IBM Software Group

WebSphere® Business Monitor V6.0

Basics

@business on demand.

This presentation will provide a look at some of the basics of WebSphere Business Monitor V6.

Goals

- Introduce WebSphere Business Monitor V6.0 features
- Review the business measures modeling
- Summarize Security

This presentation introduces the process involved in modeling and monitoring business processes with WebSphere Business Monitor V6.

IBM Software Group

Agenda

- Business Measures Model
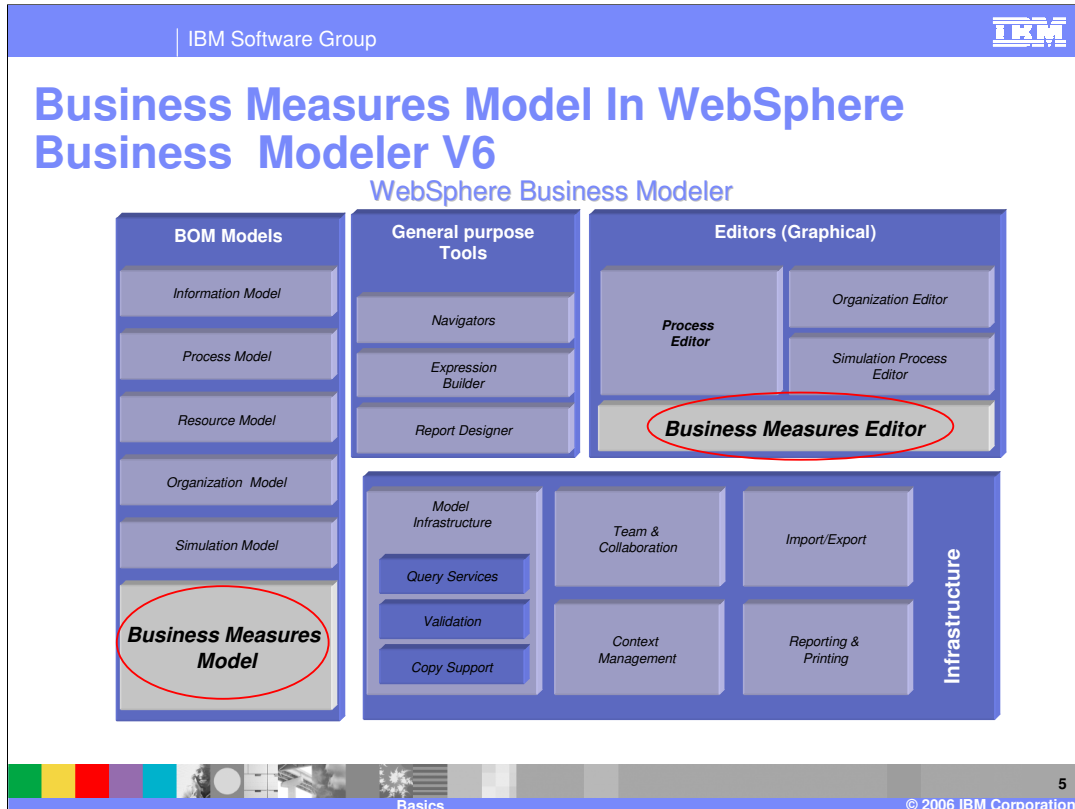- Databases
- Security

3

© 2006 IBM Corporation

This is the agenda for the presentation includes a review of the business measures model as it is defined in the WebSphere Business Modeler product. Next, it will cover the WebSphere Business Monitor databases and replication strategy. Finally it will review the security components in WebSphere Business Monitor.

IBM Software Group

Section

Business Measures Model

This section covers the Business Measures Model.

Business Measures Model In WebSphere Business Modeler V6

WebSphere Business Modeler

BOM Models	General purpose Tools	Editors (Graphical)
Information Model	Navigators	Organization Editor
Process Model	Expression Builder	Process Editor / Simulation Process Editor
Resource Model	Report Designer	Business Measures Editor
Organization Model		
Simulation Model		
Business Measures Model		

Infrastructure: Model Infrastructure, Query Services, Validation, Copy Support, Team & Collaboration, Context Management, Import/Export, Reporting & Printing

Basics

5

© 2006 IBM Corporation

The Business Measures Model is one of several models defined in WebSphere Business Modeler V6.

