TransactionVision®

Programmer's Guide Version 4.2.1

Bristol Technology Inc. 39 Old Ridgebury Road Danbury, CT 06810-5113 USA (203) 798-1007 Bristol Technology BV Plotterweg 2A 3821 BB Amersfoort The Netherlands +31 (0)33 450 50 50 Printed April 3, 2005

This manual supports TransactionVision Release 4.2.1 SupportPac A. No part of this manual may be reproduced in any form or by any means without written permission of:

Bristol Technology Inc. 39 Old Ridgebury Road Danbury, CT 06810-5113 U.S.A.

Copyright © Bristol Technology Inc. 2000 - 2004

RESTRICTED RIGHTS

The information contained in this document is subject to change without notice.

For U.S. Government use:

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at 52.227-7013.

All rights reserved. Printed in the U.S.A.

The information in this publication is believed to be accurate in all respects; however, Bristol Technology Inc. cannot assume responsibility for any consequences resulting from its use. The information contained herein is subject to change. Revisions to this publication or a new edition of it may be issued to incorporate such changes.

Bristol Technology® and TransactionVision® are registered trademarks of Bristol Technology Inc. The IBM e-business logo, zSeries, z/OS, S/390, OS/390, OS/400 and WebSphere MQ are all trademarks of IBM Corporation. All other trademarks herein are the property of their respective holders.

General Notice: Some of the product names used herein have been used for identification purposes only and may be trademarks of their respective companies.

Part No. TV15050331 SupportPac A

Contents

1.	Introduc	ction	1
	1.1. Cha	nges in TransactionVision 4.2.1 SupportPac A	1
	1.2. Cha	nges in TransactionVision 4.2.1	1
	1.3. Cha	nges in TransactionVision 4.2	2
	1.4. Cha	nges in TransactionVision 4.1	2
	1.5. Prer	equisites	3
2.	Archite	cture Overview	5
	2.1. Syst	em Components	5
	2.2. Data	abase	6
	2.3. Use	r Interface Framework	7
3.	Tutorial	- Extending the Analyzer	9
	3.1. Hov	v to Handle XML Message Data in Events	9
	3.1.1.	Step 1: Modify the Beans.xml file to use the DefaultModifierBean	10
	3.1.2.	Step 2: Verify that XML data is extracted correctly	10
	3.2. How	v to Handle Custom Message Data Formats in Events	10
	3.2.1.	Step 1: Document message format(s) layout	10
	3.2.2.	Step 2: Document the target XML format	11
	3.2.3.	Step 3: Implement the bean to do the format conversion	11
	3.2.4.	Step 4: Modify the Beans.xml file to use the custom bean	15
	3.2.5.	Step 5: Test the custom bean in the Analyzer environment	16
	3.3. Ove	rview of XDM Files	16
	3.4. Hov	v to Map Custom Message Data Fields to Database Tables	16
	3.4.1.	Step 1: Determine which fields in the XML event document need to be mapped	to
	database	columns	16
	3.4.2.	Step 2: Determine the database column names for these fields	17
	3.4.3.	Step 3: Construct XDM file entries	18
	3.4.4.	Step 4: Recreate your project database schema	19
	3.4.5.	Step 5: Verify that the XDM mapping works correctly	19
	3.5. Add	litional XDM File Examples	19
	3.6. Hov	v to classify business transactions and map transaction attributes to database table	S
			21
	3.6.1.	Overview of Transaction Classification:	22
	3.6.2.	Task Description:	22
	3.6.3.	Implementation:	23
	3.6.3.	1. Step 1: Determine the event attributes that apply to a business transaction	23
	3.6.3.	2. Step 2: Determine database column names for these fields	23
	3.6.3.	3. Step 3: Extract transaction attributes from event data	24
	3.6.3.	4. Step 4: Construct XDM file entries for transaction attributes	26
	3.6.3.	5. Step 5: Determine the transaction classes and their classification criteria	27
	3.6.3.	6. Step 6: Implement classification rules	27
	3.6.3.	7. Step 7: Recreate the project database schema	28
	3.6.3.	8. Step 8: Enable classification in the Analyzer	28

	3.6.3.9. Step 9: Verify that the transaction classification works correctly and the	
	transaction attributes are written correctly	29
	3.7. How to perform custom correlation of related events	29
	3.7.1. Overview of Custom Event Correlation:	29
	3.7.2. Task Description:	30
	3.7.3. Implementation:	30
	3.7.3.1. Step 1: Determine correlation requirements	30
	3.7.3.2. Step 2: Determine which events need to be correlated and common	
	correlation data between the events	30
	3.7.3.3. Step 3: Implement the correlation bean	30
	3.7.3.4. Step 4: Enable the Analyzer to invoke the correlation bean	35
	3.7.3.5. Step 5: Test the correlation bean	35
4.	Reference - Extending the Analyzer	36
	4.1. Using the Beans.xml File	36
	4.2. Unmarshalling Message Data	36
	4.2.1. The Default Modifier Bean	36
	4.2.2. Adding a Message Data Unmarshal Bean	37
	4.2.3. IEventModifier Interface	37
	4.2.4. XML Related Classes	37
	4.2.4.1. Class XMLEvent	37
	4.2.4.2. Class XPathSearch	39
	4.2.4.3. Class XMLParser	41
	4.2.4.4. Other Utility Classes	42
	Interface DOMElement	43
	Class EventElement	43
	Class TextElement	43
	Class ByteElement	44
	Class ByteStringElement	45
	Class IntElement	45
	Class Intelement	45
	Class Intrexclement	40
	• Class LongElement	40
	• Class LongHexElement	47
	Class StringElement	47
	Class RawStringElement	48
	4.2.5. Sample Usage of the IEventModifier Interface	48
	4.3. Trimming Data From an Event	51
	4.3.1. Interface IDBWriteExit	51
	4.4. XML-Database mapping Using XDM Files	52
	4.5. Performing Event Analysis	54
	4.5.1. Event Analysis Utility Classes and Interface	55
	4.5.1.1. Interface Cache	55
	4.5.1.2. Class ConnectionInfo	57
	4.5.1.3. Class EventID	57
	4.5.1.4. Class TechEventID	58
	4.5.2. Event Analysis Classes	58
	4.5.2.1. Interface IAnalyze	58
	4.5.2.2. Class AnalyzeEventCtx	59
	4.5.2.3. Class AnalyzeEventBean	59
	4.5.3. Adding Custom Correlation Analysis Beans	59
	4.5.5.1. Interface IEventCorrelation	62

	4.5.3.2	2. Class CorrelationTechHelperBean	63
	4.5.3.3	3. Class MQCorrelationData	63
	4.5.3.4	I. Class JMSCorrelationData	. 64
	4.5.3.5	5. Class LookupKey	. 65
	4.5.3.6	5. Class EventRelation	. 65
	4.5.3.7	7. Class MQRelationDBService	. 66
	4.5.3.8	3. Class JMSRelationDBService	. 67
	4.5.3.9	O. Sample Custom Event Correlation Bean	. 68
	4.5.4.	Custom Business Transaction Attributes and Classification	. 70
	4.5.4.1	. Transaction Classification	. 72
	4.5.4.2	2. Transaction Classification with the Standard Classification Bean	. 72
	4.5.4.3	3. Classification Action Rules	. 75
	4.5.4.4	4. The ClassifyTransactionCtx and the IClassifyTransaction Interface	. 76
	4.5.4.5	5. Writing a Custom Classification Bean	. 77
	4.5.4.6	5. The Transaction Class Table	. 78
	4.5.4.7	7. Business Groups	. 80
	4.6. Exte	nding the System Model	. 81
	4.7. Gene	erating Application Events to Tivoli Enterprise Console (TEC)	. 82
	4.7.1.	Class MonitoringEvent	. 82
	4.7.2.	SlotMap.properties	. 83
	4.7.3.	Example Usage:	. 83
	4.7.4.	BTV Class Definitions and Rulebase	. 83
5.	Using th	e Query Services	. 85
	5.1. Sam	ple Usage	. 85
	5.2. Class	s QueryServices	. 86
	5.2.1.	Methods:	. 87
	5.3. Clas	s QueryDoc	. 95
	5.3.1.	Constructors	. 96
	5.3.2.	Methods	. 97
	5.4. Clas	s QueryDoc.WhereClause	104
	5.4.1.	Fields	104
	5.4.2.	Constructors	104
	5.4.3.	Methods	106
	5.4.4.	Example	106
	5.5. Inter	face Query	107
	5.5.1.	Methods	107
	5.6. Inter	face Cursor	107
	5.6.1.	Methods	107
	5.7. Clas	s DataManagerException	112
	5.7.1.	Constructors	112
	5.7.2.	Methods	113
6.	Extendir	ng the User Interface	115
	6.1. Writ	ing TransactionVision Reports	115
	6.1.1.	Report Interfaces	117
	6.1.1.1	I. IReportData	117
	6.1.1.2	2. IReportAction	117
	6.1.1.3	BaseReportBean	118
	6.1.2.	TransactionClass	119
	6.1.3.	JSP Custom Tag Library	119
	6.1.3.1	The Report Tag	119
	6.1.3.2	2. The Form Tag	120

