includes human task manager and business flow manager
- Applications deployed and buses configured:
 - ▶ Business process container
 - BPEContainer_<node>_<server> application
 - BPC.<cell>.Bus Service Integration Bus
 - ▶ Human task container
 - TaskContainer_<node>_<server> application
 - ▶ Business process choreographer explorer
 - BPCExplorer_<node>_<server> application
 - ▶ Business process choreographer observer
 - BPCObserver_<node>_<server> application
 - ▶ Business process choreographer event collector
 - BPCECollector_<node>_<server> application

The sample configuration that is created includes the Business Process Container, the Human Task Container, the Business Process Choreographer Explorer, the Business Process Choreographer Observer and the Business Process Choreographer Event Collector. This includes some applications and a Service Integration Bus as shown on the slide.

Stand-alone response file – common database

- WPRCSDB database configuration

```
dbName=$SQLDB
dbStorageGroup=$SQLSTO
dbType=$DBPRODUCT
dbConnectionLocation=$DBLOCATION
dbJDBCProperties=$DBPROPERTIES
dbJDBCClasspath=$DBJDBCCLASSPATH
dbUserId=$DBUSER
dbPassword=$DBPASS
dbHostName=$DBHOSTNAME
dbServerPort=$DBSERVERPORT
dbSchemaName=$SQLID
```



- dbJDBCProperties file must point to DB2JccConfiguration.properties if specified

This slide shows some values you must specify for the common database configuration in order to create JDBC resources, authentication aliases and the SQL needed to configure the tables that are used. This database is needed for both the WebSphere Process Server and the WebSphere Enterprise Service Bus. Shown here are values from the DB2 response file; the Derby version has just a subset of these parameters. Again, it's important that the database name and storage group parameters match what you specified when you created the database. The ServerPort parameter is used for a Type 4 definition only. If you need to specify a JDBC properties file (dbJDBCProperties), it must be called DB2JccConfiguration.properties.

Configure stand-alone – ‘augment’

```
export LIBPATH=/db2810/jcc/lib:$LIBPATH
<app_server_root>/bin/zWPSConfig.sh
    -response standAloneProfile<DB2>.rsp
    -augment
```

- ‘augment’ will create resources and install applications needed to run the WebSphere Process Server or WebSphere Enterprise Service Bus
- Needs to be run from the <app_server_root>/bin directory as ‘WSADMIN’
- Need to export LIBPATH first if using DB2 and dbDelayConfig=false

15

z/OS simple configuration

© 2008 IBM Corporation

Now that you have completed the update of the response file, you are ready to run the augment job. If using DB2 and you set dbDelayConfig to ‘false’ in order to configure the databases automatically, you will need to export your LIBPATH first as shown on the slide. The zWPSConfig script, found in the bin directory of your app_server_root is used for the augment. You should again run the script from a WebSphere administrator user ID. The only parameter you need to specify other than ‘augment’, which takes no value, is the ‘response’ parameter to indicate where the updated copy of the sample response file can be found.

Configure stand-alone – ‘augment’

- Keep in mind that configuring WebSphere Process Server will also configure ESB
- To configure an ESB-only server use:
 - zWESBConfig.sh

```
<app_server_root>/bin/zWESBConfig.sh  
-response standAloneProfile<DB2>.rsp  
-augment
```

Recall that when you configure WebSphere Process Server, you automatically configure the WebSphere Enterprise Service Bus since that is included as part of the WebSphere Process Server Product. It is also possible to configure only the WebSphere Enterprise Service Bus. The command to do that, zWESBConfig, is shown here. This would be the only available option if you have the WebSphere Enterprise Service Bus for z/OS V6.1 product.

Configure stand-alone – ‘augment’ JCL

```

//AUGMTDB2 JOB (ACCTNO,ROOM), 'HONKEN', CLASS=A, REGION=0M,
// NOTIFY=HONKEN, TIME=NOLIMIT
// *
// *
//*****
// * zWPSConfig.sh
//*****
//AUGMT EXEC PGM=IKJEFT01, REGION=0M
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
BPXBATCH SH +
cd /etc/sscell/ssnode1+
/AppServerNode1+
/bin; +
export LIBPATH=/usr/lpp/db2810/jcc/lib:$LIBPATH; +
./zWPSConfig.sh +
-response /u/honken/wpswork/sscell/standAloneProfileDB2.rsp +
-augment +
-trace '*=all=enabled' +
1> /tmp/zWPSConfig_40135.out +
2> /tmp/zWPSConfig_40135.err
// *
//*****
// * STEP Copy - Copy script output back to joblog
//*****
//MCFG EXEC PGM=IKJEFT01, REGION=0M
//SYSEXEC DD DISP=SHR, DSN=WAS60A.SBB0EXEC
//SYSTSIN DD *
BB0HFSWR '/tmp/zWPSConfig_40135.out'
BB0HFSWR '/tmp/zWPSConfig_40135.err'
//SYSTSPRT DD SYSOUT=*
//

```

- Note the ‘export’ in the JCL

Again, here is an example of running the augment function using JCL. Do not forget to export the native Java common client LIBPATH if dbDelayConfig is set to ‘false’ and you are using DB2.

