WebSphere® Business Monitor V6.0.2

Basics – Integrated test environment and databases

This presentations covers the integrated test environment and databases used by WebSphere Business Monitor V6.0.2

Goals

- Introduce WebSphere Business Monitor V6.0.2 test environment and databases

This presentation should give you a good understanding of the Monitor integrated test environment and the Monitor database structure.

Agenda

- Integrated test environment
- Monitor databases

3

This is the agenda for the presentation.

You will see how the integrated test environment works.

The Monitor database structure and replication facilities will be discussed.

Section

Integrated test environment

This section will explain using the Monitor integrated test environment.

Integrated test environment goals

- Rapid deployment of models

- Simplify iterative development environment

- Validation of monitored data

- Deployment and testing process is consistent with other WebSphere tools

5

The advantage of using the integrated test environment is that you will be able to quickly generate and deploy your monitor models into the server. And you can also easily verify that your monitor models are working, without having to install the Dashboard server. Since the integrated test environment is built upon WebSphere Integration Developer and the integrated test server, deploying monitor models will be familiar process if you have previously used the WebSphere application deployment process.

Steps for deployment and testing

- Build monitor model in Monitor model editor
- Add installed runtime to Server preferences in WebSphere Integration Developer
- Add server to the Server view in WebSphere Integration Developer
- Generate Monitor EAR in Monitor perspective
- Deploy EJB project in J2EE perspective
- Add project to the server
- Run Monitor model Setup Wizard in administrative console
- Submit test events
- Launch Web Dashboard from server view

6

These are steps to deploy a monitor model using the integrated test environment. First you build the monitor model using the Monitor model editor in WebSphere Integration Developer. You need to setup the server in WebSphere Integration Developer, then generate a monitor EAR using the Business Monitor perspective. You deploy the project to the server from the J2EE perspective. You run through the monitor model deployment steps using the Setup Wizard in the administrative console for the server. Finally you submit events to test the application and then using the Web interface to check that the monitor model is working properly.

Setup wizard for integrated test environment

Monitor Models > Version

Configuration that applies to the selected Model Version.

General

General Properties

Model
AutoParts

Version
2006-10-01T06:00:06

Application
AutoPartsEAR

CEI distribution mode
Inactive

Active MC instances
0

❌ Startable

Setup Status
- ○ Run Data Services Generation
- ○ Run Schema Create Scripts
- ◑ Run DMS Create Scripts
- ◑ Import DB2 Cube Definition
- ◑ Manage Alphablox Cube
- ○ Configure CEI Distribution
- ○ Confirm CEI Server Reboot

Version Properties
- ▪ Setup Wizard
- ▪ Runtime Configuration
- ▪ View Model

Basics – Integrated test environment and databases
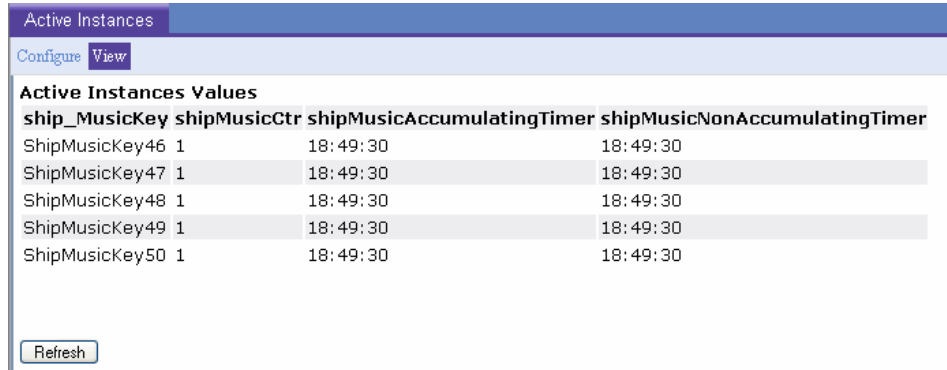
7