	6.1.3.3. Tag Reference	120
	6.1.3.4. Migrating the Form Tag from TransactionVision 4.0	122
	6.1.3.5. Deprecated Tags	123
	6.1.3.6. Report Example	129
	6.1.4. Adding a Report to the Framework	132
	6.1.4.1. Required Configuration Information	133
	6.1.4.2. Optional Configuration Information	133
	6.1.4.3. Adding Actuate Reports	134
	6.2. Adding Query Pages	136
	6.3. User Interface Utility Classes	138
	6.3.1. Class TVisionServlet	138
	6.3.1.1. Methods	138
	6.3.2. Class TypeConvService	139
	6.3.2.1. Methods	139
	6.4. Using Job Beans	141
	6.4.1. JobBean	142
	6.4.2. IJob Interface	142
	6.4.3. Creating Jobs at Project Creation	143
7.	Database Schema	145
	7.1. System Object Model Tables	145
	7.2. Event Tables	148
	7.3. Event Relationship Tables	152
	7.4. Transaction Tables	153
	7.5. Statistics Tables	156
	7.6. User Preference Tables	157
	7.7. Administration (System) Tables	159
8.	Event XML Schema	165
	8.1. Basic Types	165
	8.2. Event Schema Description	166
9.	The Data Manager	169
	9.1. Using the DataManager to Access the Database	169
	9.2. XML-Database Mapping Using XDM Files	172
	9.3. The XDM Syntax	173
	9.3.1. Creating the XDM database tables	175
	9.3.2. Properties of the TransactionVision Document Types	176
	9.3.2.1. The /Event Document Type	176
	9.3.2.2. The /Transaction Document Type	176
	9.3.2.3. The /TransactionClass Document Type	176
	9.4. The XMLDatabaseMapper Interface	176
	9.5. Extending the /Event Document Type	178
	9.6. Extending the /Transaction and /TransactionClass Document Type	179
	Adding New Document Types	179

1. Introduction

This guide provides details of how the TransactionVision platform can be extended and programmed against to achieve better control over its various functions. This manual presents an architecture overview of the TransactionVision system and documents the different methods available to use and extend the analyzer service, the query service, project manager services and the TransactionVision user interface.

1.1. Changes in TransactionVision 4.2.1 SupportPac A

A new column "timerule_status" has been added to the business_transaction table. This column indicates if a start or end time rule has fired during during transaction classification. The possible values are: 0=no rule fired, 1=start time rule fired, 2=end time rule fired, 3=both time rules fired. This column is used internally by the Analyzer and should not be modified by the user.

1.2. Changes in TransactionVision 4.2.1

• In TransactionVision 4.2.1, the key definition for the business_trans_id key column has changed. In earlier releases, this key definition was as follows:

```
<Key name="business_trans_id" type="INTEGER"
generated="true"
description="TransactionId">
<Path>/Transaction/BusinessTransId</Path>
</Key>
```

The new definition in TransactionVision 4.2.1 is as follows:

<Key name="business_trans_id" type="INTEGER" generateSequence="true" description="TransactionId"> <Path>/Transaction/BusinessTransId</Path> </Key>

If you have a custom transaction XDM file that uses the old definition, you will encounter the following error:

TransactionVision Error (XDMInconsistentDefForKeyColumn): Inconsistent Key definition in XDM file 'Transaction.XDM'.

This error indicates that the definition for a key column is not exactly the same for two (or more) XDM files.

• You can use the IAnalyzerAction interface to specify custom actions to be performed for specific classification values. For more information, see 4.5.4.3, Classification Action Rules.

• State and Result constants in the following file have changed <TVISION_HOME>/config/technologyconst/TransactionConst.xml

1.3. Changes in TransactionVision 4.2

The following changes in TransactionVision 4.2 may require existing custom beans or reports to be changed accordingly.

- The **DeleteEvents** utility and job uses an optimized fast deletion scheme based on timestamp columns, resulting in changes to the TransactionVision database schema. See Chapter 7 for updated schema table diagrams.
- In order for the new deletion scheme to delete user defined transaction XDM tables, it is necessary to write the transaction end time into each new XDM table. See Chapter 9 for more information.
- Two new options have been added to to the **DeleteEvents** utility: -threadcount and -nosplit (only valid for the -older option). For more information, see the *TransactionVision Administrator's Guide*.
- A business group table has been added to allow grouping of transaction classes into business groups. See Chapter 3 for more information.
- A CICS lookup table has been added to store CICS event related data. See Chapter 7 for more information.
- In the report framework, the TransactionClassLookup class, which gets access to the transaction classification definitions, has been replaced with com.bristol.tvision.datamgr.dbtypes.TransactionClass. This class is similar to the previous class, but its usage is slightly different. See Chapter 5 for information about this class.

1.4. Changes in TransactionVision 4.1

The following public interfaces changed in TransactionVision 4.1. Any exiting custom beans that use these interfaces must be changed accordingly. For detailed information, see section 3.5.3.

- Interface com.bristol.tvision.services.analysis.eventanalysis.IEventCorrelation
- Constructor of class com.bristol.tvision.services.analysis.eventanalysis.CorrelationTechHelperBean
- Constructor of class com.bristol.tvision.datamgr.dbtypes.LookupKey

In addition, changes have been made to the JSP custom tag library used to create TransactionVision reports. For detailed information, see section 5.1.2.

1.5. Prerequisites

To use this guide to customize TransactionVision requires knowledge of some of the following technologies and APIs:

- The Java programming language
- The Java Database Connectivity (JDBC) API
- J2EE concepts related to JSPs and servlets
- Relational database and SQL knowledge

This guide is designed to explain common programming tasks to extend TransactionVision. It is not a comprehensive guide and your customization needs may go beyond what this guide defines. In that case, please contact Bristol Technology technical support for assistance.

Important! This documentation is related to the internals of the TransactionVision product; incorrect changes could break the functioning of the product.

2. Architecture Overview

2.1. System Components

TransactionVision consists of the following logical components:

- The **Sensor** component generates events based on the technology being sensed. The sensor gets configuration and filtering messages from the configuration queue and sends events into the event queue. The event and configuration queues are represented by the "communication infrastructure" box in the diagram on the following page.
- The Analyzer component is responsible for retrieving and analyzing events from the communication link. It contains a chain of Java bean contexts, each performing a particular function on the event data. Each bean context can hold multiple chained beans to perform technology specific or other custom processing of the event data. The beans in each bean context are controlled by the <TVISION_HOME>/config/services/Beans.xml file. The main components of the Analyzer include:
 - The **EventModifier bean context** is responsible for converting raw event data from its binary format into XML. This bean context provides an environment for user message data unmarshaller beans to be plugged in.
 - The **DBWriteExit bean context** allows a custom bean to trim or cut down on the data written into the database. This gives a user flexibility to cut down on storage size. Typically this is an XSLT which processes the XML tree generated by the unmarshaller context.
 - The **Database write context** is responsible for mapping the XML tree generated by the unmarshaller and trim contexts to database tables and writing the tree into the database. This context uses the XML data mapper component to map the XML tree to relational database tables.
 - The **Analysis context** performs event correlation, local and business transaction analysis, transaction classification, statistics analysis and any other custom data analysis.
- The **XML data mapper** and **database query services** components provide a means of mapping XML data to relational tables and a XML based query to an SQL statements based on relational tables.