Configure stand-alone – post configuration tasks

- Stand-alone application server cell (with Derby)
 - ▶ No post-configuration tasks!!
 - ▶ You are done!



- Stand-alone application server cell (with DB2)
 - ▶ Create DB2 Databases and Tables

Once your augment job has run successfully, you have to look at what post configuration tasks are needed. If you configured using Derby, you are done. Everything, including your databases, has been automatically configured and you are ready to start using the WebSphere Process Server or WebSphere Enterprise Service Bus function. If you configured using DB2, you most likely have some SQL to run to in order to configure your database or databases. If you were able to specify, `dbDelayConfig=false`, you are also done at this point.

Configure stand-alone – configure DB2 databases

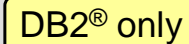
- Need to run the tailored SQL/DDL that was created if dbDelayConfig='true' was specified

```
<profile_path>/dbscripts/*
```

```
CEI_xxCELLDB
```

```
CommonDB
```

```
ProcessChoreographer
```



- Need to create and run DDL for the service integration bus message stores that were created during augmentation

➤ createDB sample script alternative:

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

If you are using DB2 in your configuration, you need to now run the tailored SQL or DDL that was created during augmentation. Each component creates a directory under the dbscripts directory where you will find the SQL or DDL needed. It has been tailored to your installation but as seen in the [z/OS DB2 configuration](#) presentation; this entails multiple files under each directory. The next slide will show you how you can easily combine the files into one but an alternative is to use the createDB sample script as mentioned in the [z/OS DB2 configuration](#) presentation. You also need to generate DDL to configure the message stores needed for the service integration bus messaging engines that were created. We'll look at these in more detail on a later slide. Keep in mind that all files ending in SQL are encoded in ASCII and all files ending in DDL are encoded in EBCDIC. All files other than the CEI ones are therefore in ASCII.

Configure stand-alone – configure DB2 databases

- CEI (EVENT and EVENTCAT)


```
cat cr* >CEI.ddl  
cat cat* >>CEI.ddl  
cat ins* >>CEI.ddl
```



EBCDIC

- CommonDB (WPRCSDB)

```
cat cr* >COMMON.sql  
cat ins* >>COMMON.sql
```



Order is
important

- ProcessChoreographer (BPEDB)

```
cat cr* >PC.sql
```

To see the actual files created for each component, you can look at the **z/OS DB2® configuration** presentation. What you see here is a sequence of commands in order to put the multiple files into one if you choose not to use the createDB sample script to generate the SQL. Again, note that only the CEI files are encoded in EBCDIC.

Configure stand-alone – configure DB2 databases

▪ Configure service integration buses' messaging engine data stores

- `xxnodey.xxsr01y-CommonEventInfrastructure_Bus`

- Default schema name : xxxCM00

- `xxnodey.xxsr01y-SCA.APPLICATION.xxbasey.Bus`

- Default schema name : xxxSA00

- `xxnodey.xxsr01y-SCA.SYSTEM.xxbasey.Bus`

- Default schema name : xxxSS00

where xxx = first 3 characters of dbSchemaName

+CM00 (SIBCEI)

+SA00 (SIBAPP)

+SS00 (SIBSCA)

Moving on to the message stores needed for the service integration bus messaging engines, you'll see three service integration buses were created using DB2 as a message store. They are SIBCEI, SIBAPP and SIBSCA. The default schema is shown here. It is formed by taking the first 3 characters of the dbSchemaName variable in the response file and concatenating it with the values shown on the slide. You may want to change these to conform to your standards. You'll see how to do that on a later slide.

Configure stand-alone – 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 create 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 choose to use the createDB sample script instead, the SQL will be found in the directory shown. The schema names in that case are shown in the **z/OS DB2® configuration** presentation.

Configure stand-alone – configure DB2 databases

- Change schema name if required!
 - ▶ Application servers > *<serverName>* > Messaging engines > *<MEName>* > Message store

The screenshot shows the configuration page for a Data store in the IBM WebSphere Configuration Console. The breadcrumb navigation is: Application servers > sssr01a > Messaging engines > ssnodea.sssr01a-CommonEventInfrastructure_Bus > Data store. The page title is 'Data store' and the description is 'The persistent store for messages and other state managed by the messaging engine.' The 'Configuration' tab is active. Under 'General Properties', the 'Schema name' field is highlighted with a red box and contains the value 'bxyyC'. Other fields include 'UUID' (0CE5869563F1A508), 'Data source JNDI name' (jdbc/com.ibm.ws.sib/ssnodea.sssr01a-CommonEventInfrastructure_Bus), and 'Authentication alias' (CEIME_ssnodea.sssr01a_Auth_Alias). A 'Create tables' checkbox is present and unchecked. On the right, under 'Related Items', there is a link for 'JAAS - J2C authentication data'.