The Business Measures Editor(BME) is one of several WebSphere Business Modeler editing tools.

The BME uses WebSphere Business Modeler infrastructure services and components such as Expression Builder and Report Designer.

BME artifacts are managed by WebSphere Business Modeler version control (CVS, Clearcase®), they are printable using WebSphere Business Modeler Reporting & Printing components, and they can be published using WebSphere Business Modeler Collaboration components

Business Measures Model Functions

- Information **gathering**, from real-time events

- Information **aggregation** to calculate higher-level business metrics

- Providing an interface for **visualization** of metrics on a business "dashboard"

- The **recognition of situations** warranting action

Basics

6

© 2006 IBM Corporation

Information that is needed for creating a business measures model can be extracted from real-time events from your process engine.

Higher-level metrics that have business relevance like the daily sales of a product are oftentimes not readily available, but must be deduced from lower-level information. This is called aggregation.

The data that is being monitored should be displayed in easily configured dashboards consisting of flexible layouts of different views of the information.

For situations that need attention, the model should provide the capability to define events which can then be used to invoke one or more actions for dealing with the situation.

Business Measures Model Elements

- Metric

- Stopwatch

- Counter

- Trigger

- KPI

- Dimension

- Situation Event

The elements of the Business Measurements Model are listed here.

A Metric is a measure that represents the value of a monitored entity and helps in assessing performance in a particular area. A metric is sometimes used directly and sometimes as an element within other metrics. It can have a numeric value, like the average cost of sales, or non-numeric values like the delivery dates of shipments. It can be of type string, boolean, integer, float, double, date, time, date/time, duration or cost. Triggers are used to identify when the metric is calculated.

Stopwatches are metrics that are used to time specific processes. You define a Stopwatch using two triggers, a start trigger and a stop trigger.

Counters are metrics that can be used to count things such as transactions in a process.

Triggers are used to initiate the calculation of the value of a metric. A trigger might be based on the state change of a task, or the input value of the process.

The Key Performance Indicator (KPI) is a type of metric that indicates the degree of business performance. It is a quantifiable measure designed to track one of the critical success factors of a business. It may have upper or lower limits, or both, forming a range or target that the business performance should lie within. A KPI could also have a target and monitor and report the percentage of achieving that target. For example a KPI could be "an average time to shipment of less than 5 business days"

IBM

Process-Centric Business Measures Modeling

- Model with a process as the context

- Auto-generated business measures model from a process

  ▸ Generated business measures model reflects process structure

  ▸ Predefined event and metric definitions allowing process tracking

  ▸ Observation points represented as metrics, allowing propagation of "business payload" to other metrics/KPIs and to user defined outbound events

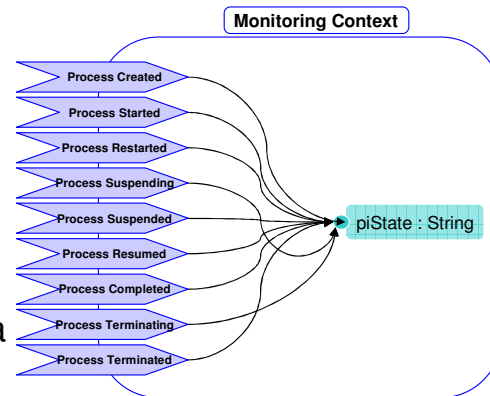 ▸ Generated model elements are runtime-independent

Basics

WebSphere Business Modeler V6 supports Business Measures Models which are based on process models. In this case the basic Business Measures Model is constructed for you automatically based on the Process Model that you have defined using the Process Editor.

Using the Business Measures Editor(BME) is optional. The business measures model that is generated will track metrics such as process start/end, activity start/end, elapsed duration, and working duration. But, using the BME, you can add additional metrics and KPIs that you would like to track.

Process To Monitoring Context

- A Monitoring Context (MC) is auto generated representing a business process

- Process MC is populated with set inbound events representing all event types

- Process MC is designed to instantiate MC instance with every process instance created

- Process instance state transition is deduced from incoming events to populate a State metric

Monitoring Context

Process Created
Process Started
Process Restarted
Process Suspending
Process Suspended
Process Resumed
Process Completed
Process Terminating
Process Terminated

piState : String

9

This slide depicts the creation of a monitoring context (MC) for a process. An MC is created for each process in the model, and it is designed to listen for all event types associated with the process. The process state is stored in a predefined metric.