The following diagram shows the TransactionVision architecture layout:

Chapter 2 • Architecture Overview *Database*



Figure 1 TransactionVision Architecture

2.2. Database

The general table organization consists of a TVISION schema, where project, communication links, filter, queries and other administration related information is stored, and project-specific schemas, where events collected by a project are stored. Each project schema consists of an event table, where the event identifier and the XML event are stored, and several lookup tables that provide indexes to the event. In addition there are several other tables in a project schema storing event correlation, local transaction, business transaction and other WebSphere MQ objects related information.

2.3. User Interface Framework

The Presentation, User and Session management components interact with the user interface, transferring requests for data into calls into the database component and returning the results of the requests back to views in the user interface. The framework consists of servlets and JSPs running under a web application server such as WebSphere. This framework also manages user login into the system and all aspects of the user's access rights to the TransactionVision system. Users can create new views and reports to be plugged into this framework.

The Security Management component gets user rights, that control what view and data is accessible to a user, from an LDAP server. This component allows administrators to assign rights to users and groups of users to aspects of TransactionVision functionality and data collected. This includes defining policies for starting and stopping data collection, changing data collection filters and access rights to collected data. Refer to the Administration Guide for setting up these user rights.

The Configuration and Administration component manages administration of the TransactionVision system. This includes starting/stopping data collection and event processing, changing data collection filters, providing system status, and administering security policies on the Analyzer service. The Analyzer is controlled using embedded RMI and can run on systems different from the application server.

This programmer's guide provides the details of extending TransactionVision. It is important to note that since this documentation is related to the internals of the product, incorrect changes could break the functioning of the product.

3. Tutorial - Extending the Analyzer

The Analyzer reads in binary event packets from the TransactionVision Event Queue and processes them through a chain of bean contexts. Each bean context performs a specific function to analyze and write data from the event into the database. Many of these operations can be extended and customized to perform transformations based on your systems or application needs. This chain of beans is defined by the Beans.xml file. The sequence of bean contexts includes:

- The event modifier context, which allows users to write custom beans to modify the incoming event, such as convert binary message data into XML.
- The data writer context, which contains beans to write the data into various relational database tables.
- The analysis context, which contains various beans to perform event analysis, transaction analysis and correlation of events to create a business transaction.

Each context holds beans that perform a default function and can be replaced or added on to perform further actions on the data being processed. The following sections document common tasks related to extending the Analyzer.

3.1. How to Handle XML Message Data in Events

Task Description:

When your message data is already composed of XML, a custom bean is not required to have the XML processed by the Analyzer. Instead, TransactionVision provides a default modifier that can be used to attach the message data XML contents to the TransactionVision event.

Implementation:

This section describes how to set up the TransactionVision DefaultModifierBean, which detects XML data in the message field and appends it to the XML event. The default event modifier bean,

"com.bristol.tvision.services.analysis.eventmodifier.

DefaultModifierBean" scans the user data for any XML data and, if found, simply adds it to the Event XML document at the position

"/Event/UserData/Chunk [@seqNo='n']" wher n is the number of the data range (defined in the data collection filter).

3.1.1. Step 1: Modify the Beans.xml file to use the DefaultModifierBean

Edit the file Beans.xml under the <TVISION_HOME>/config/services directory to uncomment the following line of XML:

```
<!--Module type="Bean"
class="com.bristol.tvision.services.analysis.eventmodifier.DefaultMo
difierBean"/-->
```

The changed line is:

```
<Module type="Bean"
class="com.bristol.tvision.services.analysis.eventmodifier.DefaultMo
difierBean"/>
```

Refer to Section 4.1, "Using the Beans.xml File" on the layout and format of the Beans.xml file.

3.1.2. Step 2: Verify that XML data is extracted correctly

Re-start the Analyzer after making the changes in Step 1. Collect events using the TransactionVision sensors from your application. For each event that generates XML message data, go to the event detail view and verify that the XML data shows up in the user data panel.

Once this task is done, the XML message data can be mapped to custom database tables based on the kind of analysis that is required to be performed on the message data. Section 3.4 describes how to implement this mapping.

3.2. How to Handle Custom Message Data Formats in Events

Task Description:

Typically, event data from applications may contain binary, text or XML data embedded within the message. This data is often in custom and proprietary formats that are not known to the TransactionVision Analyzer. A common task is to convert these custom formats into XML within the Analyzer for later use in reports, for analysis, browsing or querying. The TransactionVision Analyzer allows for embedding a Java bean that implements the IEventModifier interface to perform the format conversion. This bean can modify the event being currently processed to do any kind of format conversion on the event data.

Implementation:

This section describes the steps to write a custom event modifier bean in Java to extract and convert binary message data and insert it into the event data as XML. A custom bean that converts a text message into XML will be used as an example. The sample code used below is from the CICS Accounting sample shipped with TransactionVision. The source code can be found at "<TVISION_HOME>/samples/CICSAccounting".

3.2.1. Step 1: Document message format(s) layout

The first step in the process of writing a bean to handle custom event data is to know the layout of all message formats in the event data and document them.

Consider the sample message to contain the following text layout, with fields Account Number, Last Name, First Name and Account Type:

AccountNumber	LastName	S	FirstName	S	AccountType	•••
(5 characters)		-		-	(1 character-S/M/C)	

In the above layout, the first 5 bytes are the AccountNumber field, while the remaining fields of "LastName", "FirstName" and "AccountType" are separated by a space separator "S". The "LastName" and "FirstName" fields are variable length fields. The "AccountType" field is one character and can be either "S" (Savings), "M" (Money Market) or "C" (Checking). The remaining fields are ignored and not required to be processed by TranactionVision.

3.2.2. Step 2: Document the target XML format

First design the target XML document to be created from the above text message. The following is a sample resulting XML structure:

```
<Event sort="false">
...
<Data>
<Chunk blobType="2" ccsid="0" from="0" seqNo="0" to="382">
<Account>
<AccountNo>11111</AccountNo>
<LastName>DOE</LastName>
<FirstName>JOHN</FirstName>
<AccountType>Saving</AccountType>
</Account>
</Chunk>
</Data>
...
```

Here, an "Account" node is created under the item "/Event/Data/Chunk". This is the point where TransactionVision stores references to message data. Hence, this is a good point, though not the only point, where any XML converted message data can be added. Under the "Account" node, nodes for "AccountNo", "LastName", "FirstName" and "AccountType" are created and their values filled in.

The XPath values of each of the above fields are as follows:

```
AccountNo - "/Event/Data/Chunk/Account/AccountNo"
LastName - "/Event/Data/Chunk/Account/LastName"
FirstName - "/Event/Data/Chunk/Account/ FirstName"
AccountType - "/Event/Data/Chunk/Account/AccountType"
```

3.2.3. Step 3: Implement the bean to do the format conversion

This section describes the implementation of the Java bean to perform the format conversion described in Steps 1 and 2.

1. A Java class, CICSAccountingModifierBean, extends from the base class EventModifierBean and implements the modify method of the IEventModifier interface. This modify method is invoked by the Analyzer framework to perform any custom data conversion tasks.

```
public class CICSAccountingModifierBean extends EventModifierBean {
    ...
    /** Creates new CICSAccountingModifierBean */
    public CICSAccountingModifierBean() {
    ...
    public void modify(XMLEvent event) throws EventModifyException {
    ...
    The above code defines the CICSAccountingModifierBean class with a method
    modify. This method accepts the current XML event object as its input and is
```

modify. This method accepts the current XML event object as its input and is allowed to modify this object in any way, typically for transforming data from proprietary formats to XML.