If you decide not use the default schema names chosen for you, you must go into each messaging engine definition and change the schema name to match what you configured your database with. It is important to have the schema name here match. It is used to qualify your table names. The path to get here is shown on the slide.

Configure stand-alone – configure DB2 databases

- BPC – configured using File store

Application servers > sssr011 > Messaging engines > ssnode1.sssr011-BPC.sbase1.Bus > File store

The persistent store for messages and other state managed by the messaging engine.

Configuration

General Properties

UUID
E0D02DB7B9F6A872

Log

* Log size
100 MB

Log directory path
\${USER_INSTALL_ROOT}/filestores/com.ibm.ws.sib/ssnode1.sssr011-BPC.sbase1.Bus-244379AB91FEE07D/log

Permanent store

* Minimum permanent store size
200 MB

Unlimited permanent store size

* Maximum permanent store size
500 MB

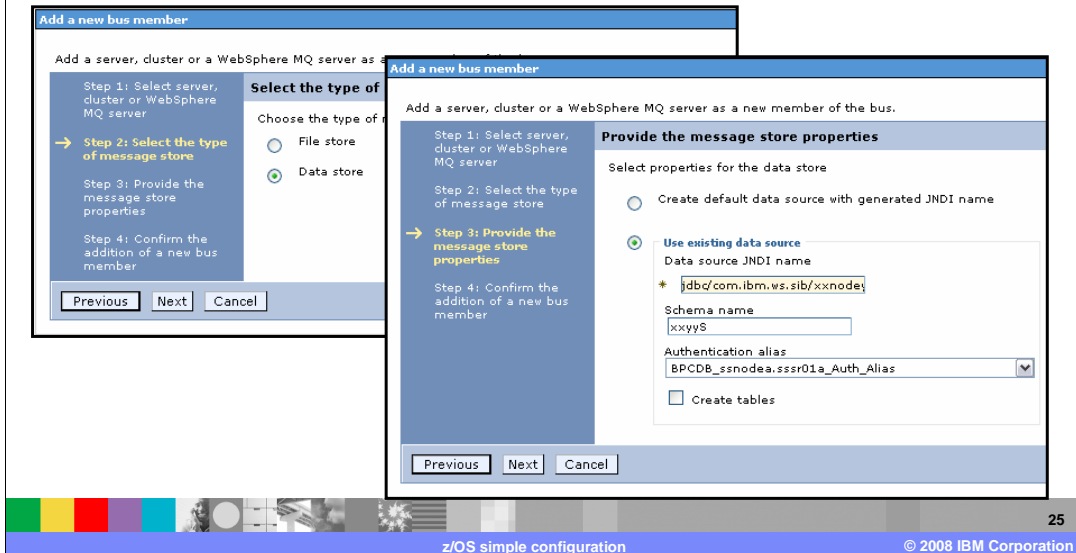
Permanent store directory path
\${USER_INSTALL_ROOT}/filestores/com.ibm.ws.sib/ssnode1.sssr011-BPC.sbase1.Bus-244379AB91FEE07D/store

May want to
change to
DB2

You'll notice that one service integration bus was missing in the list of messaging engines that required DB2 data stores. In the stand-alone configuration, the business process choreographer bus' messaging engine is configured with a File store as its message store. That is shown here on the slide. You may want to change that to DB2. To do that, however, you need to delete the current bus and re-create it.

Configure stand-alone – configure DB2 databases

- Data store chosen when adding the server as a bus member



When recreating the bus, you'll have to add the server as a bus member and it is then that you get the choice of selecting a Data store instead of a File store.

Summary

- Simple configuration of WebSphere Process Server for z/OS V6.1 and WebSphere Enterprise Service Bus for z/OS V6.1
 - ▶ Fully automated with Derby
 - ▶ Can be highly automated with DB2

The simple configuration of WebSphere Process Server for z/OS V6.1 and WebSphere Enterprise Service Bus for z/OS V6.1 is done on a stand-alone application server. If using Derby, the configuration is fully automated. If using DB2, it can be highly automated. There are some manual tasks that are necessary in the DB2 case. This presentation looked in detail at the steps necessary to configure a simple configuration on z/OS.

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_zOSWPSSimpleConfig.ppt

This module is also available in PDF format at: ..\WBIV61_zOSWPSSimpleConfig.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 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 NON-INFRINGEMENT. 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.

