



IBM Software Group

WebSphere® Business Monitor V6.0.2

Administration: Deployments



@business on demand.

© 2007 IBM Corporation
Updated April 18, 2007

This presentation covers using the Process Server administrative console to deploy models, un-deploy models and use version control for WebSphere Business Monitor Version 6.0.2.

Agenda

- Change management
- Versions
- Deploying a model
- Un-deploying a model
- Troubleshooting

This is the agenda for this presentation.

You will see how to deploy, redeploy and un-deploy monitor models. You will also review version control.

Finally, you will review various logs to use when you encounter problems.

Redeploying a model – What you can change

- Change metric string length
- Adding or deleting metrics, stopwatches, counters
- Adding or deleting monitoring contexts
- Updating correlation information for a monitoring context
- Adding trigger or updating evaluation times
- Adding or deleting a dimension (all metrics in deleted dimension must also be deleted)
- Adding a new measure in data mart model

3

This slide lists various things that you can change in a model when redeploying a new version of the model. You can update, add and delete metrics, stopwatches, counters, monitoring contexts, triggers, dimensions and measures in the Data mart model. But be careful when deleting a dimension, because you must also delete all the metrics associated with the dimension.

Redeploying a model – What you cannot change

- Changing a Stopwatch from accumulating to non-accumulating, or vice versa
- Changing which dimension a metric is assigned to
- Changing the set of key metrics for a dimension, this can happen in various ways:
 - ▶ Creating a new metric, assigning it to an existing dimension and marking it as part of the key
 - ▶ Deleting a metric which is part of the key of an existing dimension
 - ▶ Toggling the 'Is Key' check box for a metric in a dimension
- Changing a metric from being a fact attribute to dimension attribute or vice versa
- Changing data type of a metric - you should change ID
- Changing model ID – you should change MC and KC ID's also

4

This slide lists some items that, if changed, will cause errors during deployment. You should avoid changing a stopwatch accumulation type. You should watch out for moving metrics from one dimension to another. You should not change the key metrics for a dimension, nor should you move a metric from fact to dimension or vice versa. If you change the data type of a metric, you should also change the identifier on the metric so you are essentially creating a new metric. Also, you may want to make a copy of a model to use as a starter for a new model, so you can rename the model ID on the copied model, but note that the context IDs must also be unique, so you should plan to rename the monitoring context and KPI contexts as well.

Redeploying a model – DMS artifacts

- If you run life cycle step 1 and DS_DMS_setup.zip is not produced
 - ▶ DMS detected that you had not changed your model in a way that impacted DMS
- If you run life cycle step 1 and DS_DMS_setup.zip is produced but no capture/apply scripts are created
 - ▶ You added some metrics, but did not add any new MCs, so there were no new capture/apply servers needed, so you just keep running the previous version scripts
- You still need to run cube life cycle steps, because a new cube is created for every MC/model/version



When you deploy a new version of a model, you run step 1 of the life cycle wizard to generate the data services scripts. If you find that the Data Movement Services zip is not produced, this is not an error. In fact, DMS did not create the zip because the model was not changed in a way that required changes to DMS.

If the Data Movement Services zip *is* created but does not contain any capture and apply scripts, this is not an error. In this case you did not add any new monitoring contexts so the scripts are not needed, and you will just keep running scripts for the previous version.

Note that you still need to continue with all the life cycle steps - including the cube steps - because a new cube is created for each version.

Model versions

- Specified as a model identifier and unique timestamp in MME on the Monitor Details Model tab for the model artifact
 - ▶ Automatically set to current timestamp when you create a new model
 - ▶ Need to specify unique EAR application name during install of later version
- Each version is separate EAR file containing mm and diagrams
- You can only import versions into Monitor with the timestamp version increasing, else you get a log error – you must update the timestamp manually
- Timestamps on queues (and other internal artifacts) do not match model version timestamp, but is milliseconds since 1970
 - ▶ Separate queues for each version
- Timestamp in schemagen folder is not related to model timestamp – it is current timestamp
- All versions are in the database but only the latest version is capable of creating new monitoring contexts in the monitor server. Monitoring context instances that were created under older versions, but have not terminated, may continue to receive events.

A new model version is designated by a unique timestamp specified in the MME. It is automatically set to the current timestamp for new models, but for new versions you will need to increase the timestamp. Each model version is a separate EAR file that is deployed using the Monitor model life cycle steps. Note that the timestamp on many internal artifacts including JMS queues have a unique timestamp that is related to the version timestamp, but it shows the number of milliseconds since 1970. Also, only the latest version is capable of creating new monitoring context instances. The older versions will process events related to existing monitoring context instances.