2. The following code fragment from the modify method first verifies whether this is the right CICS event whose message data needs to be processed, and then processes chunks of message data.

```
00055
                  String tech =
event.getDocumentValue(XPathConstants.TECH NAME);
                  if (tech == null || !tech.equalsIgnoreCase("CICS"))
00056
00057
                      return;
00058
00059
                  String eventType =
event.getDocumentValue(XPathConstants.CICS_COMMON_EVENTTYPE);
                 if (eventType == null)
00060
00061
                      return;
00062
00063
                  int type = Integer.parseInt(eventType);
                  if (type != CICSConstants.B CICS TYPE FC)
00064
00065
                      return;
00066
```

In the above code, the constant XPathConstants.TECH_NAME contains the value to the technology XPath expression. The XPathConstants class contains various other commonly used XPath expression values. Hence, line 55 extracts the value of the technology name from the event document. Line 56 ignores all events that are not from the CICS sensor (ie. are not of technology CICS). Line 59-64 lookup the event type of the CICS event. Only file control APIs (CICSConstants.B_CICS_TYPE_FC) are considered for further processing. The method getDocumentValue returns the value of any XPath location in the DOM tree in the XMLEvent object. This method is further documented in Section Error! Reference source not found..

3. The following code fragment shows how to obtain pieces of user data from the XMLEvent object.

XPathSearch lookup = new XPathSearch(event);

00067

00068	/*
00070	* get user data chunks by retrieving all
/Event/Data/Chunk	nodes
00071	*/
00072	VodeLigt dataChunkg -
lookup getNodeg (XI	
100kup.gechodes (M	if (dataChunka null dataChunka gotIongth()
0)	II (datachuliks == hull datachuliks.gethength() ==
0)	
00074	recurn;
00075	
00076	int ChunkNum = dataChunks.getLength();
00077	
00078	/*
00079	* process each user data chunk
00080	*/
00081	for (int i = 0; i < chunkNum; i++) {
00082	try {
00083	processUserData(event, lookup,
(Element) (dataChur	nks.item(i)));
00084	}
00085	catch (XMLException xmlEx) {
00086	throw new EventModifyException(xmlEx):
00087	}
00088	}
	J

Line 67 creates an XPathSearch object, whose function is to perform lookups on the XMLEvent document. The getNodes and getValues methods on the XPathSearch class enable lookups based on given XPath expressions. Section 4.2.4.2 has the documentation on this class and its methods.

The message data in a TransactionVision event is stored as a series of chunks. This is done since the message data from TransactionVision Sensors can be broken up based on data ranges specified in the data collection filter. The XMLEvent document contains the location and byte count of each of these chunks and can be looked up using the XPath expression "/Event/Data/Chunk". Typically, if no data range is specified in the data collection filters, only one chunk is created.

Line 72 gets a list of data chunk nodes. For each of the data chunks, the method processUserData is called to perform the format conversion.

4. The following code fragment is from the processUserData method which converts the text message into XML.

```
private void processUserData(XMLEvent event, XPathSearch
00104
lookup, Element owner) throws XMLException {
00105
              int chunkId =
00106
Integer.parseInt(owner.getAttribute(ATTR CHUNK ID));
00107
00108
               * find the XMLEvent.Blob which has the same chunk ID
00109
               * /
00110
              XMLEvent.Blob chunkBlob = null;
00111
00112
              Iterator it = event.blobIterator();
00113
              while (it.hasNext() && (chunkBlob == null)) {
00114
00115
                  XMLEvent.Blob blob = (XMLEvent.Blob)it.next();
                  if (blob.id == chunkId)
00116
```

Chapter 3 • Tutorial - Extending the Analyzer How to Handle Custom Message Data Formats in Events

00117 chunkBlob = blob; 00118 } 00119 00120 if (chunkBlob == null) 00121 return; 00122 00123 int ccsid = chunkBlob.ccsid; 00124 00131 * if ccsid <=0 use ccsid in standard header 00132 */ 00133 00134 if (ccsid <= 0) { 00135 ccsid = Integer.parseInt(lookup.getValue(XPATH_EVENT_CCSID)); 00136 } 00137 00138 String strBuf = Translator.instance(ccsid).translate(chunkBlob.blob); 00139 00142 00143 * parse XML document, and if succeeds, append it to 00144 owner node, 00145 * then change the data type of the chunk. 00146 */ 00147 Element acctRoot = event.createElement("Account"); 00148 owner.appendChild(acctRoot); 00149 StringElement acctNo = new StringElement("AccountNo"); 00150 00151 StringElement lastName = new StringElement("LastName"); 00152 StringElement firstName = new StringElement("FirstName"); StringElement acctType = new 00153 StringElement("AccountType"); 00154 00155 StringTokenizer tokens = new StringTokenizer(strBuf); 00156 int cnt = tokens.countTokens(); if (cnt < 3)00157 00158 return; 00159 00160 String[] tokenStrs = new String[cnt]; 00161 for (int i = 0; i < cnt; i++) 00162 tokenStrs[i] = tokens.nextToken(); 00164 } 00165 00166 // check for numeric value 00167 char c = tokenStrs[0].charAt(0); 00168 if (!Character.isDigit(c)) 00169 return; 00170 acctNo.value = tokenStrs[0].substring(0, 5); 00171 00172 lastName.value = tokenStrs[0].substring(5); firstName.value = tokenStrs[1]; 00173 00174 00175 int idx = strBuf.indexOf("FSR"); 00176 acctType.value = strBuf.substring(idx + 3, idx + 4); 00177 if (acctType.value.equalsIgnoreCase("M")) acctType.value = "Money Market"; 00178 00179 else if (acctType.value.equalsIgnoreCase("S")) 00180 acctType.value = "Saving"; else 00181 00182 acctType.value = "Checking"; 00183 00184 acctNo.toDOM(event, acctRoot);

00185		<pre>lastName.toDOM(event, acctRoot);</pre>
00186		<pre>firstName.toDOM(event,acctRoot);</pre>
00187		<pre>acctType.toDOM(event, acctRoot);</pre>
00188		
00189		chunkBlob.type = TVisionCommon.XMLEVENT BLOB XML;
00190		
00191		owner.setAttribute(ATTR_BLOBTYPE,
00192		Integer.toString(TVisionCommon.XMLEVENT BLOB XML)
00193);
00194	}	
00195 }	-	

The method processUserData converts one chunk of message data text into XML. Line 106 obtains the id of the current chunk of message data being processed.

Lines 111-119 access the array of message data binary objects (BLOB) that are stored in a separate table in the same sequence as the chunks in the XML document. Typically, there is just one object, but there could be more depending on whether data ranges have been set in the data collection filter. Hence, a chunk in the XML document with the ID 1 will have its equivalent BLOB in the user data table at the sequence number 1. The chunk id from the XML document is matched with the message data BLOB ID.

Lines 134-138 find the codepage of the message data and convert it to the code page the Analyzer is using. This is required because the message data is from CICS and needs to be converted from EBCDIC to ASCII.

Lines 147-148 create new nodes in the XMLEvent document to hold the Account related data.

Lines 150-154 create new objects of type StringElement. This class is a TransactionVision utility class that has the ability to generate XML DOM nodes from input values. Refer to Section 4.2.4.4 for details on this class. The toDOM method of this class creates and appends XML DOM nodes to a DOM tree at a specified location.

Lines 155-183 is Java code which parses the message data string buffer and extracts values for Account, LastName, FirstName and AccountType based on the format defined in Step 2.

Lines 184-188 convert the parsed message data from their StringElement values into DOM nodes attached to the XMLEvent DOM tree at location "/Event/Data/Chunk/Account".

3.2.4. Step 4: Modify the Beans.xml file to use the custom bean

The event modifier bean implemented in the previous steps needs to be enabled in the event modifier context of the Beans.xml file. Change the Beans.xml file to add the following line:

The Analyzer needs to be re-started after this change.

3.2.5. Step 5: Test the custom bean in the Analyzer environment

To verify that the above data extraction is working correctly, check the right events user data buffer in the event detail view. In the example above, check the user data for the file control READ API.

3.3. Overview of XDM Files

Certain pieces of information in the message data may be useful to be queried upon by custom reports or analysis modules. In that case, these fields need to be extracted from the message data and mapped to database columns by the Analyzer. Before these fields can be written to a database column by the Analyzer, they need to be extracted from the message and converted to XML (if not already in the XML format). Section 3.2 describes how to extract binary message data and convert it to XML and Section 3.1 describes how to handle XML message data.