© 2007 IBM Corporation

This is the monitor model version panel in the administrative console for the integrated test environment. Notice that three of the steps are disabled for setup in the integrated test environment, because these steps are related to the Dashboard server which is not part of the integrated test environment. So you will not need to run Data Movement Services nor create the cubes.
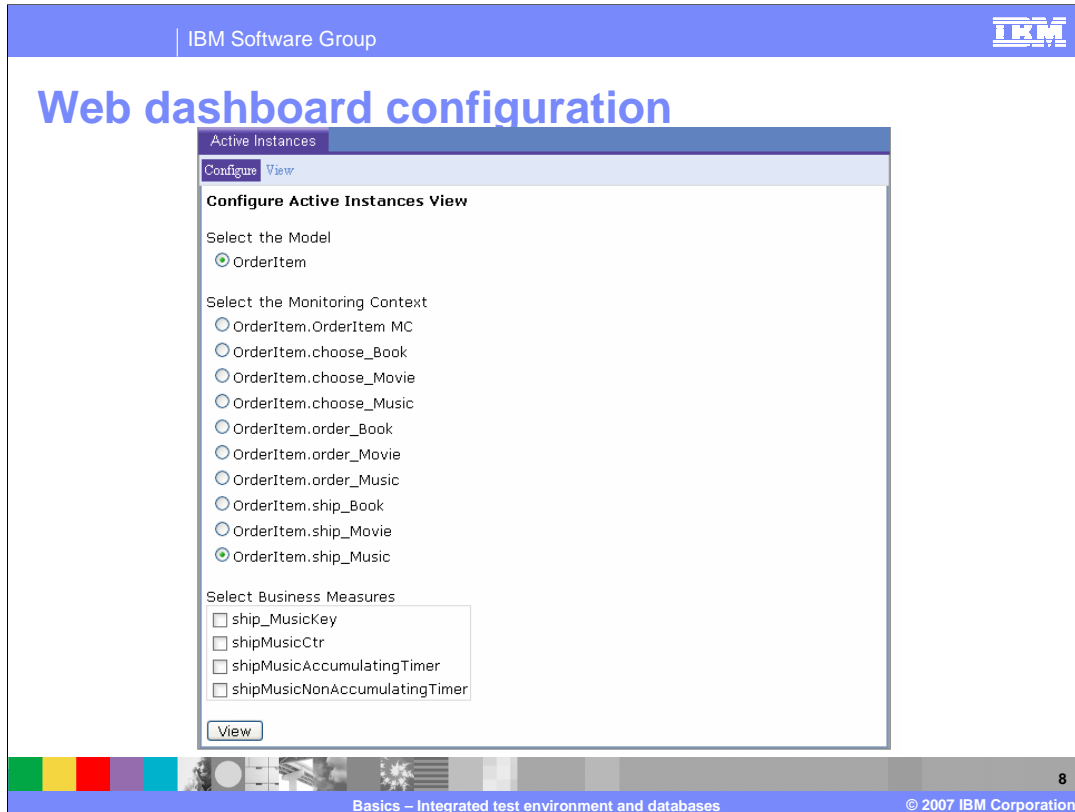
In the Web user interface for validation of your monitored data, there are two tabs, Configure and View. In the Configure tab, you select the model, monitoring contexts and business measures that you want to see on the View tab.

On the view tab of the Web interface for the integrated test environment, you can see what data is being collected for your monitoring context instances. This information is being accessed directly from the Monitor database, so the Datamart database is not used and the Dashboard components are also not used.

Section

Monitor databases

This section will describe the WebSphere Business Monitor databases.