Model versions

- Each monitoring context should have a single unique common base event definition that creates monitoring contexts. All other event definitions for that monitoring context should not create monitoring contexts and are viewed as "updating" the monitor context".
- The single unique "create" event definition should not be used subsequently in the monitoring context to update metrics.



To ensure that your model works when deploying multiple versions of the model, you should define a unique event with 'create monitoring context' semantics. This event should not be used elsewhere in the model to update metrics. This will prevent the newest version from creating MC instances for existing instances that are being managed by the older versions.

Model versions

- No 'valid from' – version with greatest timestamp is the active version (even if it is a future timestamp)
 - ▶ However, the 'valid from' on the BPEL process is used in the model correlation criteria, so a model will start processing events only after the validFrom date is satisfied
- CEI distribution mode
 - ▶ Active – latest version
 - ▶ Active (no new MC instances) – older versions are set to this automatically by life cycle
 - ▶ Inactive – after model install, but before life cycle steps
 - ▶ Inactive (event queue recoverable)
- KPIs are calculated based on fact table, which can contain data for multiple versions of the model

8

Unlike Monitor 6.0.1, there is not a capability to use 'valid from' in the monitor versions. The version with the greatest timestamp is active even if the timestamp is in the future.

There are four CEI distribution modes. Active mode is automatically set for the latest version. Mode 'Active (no new MC instances)' is automatically set for older versions of a model. A model in Inactive mode has been installed, but the life cycle steps have not completed. Mode 'Inactive (event queue recoverable)' is used for recovery purposes when a bad version of a model has been deployed.

Note that KPIs are based on the underlying fact table which is not cognizant of the versions, so KPI values will automatically span multiple versions of a model.

Deploy a new version of a new model

- Create EAR in MME
- Install the Monitor model using administrative console
- Run setup wizard
 - ▶ Run data services generation
 - ▶ Execute create schema scripts
 - ▶ Execute create DMS scripts
 - ▶ Import DB2® cube definition
 - ▶ Create Alphablox® cube
 - ▶ Configure CEI distribution
 - ▶ Reboot CEI server
- Start replication manager daemons
- Create dashboards

9

Here are the steps used to deploy a new version of a new model. First you create the EAR in the MME, then install it using the administrative console of the Monitor Server. Then you run the seven steps of the model life cycle. You start the replication daemons, then you build the dashboards in Portal Server.

Setup wizard notes

- **Configure CEI distribution**
 - ▶ Initiates potentially long-running background process
- **Confirm CEI server reboot**
 - ▶ If confirm button is enabled – time to reboot CEI server
 - ▶ Initiates potentially long-running background process
 - ▶ May involve draining event queues of previous versions (3 minutes each)
 - ▶ Be patient

Note that the life cycle steps to configure CEI distribution and to reboot the server both initiate potentially long running processes, so you may have to wait for them to complete. The button to confirm reboot of the server will not be available so then you will know that processes are not complete.

Deploy a new version of a copied model

- To avoid name collisions you will need to rename:
 - ▶ Monitor model ID
 - ▶ Monitoring context ID
 - ▶ KPI context ID

Note that there are three unique identifiers that you must change when attempting to make a copy of a model and deploy it. In the copied model you will need to rename the model ID, rename all monitoring context IDs and rename all KPI context IDs.

Deploy a new version of an existing model

- Same process as for new model
 - ▶ Update model in MME
 - ▶ Update version timestamp in MME
 - ▶ Create EAR in MME
 - ▶ Run setup wizard
 - ▶ Start RM daemons
 - You must start all daemons for all versions of the model
 - To generate a script to do this for you, see
C:\IBM\WebSphere\Monitor\tools\DataMovementServices\StartStopScriptGen.zip
 - ▶ Update dashboards as needed
- Old versions are 'Active (no new MC instances)'

To deploy a new version of an existing model, the process is essentially the same as for a new model. You will update the model using the Monitor Model Editor, increase the timestamp on the model in the MME and export the EAR. Then you run the setup wizard in the administrative console of Monitor Server, start the replication daemons, and update the dashboards.

Note that you need to start all daemons for all versions of the model. There is a supplied script listed here that can automatically generate a single script to start all daemons for all versions of a model.