The life cycle of a monitoring context matches that of the underlying process instance.

Process Monitoring Context Predefined Metrics

- The auto generated process monitoring context contains the following predefined metrics which are directly extracted from suitable events
  - ▸ Process State
  - ▸ Process Instance ID
  - ▸ Process Instance Name
  - ▸ Process Instance Description
  - ▸ Process Start Time
  - ▸ Process Is Delayed
  - ▸ Process Elapsed Duration
  - ▸ Process Working Duration
  - ▸ Similar metrics are predefined for activities in the process

10

This chart lists some of the other predefined metrics that are automatically tracked in the generated business measures model for the process. Similar metrics are also created for tasks or activities that are defined in the process.
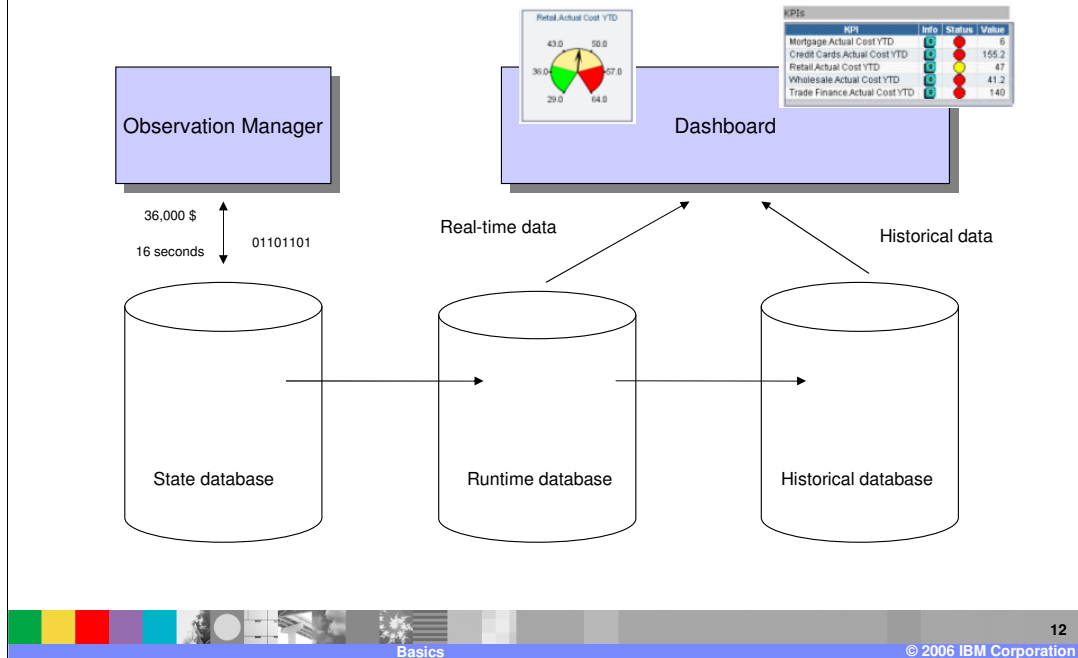
IBM

Section

Databases

This section covers the databases used by the WebSphere Business Monitor V6

For performance reasons, the event processing data store (State Data Store) will be separated from the runtime data store, and the historical database. Schema Generators will generate replication scripts that are used to replicate data between state, runtime, and historical data stores.

Databases Overview

- LaunchPad
    - Creates the databases

- Schema Generator
    - Analyzes the business measures model after you have annotated the business measures model with the Business Measures Editor (BME)
    - It will produce DDL and Replication Manager scripts for a DBA to review, possibly modify, and to run

13

The LaunchPad for WebSphere Business Monitor will collect attributes associated with database names along with the various attributes such as user ID and password. It will also give you an option to create the databases.

The Schema Generator generates the DDL scripts, Replication Manager scripts, and CubeViews meta-data. The scripts are used by the DBA to create the Star Schema of the Data Mart, and the files necessary for the cube creation in AlphaBlox, and replicating data between state, runtime, and historical data stores.