Monitor databases

| Observation Manager | | Dashboard |

36,000 $

16 seconds   01101101

Real-time data   Historical data
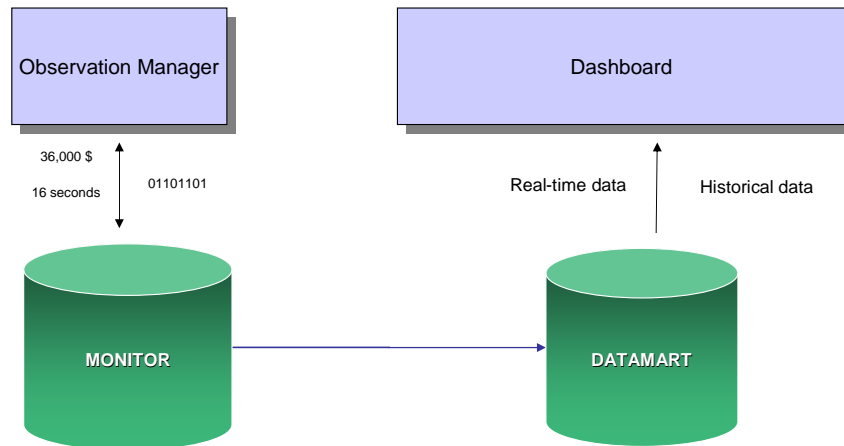
MONITOR → **DATAMART**

For performance reasons, the event processing database (Monitor) will be separated from the historical database (Datamart). The observation manager component of Monitor will extract data from the common base events and perform calculations as defined in the monitor model and store that information into the Monitor database. Schema Generators will generate replication scripts that are used to replicate data between the Monitor and Datamart databases. The Datamart database is used by the Dashboard server to display real-time data and historical data on the Monitor dashboards.

Data services generation (DSG)

- Schema generator
  - ▶ DDLs

- Cube views generator
  - ▶ Model_cv.xml

- Data movement services (DMS)
  - ▶ Replication manager setup/runtime
  - ▶ Leverages DB2® Replication Manager product

Basics – Integrated test environment and databases

12

© 2007 IBM Corporation

The Data Services Generation component of Monitor is responsible for the creation of scripts which you run for each monitor model version. These scripts will create the tables for the models, along with the cube information for Cube Views and Alphablox®, and then create the DB2 replication manager information that is necessary to replicate data between the two databases.

Replication manager

- Setup infrastructure that supports copying of Monitor model data

- Movement from Monitor database to Datamart database

Replicate every 5 minutes

(Admin settable)

Monitor database   Datamart database
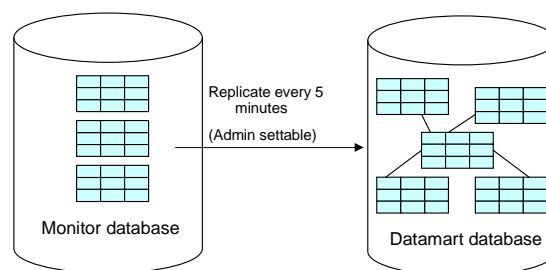
Replication manager is responsible for the efficient and accurate movement of information between the databases used by Monitor.

The Datamart database will contain one or more star schemas which will be more conducive to business integration style reporting tools such as Cube Views and Alphablox.

Replication manager loads the Datamart every 5 minutes, but this interval can be set using the Process Server administrative console.

The Monitor database will contain active process instances. The Datamart database will have active process instances, completed process instances with latency defined by the replication interval setting. The Datamart database will also contain all historical metrics, KPIs, dimensions and facts.

Model generation artifacts for database

- **state.ddl** – define the tables used by monitor server
- **datamart.ddl** –define the tables used by dashboard
- **stateDrop.ddl** –drop all tables/objects created in state.ddl
- **datamartDrop.ddl** –drop all tables/objects created in datamart.ddl
- **stateMapping.html** –mapping of model artifacts to database tables/columns defined in state.ddl
- **datamartMapping.html** –mapping of model artifacts to database tables/columns defined in datamart.ddl.
- **model_cv.xml** – defines the cube model for Cube Views
- **drop_model_cv.xml** – drops the cube model
- **DS_DMS_setup.zip** – data movement services

Basics – Integrated test environment and databases

© 2007 IBM Corporation

Data Services Generation creates many files for each model. There are DDL files for defining and deleting the tables. There are mapping files which show all the database artifacts that are created for each model. There are cube files for defining and dropping the cubes. And there is a file which is used for setting up the movement of information using Replication Manager.