Error recovery – Remove a bad version

- Change CEI Distribution state to ‘Inactive (event queue recoverable)’
- Repackage previous ‘good’ version with timestamp greater than ‘bad’ version
- Deploy ‘good’ version
 - ▶ Events for ‘bad’ version will be processed by ‘good’ version (assuming same filters, correlation keys)
- ‘Bad’ version can be uninstalled after it goes to ‘Inactive’ status

If you find that you have deployed a new version of a model and it is a bad version that needs to be removed then this slide shows you the procedure. You change the CEI distribution mode to ‘Inactive (event queue recoverable)’ for the bad version, then repackage a good version to deploy. You deploy the good version and assuming it has the same filters and correlation keys, then the new version will process the events queued up for the bad version. After the events are processed, then the bad version can be uninstalled.

Un-deploying a model and monitored data

- **Note:** this will undeploy all versions of the model, you cannot undeploy just one version; to keep monitored data, skip Drop DDLs & cube removal
- Stop all the model version EARs
- Stop the replication daemons
- Run stateDrop.ddl and datamartDrop.ddl in schemagen folder
 - ▶ This drops the tables containing monitored data and repository information
 - ▶ If you have more than one version, then just run the drops for the last version
 - remove the comments from both ddl files
 - db2 connect to *MONITOR*
 - db2 -tvf stateDrop.ddl
 - db2 disconnect *MONITOR*
 - db2 connect to *DATAMART*
 - db2 -tvf datamartDrop.ddl
 - db2 disconnect *DATAMART*
 - ▶ Run data services generation drop DDL scripts before uninstalling EAR file to avoid having to run purge
 - You know you need to run purge if you uninstall application, but it still shows in console as monitor model
- Un-install the enterprise application (queues are removed)

This is the procedure to undeploy all versions of a model. First you stop the model version applications using the administrative console, then you stop the replication daemons. Then you run the drop scripts which were created for you in the schemagen output folder, but note that you should just run the scripts for the most recent version. Note also that you should run the drop scripts before uninstalling the monitor model EAR to avoid having to run the purge function in the Monitor administrative console. Then uninstall the EAR which will also remove the JMS queues.

Un-deploying a model (Cont.)

- Run RM undeployment scripts
 - ▶ Unzip
C:\IBM\WebSphere\Monitor\tools\DataMovementServices\DMSCleanup.zip
 - ▶ Cd ... \dmscleanup
 - ▶ dmscleanup -model <yourModelName>
 - When prompted, enter DATAMART, MONITOR
- Drop cubes in Cube Views
 - ▶ cd <yourPath>\schemagen\<timestamp>
 - ▶ db2mdapiclient -d DATAMART -u userid -p password -i drop_model_cv.xml -o drop_response_model_cv.xml
- Alphablox > Administration > Cubes > for each cube, select it > Delete
- Portal > Administration > Portlet Management > Portlets > for each portlet based on the model, click Delete icon – they are also deleted from the Dashboards

Next, you run the replication un-deployment scripts which are supplied in the supplied DMSCleanup.zip file. You will drop the cubes in Cube Views using a supplied script which can be found in the schemagen folder. You will also drop cubes in Alphablox. Finally you will remove portlets from your dashboards that are based on the model that is being removed.

LifecycleServices MBean

- Deployment methods:
 - ▶ runDSG
 - ▶ confirmRunDMSCreateScripts
 - ▶ confirmImportDB2CubeDefinition
 - ▶ createAlphabloxCube
 - ▶ enableDashboards
 - ▶ configureCEIDistribution
 - ▶ confirmCEIServerReboot
- Other methods:
 - ▶ listModels and listModel
 - ▶ listVersions and listVersion
 - ▶ purge
 - ▶ readAlphabloxConfig, applyAlphabloxConfig and removeAlphabloxCube
 - ▶ readCEIConfig, applyCEIConfig and applyCEICredentials
 - ▶ readVersionGeneralProperties and applyVersionGeneralProperties
 - ▶ start and stop
 - ▶ countActiveMCIs

Here is a list of the MBean operations that can be added to scripts to use for command line access to the life cycle services. You see from the list that there is an operation for each of the seven steps of life cycle and other methods, such as methods to list models and versions and purge models.

LifecycleServices MBean - Usage

- JACL Example:
 - ▶ set mbean [\$AdminControl completeObjectName type=LifecycleServices,*]
 - ▶ \$AdminControl invoke \$mbean start { "Claim_id" 1154476800 }
- See information center for more information



Here is an example of using a JACL script to run one of the LifecycleServices MBean interfaces.