Replication Manager doesn't run the created DDL and scripts, and therefore doesn't create/modify/drop database objects. The DBA takes these outputs and makes changes if necessary and runs them. The DBA can modify the DDL, but there are limitations to what a DBA can change. No table names or Column names can be changed in the DDL
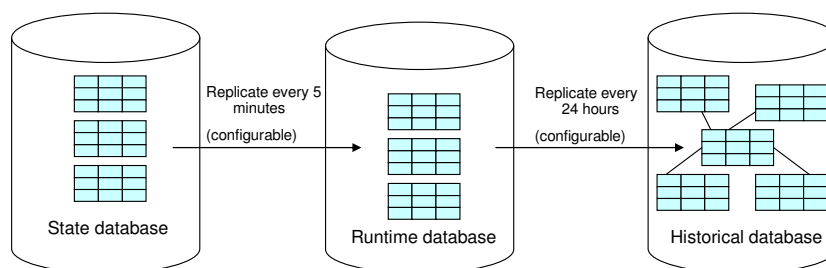
Replication is the copying of data from one place to another. Data can be extracted by programs, transported to some other location, and then loaded at the receiving location. For updated models, only the changes made since the last processing cycle are applied.

At a high level, you can think of the a row being a Monitoring Context(MC) and the columns as metrics within the MC.

WBMonitor06_Basics.ppt A BI model is generated for each monitoring context. This is a star schema that is Page 13 of 42 based on the facts, dimensions and aggregations that are identified using the BME. For

IBM

Replication Manager

- Setup infrastructure that supports copying of Business Measures Model process instance data
- Movement from
    - ▶ State to Runtime database
    - ▶ Runtime to Historical database
- Transformation (mapping) and life cycle of monitoring contexts

State database → Replicate every 5 minutes (configurable) → Runtime database → Replicate every 24 hours (configurable) → Historical database

14

Replication Manager is responsible for the efficient and accurate movement of information between the State, Runtime, and Historical databases used by WebSphere Business Monitor.
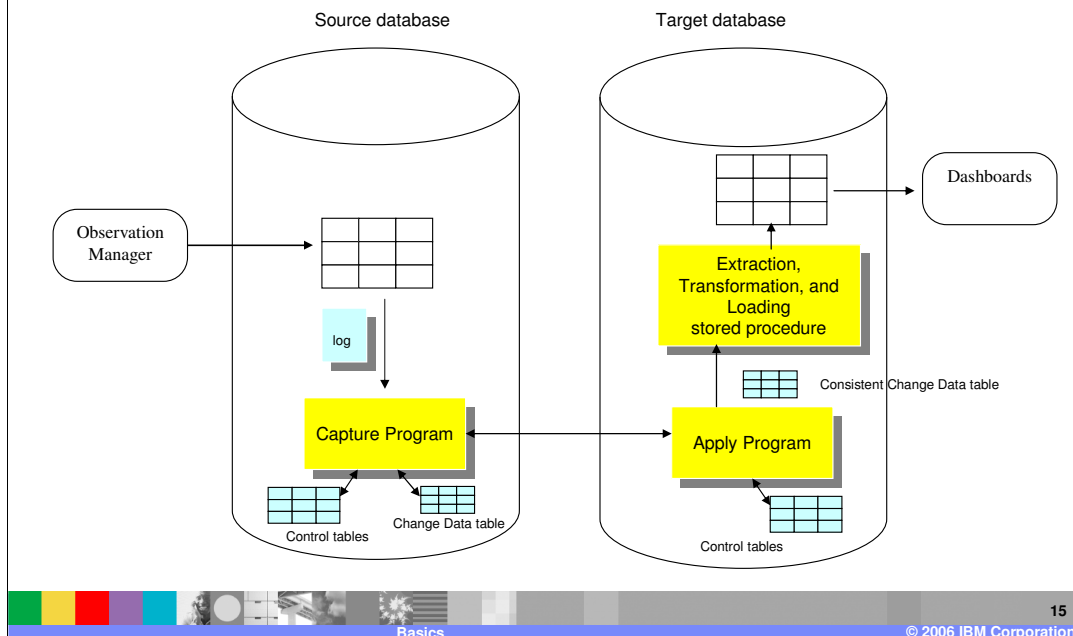
The **Runtime** database is almost identical in structure to the **State** database. The **Historical** database will contain one or more star schemas which will be more conducive to business integration style reporting tools such as Cube Views and Alphablox.