MONITOR database

- Contains Repository tables for storing model metadata
- One table per monitoring context
  - ▸ CTX_<monitoringContextName>_MC
  - ▸ Contains columns for metrics, counters, stopwatches, triggers, trigger evaluation times
  - ▸ Note that table and column names are unmangled for 6.0.2
- Database location
  - ▸ DATABASE_CHARACTERISTICS table
- Deployment steps for models
  - ▸ LC_STEP table
- Alphablox server name and port
  - ▸ PROPERTY table

The MONITOR database contains repository tables which is the metadata associated with each monitor model. It also contains one table for each monitoring context in the models. The table names start with CTX followed by the monitoring context name, and inside the table there will be columns for the metrics that are defined in the model. Note that for version 6.0.2, these names are in human readable form so you can easily find the columns that are associated with the metrics that are defined in your models.

The DATABASE_CHARACTERISTICS table contains information related to your databases.

The deployment steps are logged in the LC_STEP table.

The Alphablox server information is stored in the PROPERTY table.

DATAMART database

- One fact table per monitoring context
  - FCT_<monitoringContextName>_MC
- One table per non-time dimension
  - DIM_<dimensionName>
- Action manager alerts
  - ACTIONMGR_ALERTS table
- "runtime view" – flat (non-star) view of data
  - CTR_<modelName>_MC
  - Database view – not a real table
  - Used by dashboard instances view
- DIM_TIME table – shared by all cubes
  - Referenced by "time-based" dimensions
    - Dimension containing field of type date or datetime
  - Contains day/month/year
  - Original timestamp value preserved and seen in runtime view

These are the tables that are defined in the DATAMART database. There is one fact table per monitoring context and the name starts with FCT followed by the monitoring context name. There is one table for each non-time dimension in the model. The alerts that are being produced by the Action Manager component are stored in the ACTIONMGR_ALERTS table. There is also a view that is used to display a flat view of the data for the dashboard instances view. This view is based on the model, so there is still only one view even if you have multiple monitoring contexts. And there is the DIM_TIME table which stores day, month and year information for all time based dimensions. This table does not contain the hours, minutes and seconds but that information is kept in the fact tables and is available in the runtime view.

Start or stop data movement services (DMS)

- Run start capture/apply scripts from DS_DMS_Setup.zip
  - ▸ Must run all start/stop for all versions of the model, not just the most recent
  - ▸ Most recent may not have any start/stop
- Can run consolidated start/stop instead
- One window opened per capture/apply server started – Windows® only
- Can also manage capture/apply servers using the DB2 replication center

Data movement services scripts are stored in a file which is produced by the first step of deployment. When you run the start capture scripts for a model version, this starts the replication daemons to control movement of the data from the Monitor database to the Datamart database. The DMS file will be produced for every version of a given model, so you will need to make sure that all of the replication daemons for all the versions of a model are running.   In some cases, the scripts are not produced if the model version was not changed significantly. To help you start and stop all versions of the DMS scripts, there is a consolidated script which you can run to more easily manage this.

In a Windows environment there is one command window opened for each capture and apply server, and these command windows cannot be closed. However, if you would like to avoid opening these command windows, you can manage them using the DB2 Replication Center.

References

- DB2 Replication Redbook
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg246828.pdf
- XPath 2.0, http://www.w3.org/TR/xpath20/
- Information on BPEL/HTM event structure and format
  - http://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsps.mon.doc/doc/cmon_BPC_events.html
- WebSphere Business Integration Server Foundation 5.1 and CEI RedPaper
  - http://www.redbooks.ibm.com/abstracts/redp3915.html
- Monitor Web site for samples and best practices
  - http://www-306.ibm.com/software/integration/wbimonitor/library/tutorials.html

18

© 2007 IBM Corporation

This slide shows some references including for DB2 replication, XPath, BPEL Common Base Event formats and CEI. The CEI red paper is based on Server Foundation 5.1 but is useful to understand the CEI APIs. Also, a Web site is available that contains Monitor samples and best practices.

Summary

- You reviewed Monitor integrated test environment and Monitor databases

In this presentation you have reviewed the Monitor integrated test environment and Monitor databases.

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.

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

EJB, J2EE, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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 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.

21