The TransactionVision database schema is made extensible through the XML to Database Mapping (XDM) files. As message data specific columns are added to the database, the XDM files can be updated to describe the new schema. Hence XML to Database mapping serves several purposes:

- To describe to the CreateSqlScript program the layout of the project database schema tables.
- To describe to the Analyzer the fields that are to be extracted from the XML event data and stored in event lookup tables for fast searching and retrieval.
- To describe to the Analyzer the fields that are to be extracted from the transaction XML document and stored in the transaction lookup tables.
- To describe the database schema to the query services for use in TransactionVision user interface views and reports.

3.4. How to Map Custom Message Data Fields to Database Tables

Task Description:

The task in this section describes how to map event XML data to database fields using TransactionVision's XDM (XML to Database Mapping) module.

3.4.1. Step 1: Determine which fields in the XML event document need to be mapped to database columns

Consider a WebSphere MQ MQPUT request event which has the following XML segment in its message data:

```
<Event>

<Data>

<Order>

<ID>123456</ID>

<Branch>Danbury</Branch>

<Account></Account>

<Ticker>MSFT</Ticker>

<Price>88.88</Price>

<Shares>1000</Shares>
```

```
</Order>
</Data>
</Event>
```

Consider a WebSphere MQ MQPUT reply event in response to the above request that contains the following XML segment in its message data:

```
<Event>

<Data>

<Result>

<ID>123456</ID>

<Type>Stock</Type>

<Status>Success</Status>

</Result>

</Data>

</Event>
```

3.4.2. Step 2: Determine the database column names for these fields

The mapping of message data to database columns enables custom business reports and queries to be written to view and analyze the contents of the message data.

Consider that the following fields need to be mapped to database columns from the message data described in Step (1).

For the MQPUT request message data, a TRADE_ORDER table can be defined as follows:

Field Name	SQL Type	Length
ORDERID	VARCHAR	16
BRANCH	VARCHAR	16
ACCOUNT	VARCHAR	8
TICKER	VARCHAR	8
PRICE	VARCHAR	8
SHARES	VARCHAR	8
PROGINST_ID	INTEGER	4
SEQUENCE_NO	INTEGER	4

For the MQPUT reply message data, a TRADE_RESULT table can be defined as follows:

Field Name	SQL Type	Length
ORDERID	VARCHAR	16
TYPE	VARCHAR	8
STATUS	VARCHAR	12
PROGINST_ID	INTEGER	4
SEQUENCE_NO	INTEGER	4

In both the above tables, PROGINST_ID and SEQUENCE_NO are event identification fields that are required to join with the TransactionVision EVENT table,

while the remaining columns contain business content to be extracted from the message data.

3.4.3. Step 3: Construct XDM file entries

Now that we have determined the format and contents of the message data in Step 1 and which database tables need to be populated in Step 2, a mapping can be created from the XML message data contents to the database columns.

Consider the following XML segment:

```
<Event>
<Data>
<Order>
<ID>123456</ID>
...
</Order>
</Data>
</Event>
```

The XPath to the Order ID field can be written as: "/Event/Data/Order/ID".

The value at this XPath needs to be written to the ORDERID column of the TRADE_ORDER table.

This mapping can be done in an XDM file as follows:

```
<Table name="TRADE_ORDER" category="MQSERIES,JMS">
    <Column name="ORDERID" type="VARCHAR" size="16"
    description="OrderID">
        <Path>/Event/Data/Order/ID</Path>
    </Column>
    ...
```

The above XDM file segment defines a table name TRADE_ORDER in the "Table" element. The table contains a column ORDERID, defined by the "Column" element, of type VARCHAR and size 16 bytes. The "Column" of name ORDERID has an XPath mapping, defined by the "Path" element to be "/Event/Data/Order/ID".

The table definition part of the XDM segment is applied when a new project schema is created either by **CreateSqlScript** or the project creation web pages. The XPath mapping part of the XDM segment is applied by the Analyzer when processing events. When an event contains data at the XPath value "/Event/Data/Order/ID", the Analyzer extracts the value and writes a row to the mapped column ORDERID belonging to table TRADE_ODER for that event. The "category" attribute for the "Table" element, indicates that this mapping is applied only to MQSeries and JMS events.

The complete mapping of the MQPUT request message to the TRADE_ORDER table is as follows:

```
<Table name="TRADE_ORDER" category="MQSERIES,JMS">
<Column name="orderid" type="VARCHAR" size="16"
description="OrderID">
<Path>/Event/Data/Order/ID</Path>
</Column>
```

```
<Column name="branch" type="VARCHAR" size="16"
description="Branch">
               <Path>/Event/Data/Order/Branch</Path>
         </Column>
         <Column name="account" type="VARCHAR" size="8"
description="AccountNumber">
               <Path>/Event/Data/Order/Account</Path>
         </Column>
            <Column name="ticker" type="VARCHAR" size="8"
  description="Ticker">
                     <Path>/Event/Data/Order/Ticker</Path>
         </Column>
         <Column name="price" type="VARCHAR" size="8"
  description="Price">
                     <Path>/Event/Data/Order/Price</Path>
         </Column>
         <Column name="shares" type="VARCHAR" size="8"
  description="NumberOfShares">
                     <Path>/Event/Data/Order/Shares</Path>
         </Column>
   </Table>
```

3.4.4. Step 4: Recreate your project database schema

The TransactionVision Analyzer and Web component need to be restarted for the modified XDM files to have effect. Once the Web component is restarted, when new project schemas are created, they will contain the newly defined tables or columns. However, existing database project schemas need to be updated to create the newly added tables or columns. This can be done using options in the **CreateSqlScript** utility.

For example:

CreateSqlScript -c -e -n -p TEST -t TRADE ORDER

The above command creates the table TRADE_ORDER as defined in the XDM file in the TEST database schema.

3.4.5. Step 5: Verify that the XDM mapping works correctly

Start Analyzer collection for the project that has the custom XDM mapping. Generate events containing the message data with the expected XPath entries. Verify that rows are written into the TRADE_ORDER table for every event containing the expected message data.

3.5. Additional XDM File Examples

The XDM mappings can be technology or platform specific. The common mapping defined in the file <TVISION_HOME>/config/xdm/Event.xdm (data in the standard event header) will be written for every event, but the mappings defined in the other XDM files will only be applied if the current event matches the mapping's "category' (technology or platform) definition. The XML schema format of XDM files is defined in <TVISION_HOME>/config/xmlschema/XDM.xsd. The following code is an extract from the file Event.xdm.

<?xml version="1.0"?>

```
<Mapping documentTable="event" documentColumn="event_data">
        <Key name="proginst_id" type="INTEGER"
description="ProgramInstanceId">
                <Path>/Event/EventID/@programInstID</Path>
        </Kev>
        <Key name="sequence no" type="INTEGER"
description="SequenceNumber">
                <Path>/Event/EventID/@sequenceNum</Path>
        </Key>
        <Table name="EVENT_LOOKUP" category="COMMON">
            <Column name="host_id" type="INTEGER" description="Host"
isObject="true">
                <Path>/Event/StdHeader/Host/@objectId</Path>
            </Column>
            <Column name="program_id" type="INTEGER"
description="Program" isObject="true">
    <Path>/Event/StdHeader/ProgramName/@objectId</Path>
            </Column>
        </Table>
</Mapping>
```

The above snippet from Event.xdm defines a table EVENT containing the XML document and a table EVENT_LOOKUP, containing various indexed columns of data from the XML document. The key columns proginst_id and sequence_no are integer types and mapped to XPath expressions /Event/EventID/@programInstID and /Event/EventID/@sequenceNum. These key columns are primary keys common to the EVENT and EVENT_LOOKUP tables. Similarly, the columns host_id and program_id are mapped to XPath expressions /Event/StdHeader/Host/@objectId and /Event/StdHeader/ProgramName/@objectId respectively.

The preceeding XDM file specifies that when an XML event is written to the database by the DBWrite module in the Analyzer, these fields are extracted and written into the database columns mapped to in the XDM file. Similarly, when the database is queried using the QueryServices XML interface, these XDM files are used to construct the corresponding SQL query.