Check the Monitor information center for more information on these functions.

Troubleshooting

- Check the Process Server logs
 - ▶ New profile name:
`\IBM\WebSphere\ProcServer\profiles\wbmonitor\logs`
- Check the Portal Server logs
 - ▶ `/IBM/WebSphere/PortalServer/log`
- Check the Monitor logs and trace
 - ▶ Install logs in new location: `\IBM\WebSphere\Monitor\install\logs`
 - ▶ For Monitor installation problems, check `launchpad-monitor.log`.
 - ▶ To get trace details use the Administrative Console menu Troubleshooting then select Logs and Trace -> server1 -> Change Log Detail Levels. Specify log level for sub-components of `com.ibm.wbmonitor`.

This slide shows some areas to check when you are troubleshooting.

Check the logs on Monitor Server at the location shown here, or use the Process Server administrative console to view them using the Troubleshooting menu.

Troubleshooting (cont.)

- Check the replication manager logs
 - ▶ \IBM\WebSphere\Monitor\rm\logs
- Check the events
 - ▶ Make sure you have turned on CEI logging on the runtime server
 - In the administrative console: Select Application servers -> server1 -> Business process container settings -> Business process container -> Enable Common Event Infrastructure logging
 - ▶ Common base event browser
- Check WebSphere Integration Developer log
 - ▶ <workspacePath>\.metadata\.log
- Check data services generation logs
 - ▶ <generationFolderPath>\schemagen\<timestamp>
- Check task manager for Portal Server and Process Server

The Replication Manager logs are in found in the rm/logs folder.

You can turn on CEI logging in the administrative console menu for servers using the menu options shown here.

The Common Base Event Viewer is available in the administrative console in the Integration Applications menu, then select Common Base Event Browser

Also, when all servers are stopped, use the task manager to look for any large java.exe processes that may be running, which may indicate unresponsive servers.

Troubleshooting (cont.)

- Check “valid from” date in WebSphere Integration Developer process
- Check tables in MONITOR database to see if it is getting any data
 - ▶ context table: named ‘CTX_<modelName>_MC’
- Check tables in DATAMART database to see if replication manager is running
 - ▶ fact table: named ‘FCT_<yourModelName>_MC’
 - comment on table shows model name
 - list of table names: stateMapping.html, datamartMapping.html

Check the “valid from” date on the BPEL process if you are having trouble getting the process to run.

To check the Monitor database for data, you can look at the monitoring context table; the comment on the context table shows the model name. For a complete list of table names, look at files stateMapping.html and datamartMapping.html in the generation folder. Note that if the replication manager is working, completed instances are deleted here and moved to the Datamart database, so if they do not show up here, check in the Datamart database.

To check the Datamart database for data, look at the fact table; the comment on the fact table shows the model name. The alert table is located in the Datamart database and is called ACTIONMGR_ALERTS.

Troubleshooting (cont.)

- Restart Portal – if you see view errors or have refreshing problems for process data or models
 - ▶ Normally you do not start underlying WebSphere Application Server, just Portal
- Need an SVG Viewer for Dashboard
 - ▶ Tested with Adobe SVG Viewer 3.03
- Dashboard server: use IE 5.5 and up
 - ▶ Firefox will not work

You only need to start the WebSphere Application Server underlying Portal to do administrative console work. Normally, you can just stop and start Portal.

For Dashboards, you will need an SVG viewer. You can download an SVG viewer from www.adobe.com. Search for 'svg viewer' and download the latest one.

Troubleshooting (cont.)

- Manage Alphablox Cube connectivity problems
 - ▶ Check firewalls
 - ▶ Check for IP address in Portal server configuration
 - Use a machine with static IP addresses not dynamic IP addresses

For step 5 of the life cycle wizard, if you receive error messages indicating connectivity problems, you should check to see if you have firewalls that are causing this. You should also ensure that you are using a machine with static IP addresses rather than dynamic IP addresses. The dynamic IP addresses can cause invalid IP addresses to be used in the Portal server configuration.

Summary

- Covered model deployments, un-deployments, model versions, and troubleshooting

In summary, this presentation has covered many different aspects of administering WebSphere Business Monitor version 6.0.2.

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](#)



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

Alphablox is a registered trademark of Alphablox Corporation 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 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 2007. 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.