When a monitoring context has been completed in the State database, the column 'ready for delete' is set by monitor components. This completed monitoring context is removed by replication manager when it has been marked 'ready for delete'.

The Fast Replicator is loading the Runtime database in near realtime (the default is every 5 minutes but it is administratively configurable). The Slow Replicator loads the Historical datamart every 24 hours, but this is also configurable.

The State database will contain active process instances. The Runtime database will have active process instances and completed process instances, but only those that were completed within the last 24 hours. The Historical database will contain all historical metrics, KPIs, dimensions and facts.

Conceptual View – DB2® SQL Replication

Source database

Target database

Observation Manager

Dashboards

Extraction, Transformation, and Loading stored procedure

log

Consistent Change Data table

Capture Program

Apply Program

Control tables

Change Data table

Control tables

15

Basics

© 2006 IBM Corporation

This slide depicts replication using the DB2 Replicator which uses SQL-based replication architecture that maximizes flexibility and efficiency in managing scheduling, transformation, and distribution topologies. In SQL replication, DB2 captures changes using the log and inserts them into a relational staging table. An apply process asynchronously handles the updates to the target systems.

Changes to DB2 source tables are captured by a Capture program running at the source server. The captured changes are stored in a

table local to the source table and are automatically deleted after they have been applied. When changes are made to the source table, DB2 writes log (journal) records. These log records are used for database recovery and for replication. Each source table has a corresponding Change Data table where the captured changes are stored.

Captured changes are applied to target tables by Apply programs. The changes are stored in the Consistent Change Data table and then after the Apply program stores a set of changes, it runs a stored procedure to perform the Extraction, Transformation, and Loading to the target table.

The control tables are created when the Replicator is configured and these are used by the Capture and Apply programs to control the replication activities. The Capture program uses a set of DB2 Capture control tables that contain information about replication sources and the current position of the Capture program in the DB2 recovery log. The Apply program uses a set of DB2 UDB Apply control tables that contain information about the targets and where their corresponding sources are located.

Section

Security

The next section will address security.

Security

- Authentication using WebSphere Application Server and a user registry

- Authorization to portlets using Portal Security
    - Portal Security Administration portlets
    - Stored in Portal database
    - WebSphere Member Manager

An *authentication mechanism* defines rules about security information (for example, whether a credential is forwardable to another Java™ process), and the format of how security information is stored in both credentials and tokens. Authentication is the process of establishing whether a client is valid in a particular context. A client can be either a user, a machine, or an application. An authentication mechanism in WebSphere Application Server typically collaborates closely with a *user registry*. The user registry is the user and groups account repository that the authentication mechanism consults with when performing authentication.

Authorization is sometimes referred to as access control. Authorization determines what interactions a user is permitted to have with a resource. Administrators configure access to Portal resources by assigning roles to users and groups. A role combines a set of permissions with a specific WebSphere Portal resource. Roles are assigned to users and groups that are contained in the user registry.

Resources that can be protected by user include pages, portlets, and web module. The User and Group Permissions portlet allows you to view and modify the roles that users and groups have on WebSphere Portal resources. The Manage Users and Groups portlet allows you to view, create, and delete users and user groups. You can also change group memberships. The Resource Permissions portlet allows you to manage permissions on resources and assign roles and users/groups to resources.

Member Manager is a component of WebSphere Portal that manages data for users and groups. Users and groups are referred to as "members" in Member Manager. Member Manager keeps track of the overall attribute set of the users and groups within the system and the values of those attributes for individual users and groups.

Summary

- This presentation reviewed the concepts of business measures modeling, the WebSphere Business Monitor Databases, and the security features of WebSphere Business Monitor

18

This presentation reviewed the concepts of business measures modeling, the WebSphere Business Monitor Databases, and the security features of WebSphere Business Monitor V6.

References

- DB2 Replication Redbook
  - http://www.redbooks.ibm.com/redbooks/pdfs/sg246828.pdf

- Alphablox documentation – user and admin
  - C:\IBM\DB2Alphablox\system\documentation\help\Intro.html

Template Revision: 11/22/2005 12:10 PM

Trademarks, Copyrights, and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

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

Linux is a registered trademark of Linus Torvalds.

Other company, product and service names may be trademarks or service marks of others.

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 and/or changes in the product(s) and/or program(s) 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 (e.g., 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 2005,2006. 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.

20