The isObject attribute for a Column tag in the XDM file refers to that column being an identifier for an object in the system model table. The documentTable and documentColumn tags are the table and column where the actual XML document is stored. The key is the primary key and is common to the document table and the lookup tables. Each lookup column is indexed.

The queryOnly attribute for a Column tag indicates that the value is not written by the Analyzer in the DBWrite module, but maybe written in the analysis phase of the Analyzer or by some other application. Hence, this field is for queries only.

The generated attribute for a Column tag means that column is a database generated id.

The conversionType attribute for a Column tag means that field requires a formatting conversion before writing to the database. The TypeConvService is called into before writing that field into the database. This is typically used for writing date/time or enumeration fields.

The category attribute on the Table tag contains either COMMON or the technology string or the platform string for the event data that should be written into this table. The string COMMON indicates that this table contains data common to every event and should be written for every event going through the Analyzer. A technology or platform name like "MQSERIES" or "OS390_BATCH" used in the category field indicates that this table should only be filled for events of that technology or platform.

```
<Table name="EVENT_LOOKUP" category="COMMON">
...
</Table>
<Table name="OS390_LOOKUP"
category="OS390_BATCH,OS390_CICS,OS390_IMS">
...
</Table>
```

A column can map to multiple XPath expressions, as in the following sample code. This assumes that only one of the XPaths will exist in a given event document.

```
<Column name="datasize" type="INTEGER" description="DataSize">
    <Path>/Event/Technology/MQSeries/MQGET/MQGETExit/DataLength
    </Path>
    <Path>/Event/Technology/MQSeries/MQPUT/MQPUTExit/BufferLength
    </Path>
    <Path>/Event/Technology/MQSeries/MQPUT1/MQPUT1Exit/BufferLength
    </Path>
</Column>
```

Additionally, business transaction attributes (as opposed to event attributes) can also be mapped to transaction based XDM files. Section 3.6 describes how to map transaction attributes to transaction XDM tables.

Refer to Section 9.2 for details on the XDM file layout.

3.6. How to classify business transactions and map transaction attributes to database tables

3.6.1. Overview of Transaction Classification:

Transaction classification allows users to partition their business transactions into different transaction classes and set transaction attributes based on event data. These classes may be created based on data in the messages flowing through the business system. A transaction is classified to a transaction class when attributes in one or more events in the transaction match the criteria defined in the TransactionDefinition.xml. This file supports creating simple rules to classify transactions. This file also allows setting of attributes on transactions. These attribute values can be extracted from one or more events belonging to that transaction. These attributes then can be mapped to database tables using XDM files.

Consider a business system consisting of a JSP/servlet based user interface, a middletier based on EJBs and a mainframe based backend. The following sample classification criteria may be applied to such a system:

- Based on the types of business systems these transactions involve. For example, if the 3-tier system described above supports financial transactions such those dealing with stocks and bonds, transaction classes may be created based on this.
- Based on statistics that need to be collected for each class. Such statistics may include service level and response time requirements for different classes of transactions. In the 3-tier system described earlier, aggregate response times could be measured for each tier of the system.
- Actions or rules fired for different classes of transactions. In the 3-tier system described, email alerts may need to be fired to different administrators based on response times exceeding a threshold. Once, transaction classification has been performed, these kind of alerts can be fired based on which class a transaction belongs to.

The Transaction Tracking Report lists transaction classes and attributes automatically along with common attributes such as start time, response time etc. For more information about this report, see Chapter 7, "Using Reports," in the *TransactionVision User's Guide*.

3.6.2. Task Description:

The task in this section describes the following:

- How to extract event data and map that data to transaction attributes.
- How to map transaction attributes to database tables using transaction XDM files.
- How to write rules in the TransactionDefinition.xml file to perform transaction classification.

The sample message data used in this section is from the TRADE demo system, for which the project and event databases are shipped with TransactionVision. Refer to the *TransactionVision Administration Guide* on how to setup the TRADE demo database.

The previous sections in this chapter have discussed mapping event attributes to database tables. This section describes how to map business transaction attributes to database tables. This involves extracting attributes from events that apply to the

business transaction the event belongs to and writing them to business transaction XDM tables.

3.6.3. Implementation:

3.6.3.1. Step 1: Determine the event attributes that apply to a business transaction

Consider a request event which has the following XML segment in its message data:

```
<Event>
   <Data>
         <Order>
           <Account>123456</Account>
           <Transaction>Danbury</Transaction>
           <Type></Type>
           <Product>MSFT</Product>
           <Quantity>88.88</Quantity>
           <!-- present in FX transactions -->
           <Currency>1000</Currency>
           <RecvAccount>1000</RecvAccount>
           <!-- present in Bond transactions --> <Maturity>1000</Maturity>
           <Issue>1000</Issue>
           <!-- present in Equity transactions -->
           <Symbol>1000</Symbol>
         </Order>
   </Data>
</Event>
```

Three kinds of transactions flow through this TRADE system: Bond, Equity and FX (foreign exchange). Besides a common header, each transaction type has data specific to that transaction.

Consider the reply event in response to the above request that contains the following XML segment in its message data:

```
<Event>

<Data>

<Order>

<ID>123456</ID>

<Region>Stock</UnitPrice>

<Status>Success</Status>

<Reason>Success</Reason>

<!-- present in Bond transactions -->

<Yield>5.94</Yield>

</Order>

</Data>

</Event>
```

3.6.3.2. Step 2: Determine database column names for these fields

The mapping of message data to transaction database columns enables custom business reports and queries to be written to view and analyze the contents of the business transaction. Consider that the following fields need to be mapped to database columns from the message data described in Step 1.

Field Name	SQL Type	Length
ORDERID	VARCHAR	20
REGION	VARCHAR	12
ACCOUNT	VARCHAR	12
TRADETYPE	VARCHAR	12
TRADEACTION	VARCHAR	12
AMOUNT	DOUBLE	8
STATUS	VARCHAR	12
REASON	VARCHAR	32
BONDISSUE	VARCHAR	12
BONDMATURITY	INTEGER	4
EQUITYSYMBOL	VARCHAR	8
VALUE	DOUBLE	8
CUSTOMER	VARCHAR	32
BUSINESS_TRANS_ID	INTEGER	4

The TRADE_BUSINESS_TRANSACTION table is defined as below:

In the above table, the BUSINESS_TRANS_ID column is a transaction identification field that is required to join with the TransactionVision BUSINESS_TRANSACTION table, while the remaining columns contain business content that are extracted from the message data.

3.6.3.3. Step 3: Extract transaction attributes from event data

Now that we have determined the format and contents of the message data in Step 1, these event fields need to be set as transaction attributes. This is done in the TransactionDefinition.xml rules file with the help of "Attribute" elements with "ValueRule" elements to set values into attributes. A transaction XML document is created by the Analyzer in memory as attributes are set and this document is then mapped to database tables defined in the transaction XDM file.

Consider the mapping rule below from the TransactionDefinition.xml file for the TRADE sample database:

Here a transaction attribute called "OrderID" has been defined, with an XPath location of "/Transaction/OrderID". A "ValueRule" of name "SetOrderID" sets the value of the transaction attribute at XPath "/Transaction/OrderID" from the attribute value in the event data at XPath

"/Event/Technology/Servlet/Response/Headers/Header[@name='orderid']".

The two important pieces of information in the above attribute rule are the event XPath, which is the source of the data, and the transaction XPath, which is the destination to which the source data is copied into.

Value rules can also set constant values into transaction attributes. In the following XML snippet, a constant value of "Completed" is set into the transaction attribute "State" at XPath location "/Transaction/State".

The attribute rules can be used in the context of class rules, which determine that the attribute rules are applied only for certain classes. Consider the example below:

```
<?xml version="1.0"?>
<TransactionDefinition>
   <Class name="Bond" dbschema="TRADE">
      <Classify id="1">
         <Match xpath="/Event/StdHeader/ProgramName"
operator="EQUAL" value="TradeServlet"/>
         <Match
xpath="/Event/Technology/Servlet/Request/Parameters/Parameter[@name=
'product']/@value" operator="EQUAL" value="Bond"/>
         <Attribute name="OrderID">
            <Path>/Transaction/OrderID</Path>
            <ValueRule name="SetOrderID">
               <Value type="XPath">
/Event/Technology/Servlet/Response/Headers/Header[@name='orderid']/@
value</Value>
            </ValueRule>
         </Attribute>
. . .
```

Here, the attribute rule of name "OrderID" is applied only for already classified transactions of class "Bond".

Attribute rules also can have match criteria such that the rules are applied to every event when a match criteria is successful. Consider the XML snippet below:

Here, the value rule to set the value of "Amount" at XPath location "/Transaction/Amount" from the event XPath "/Event/Data/Chunk/Order/Amount", is fired when the logical AND of the Match criteria evaluate to True.

Refer to Section 4.5.4 for details on the syntax of the classification rules.

3.6.3.4. Step 4: Construct XDM file entries for transaction attributes

Now that we have determined the contents of the transaction attributes and extracted them from the event data as in Step (1) and (3) and determined which database tables need to be populated as in Step (2), a mapping can be created from the XML transaction attributes to the database columns.

Consider the below transaction document created by rules set XML segment: <Transaction>

```
<OrderID>123456</OrderID>
<Account> </Account>
<Region> </Region>
<TradeType> </ TradeType >
<TradeAction> </TradeAction>
<Amount> </Amount>
...
```

```
</Transaction>
```

The XPath to the OrderID field can be written as: "/Transaction/OrderID".

The value at this XPath is to be written to the ORDERID column of the TRADE_BUSINESS_TRANSACTION table for the business transactions for which this value is set.

This mapping can be done in an XDM file as follows:

```
<Mapping documentType="/Transaction">
  <DBSchema>Trade</DBSchema>
  <Key name="business_trans_id" type="INTEGER"
generateSequence="true" description="TransactionId">
        <Path>/Transaction/BusinessTransId</Path>
  </Key>
  <Table name="TRADE BUSINESS TRANSACTION">
        <Column name="orderid" type="VARCHAR" size="20"
description="OrderID">
              <Path>/Transaction/OrderID</Path>
        </Column>
        <Column name="account" type="VARCHAR" size="12"
description="Account">
              <Path>/Transaction/Account</Path>
        </Column>
        . . .
```

The above XDM file segment, the "Table" element defines a table name TRADE_BUSINESS_TRANSACTION. The table contains a column ORDERID, defined by the "Column" element, of type VARCHAR and size 20 bytes. The "Column" of name ORDERID has an XPath mapping, defined by the "Path" element to be "/Transaction/OrderID". The key for the TRADE_BUSINESS_TRANSACTION is defined by the "Key" element to be business_trans_id column of type INTEGER. The table definition part of the XDM segment is applied when a new project schema is created either by the CreateSqlScript or the project creation web pages. The XPath mapping part of the XDM segment is applied by the Analyzer when processing events. When a transaction contains data at the XPath value "/Transaction/OrderID" set by the classification rules, the Analyzer extracts the value from the transaction document and writes a row to the mapped column ORDERID belonging to table TRADE_BUSINESS_TRANSACTION for that transaction. The "DBSchema" attributes indicates that this mapping is applied only to transactions being written to the "Trade" schema.

3.6.3.5. Step 5: Determine the transaction classes and their classification criteria

Transaction classification can be based on a variety of different criteria based on the transactions flowing through your business systems. In the sample TRADE system, transaction classification is performed based on the type of financial transactions flowing through the system, namely Equity, Bonds and FX (Foreign Exchange). Hence, the next step would be to identify fields in the message data which identify the event and its transaction to be one of these three types. For this system, this field is an attribute "Product" in the XPath element

"/Event/Technology/Servlet/Request/Parameters/Parameter". The next section describes how to build a classification rule using this XPath value.

3.6.3.6. Step 6: Implement classification rules

Consider the below XML segment from the TransactionDefinition.xml sample file for the TRADE sample:

In the above segment, the element "Class" defines a transaction class called "Bond", which applies to the database schema "TRADE". Following the "Class" element is a "Classify" element, which specifies one or more classification rules for the "Bond" transaction class. The "Match" elements specify the rule criteria. The first "Match" element has a rule which evaluates to True when the XPath value of

"/Event/StdHeader/ProgramName" in an event equals the value of "TradeServlet". Multiple "Match" elements are logically AND'd together. The second "Match" criteria evaluates to True if a servlet event with the XPath element

"/Event/Technology/Servlet/Request/Parameters/Parameter" whose attribute "product" has a value of "Bond". In other words, any event with the program name "TradeServlet" and a request parameter value of "Bond" is classified to be in the "Bond" transaction class. Similarly, the below set of classification rules classify transactions to classes "Equity" and "FX".

Once a transaction is classified, attributes are attached to the transaction based on the "Attribute" rules in the TransactionDefinition.xml file. The rules for setting and writing attributes are described in Steps 3 and 4.

3.6.3.7. Step 7: Recreate the project database schema

Existing database project schemas need to be updated to create the newly added tables or columns. This can be done using options in the CreateSqlScript utility.

For example:

```
CreateSqlScript -c -e -n -p TRADE -t TRADE_BUSINESS_TRANSACTION
The above command creates the table TRADE BUSINESS TRANSACTION as
```

defined in the XDM file in the TRADE database schema.

3.6.3.8. Step 8: Enable classification in the Analyzer

By default, TransactionVision does not classify the business transactions it processes.

To enable transaction classification, the following steps are required:

</Module>

- Define your classification rules in the file TransactionDefinition.xml. This has been completed in the previous steps.
- Insert each class name into the database table "TRANSACTION_CLASS". This table must be populated before any transactions are processed by the Analyzer. For example, for the three transaction classes discussed in the previous steps, the strings "Bond", "Equity" and "FX" need to be inserted into this table.

where <schema> is the database schema into which these transactions are written to. Additional attributes of the transaction class, such as SLA thresholds, costs per transaction etc. can also be populated into this table.

The TransactionVision Analyzer and Web component need to be restarted, after the changes in above steps, so that when new projects are created, the XDM file changes are applied to create the TRADE_BUSINESS_TRANSACTION table.

3.6.3.9. Step 9: Verify that the transaction classification works correctly and the transaction attributes are written correctly

The results of the above steps can be verified by looking at the "Transaction Tracking Report", which can be accessed by going to the report "Where are my transactions?" from the Reports page. For each business transaction, this report will show you the class of the transaction and any custom attributes that have been set for that transaction. Other custom reports may be written based on the transaction attributes collected.

3.7. How to perform custom correlation of related events

3.7.1. Overview of Custom Event Correlation:

By default, the TransactionVision Analyzer correlates WebSphere MQ MQPUT and MQGET events or JMS send and receive events based on certain criteria such as message id, correlation id, put time and other fields in these events. However, at times these criteria may not be sufficient to perform event correlation and these criteria may either need to be expanded to include other data fields, such as those from the message data, or may need to be relaxed to exclude some of the standard fields or may need to be modified. This can be done by implementing a custom event correlation Java bean. Here are some scenarios where a custom correlation bean may need to be implemented:

• TransactionVision Sensors may not be installed on some systems, such as those belonging to external sensors. Hence, the messages going out to the un-sensored systems would need to be correlated with the replies coming back from these systems.

- Unique message ids or correlation ids are not used by the applications. In this scenario, custom fields from the message data may need to be used to correlate message PUTs and GETs.
- An application that replies to a message swaps the message id and correlation id fields and this application is not monitored by TranactionVision sensors.

3.7.2. Task Description:

This task walks through the creation of a simple event correlation bean. The requirement for the bean is to correlate WebSphere MQ events for which the message id and correlation ids have been swapped.

3.7.3. Implementation:

3.7.3.1. Step 1: Determine correlation requirements

Consider two applications A and B, where application A is monitored by a TransactionVision sensor while application B is not. The sequence of events for this system is as follows:

- Application A performed an MQPUT on a queue q1, with message id m1 and correlation id c1.
- Application B read the message using an MQGET from queue q1 and processed the message.
- Application B then placed a reply message using MQPUT on the reply-to queue q2, with message id c1 and correlation id m1. Hence, the message ids and correlation ids were swapped by application B.
- Application A performed an MQGET to read this message.

Now, because application B is not sensor'd and its MQGET/MQPUT are not received, this transaction path remains un-correlated and no message flow arc is drawn between application A's MQPUT and application A's MQGET. The custom event correlation bean seeks to complete this path.

3.7.3.2. Step 2: Determine which events need to be correlated and common correlation data between the events

For this task, the requirement is to correlate a MQPUT event from application A with a MQGET event from the same application A, which have their message id and correlation ids swapped.

3.7.3.3. Step 3: Implement the correlation bean

Writing a correlation bean involves implementing the IEventCorrelation Java interface. This interface is described in Section 4.5.3.1
Implementing this interface requires writing code which implements two methods, "createLookupKeys" and "correlateEvents". The class "CorrelateTechHelperBean" serves as the base class for any class implementing this interface.

The correlation process in the Analyzer is divided into two phases:

- The first phase involves generating lookup keys based on the characteristics of the current event. This lookup key is then inserted into the database and then used to match up with other correlated events as they arrive into the Analyzer. This phase is implemented by the "createLookupKeys" method of the event correlation bean. Hence, for application A for a MQPUT event, a lookup key comprising of the message id needs to be created, while for an MQGET event from application A, a lookup key comprising of the correlation id should be created.
- The second phase involves relation generation. Specifically, a set of events is passed as potential candidate for matching with the current event. This set is composed of the events that have the same lookup key as the current event. The purpose of this phase is to further narrow down set of event matches based on additional criteria which have not been covered by the lookup key data. For example, for application A, the correlation should only be performed between MQPUTs and MQGETs and not between APIs of the same type. This phase is implemented by the "correlateEvents" method of the event correlation bean.

The following steps walk through the implementation of the event correlation bean. Refer to the sample code shipped with the TransactionVision at

- <TVISION_HOME>/samples/stock for the complete source code listing.
- The following code fragment shows how the event correlation bean class is defined.

```
public class StockTradeRelationshipBean extends
CorrelationTechHelperBean {
    // user defined correlation types should be >= 100
    public static final int REQUEST_REPLY_TYPE = 100;
    /**
     * Creates a new instance.
     */
    public StockTradeRelationshipBean() {
        super(TVisionCommon.TECH_NAME_MQSERIES,
        REQUEST_REPLY_TYPE);
        }
        ....
}
```

The class StockTradeRelationshipBean extends from the base class CorrelationTechHelperBean. The constructor sets the technology name to be of type "MQSeries", since MQSeries events will be processed by this bean. The type of the bean can be any number greater than or equal to 100 and is defined by the constant "REQUEST_REPLY_TYPE".

• The below code fragment implements the "createLookupKeys" method.

```
00065 /**
00066 * Generates lookup key for relation lookup table.
```

Chapter 3 • Tutorial - Extending the Analyzer How to perform custom correlation of related events

00067 00068 * @param conInfo The database connection setup 00069 * @param event Event's XML doc * @param lookupKeys The generated lookup keys 00070 00071 * @throws AnalyzeEventException Pass along caught exception to the caller. 00072 public void createLookupKeys(ConnectionInfo conInfo, XMLEvent 00073 event, List lookupKeys) throws AnalyzeEventException { 00074 00075 00076 try { XPathSearch lookup = new XPathSearch(event); 00077 00078 String correlId; 00079 /* for StockTrade->MQPUT call (request event), use 00080 MQMD.MsgID as the lookup key, for StockTrade->MQGET call (reply 00081 event), use 00082 MQMD.CorrelId as the lookup key */ switch (StockTradeHelper.getEventType(lookup)) 00083 00084 case StockTradeHelper.MQSERIES REQUEST EVENT: 00085 correlId = lookup.getValue(XPathConstants.MSGID); 00086 if (correlId == null) 00087 return; 00088 break; 00089 case StockTradeHelper.MQSERIES REPLY EVENT: 00090 if (Integer.parseInt(lookup.getValue(XPathConstants.COMPCODE)) != 00091 MQDefs.MQCC FAILED) 00092 return; 00093 correlId = lookup.getValue(XPathConstants.CORRELID); if (correlId == null) 00094 00095 return; 00096 break; 00097 default: 00098 return; } 00099 00100 /* create a new lookup key and add it to the list */ 00101 00102 LookupKey key = new LookupKey(correlId, REQUEST REPLY TYPE); lookupKeys.add(key); 00103 00104 00105 catch (XMLException ex) { 00106 throw new AnalyzeEventException(ex); } 00107 00108 }

The following describes the implementation of this method:

1. The "createLookupKeys" method implements the first phase of the correlation algorithm, which is to generate lookup keys based on the event. This method accepts input parameters of a database connection object, which holds information to connect and access the database, an XMLEvent object, which holds the current event being processed and an output list of lookup keys which are created by this method based on data in the current event. The database connection object can be used to make additional SQL calls to read or write data.

- On line 77, the XPathSearch class is used to look for fields in the event XML document. This class supports XPath based searches on the XMLEvent class and on XML DOM trees in general. Refer to Section 4.2.4.2 for details on the usage of this class.
- 3. On line 83, the helper function "getEventType" performs several XPath lookup calls on the event object to determine whether the event is a request or a reply event or an event that does not need to be correlated. MQSeries MQPUT events are determined to be REQUEST events, while MQSeries MQGET events are determined to be REPLY events.
- 4. Line 85 performs a lookup of the message id in the event using an XPath search. The variable "XPathConstants.MSGID" is a constant value set to the XPath expression "/Event/Technology/MQSeries/*/*Exit/MQMD/MsgId". This expression accesses the message id field in the MQMD of the MQPUT event. Similarly on line 90, the completion code for an MQGET event is looked up and unsuccessful events are not further processed. The XPath constant "XPathConstants.CORRELID" holds the correlation id in the MQGET event, which is extracted from the event.
- 5. Lines 102 and 103 set the message id for an MQPUT event or the correlation id of the MQGET event as the lookup key for that event. This lookup key is used by the Analyzer to insert into the database and for correlating with corresponding matching events.

The XPathConstants.java file is shipped with TransactionVision and can be used for obtaining XPath values for some commonly used event fields.

• The following code fragment implements the "correlateEvents" method.

```
00110
          /**
           * Generates relations between two events if there is one.
00111
           * @param conInfo The database connection setup
00112
           * @param id
                                    Current event id.
00113
00114
           * @param idToMatch
                                    Event id needs to be matched with
current event.
           * @param eventRelations A list of relations between the two
00115
events.
00116
           * @throws AnalyzeEventException Pass along caught exception
to the caller.
00117
           */
          public void correlateEvents (ConnectionInfo conInfo,
00118
TechEventID id,
00119
                  TechEventID idToMatch, List eventRelations) throws
AnalyzeEventException {
00120
00121
              try
                  /* Retrieve data relevant for event correlation. */
00122
00123
                  Cache cache =
AnalysisCacheManager.instance().getCorrelationCache(conInfo.schema);
00124
00125
                  MQCorrelationData data = (MQCorrelationData)
cache.get(id);
00126
00127
                  if (data == null) {
00128
                      data =
MQRelationDBService.instance(conInfo.schema).getCorrelationData(
```

Chapter 3 • Tutorial - Extending the Analyzer How to perform custom correlation of related events

```
00129
conInfo.con, id);
00130
                      if (data != null)
00131
                           cache.insert(id, data);
00132
                      else
00133
                          return;
                   }
00134
00135
00136
                  MQCorrelationData dataToMatch = (MQCorrelationData)
cache.get(idToMatch);
00137
00138
                  if (dataToMatch == null) {
00139
                      dataToMatch =
MQRelationDBService.instance(conInfo.schema).getCorrelationData(
00140
conInfo.con, idToMatch);
00141
                      if (dataToMatch != null)
00142
                           cache.insert(idToMatch, dataToMatch);
00143
                      else
00144
                           return;
                   }
00145
00146
                  int apiId = data.apiCode;
00147
00148
                  int apiIdToMatch = dataToMatch.apiCode;
00149
00150
                   if (apiId != apiIdToMatch) {
00151
                      EventRelation eventRelation = new EventRelation();
00152
00153
                      eventRelation.relation =
EventRelation.MESSAGE_PATH | EventRelation.BIDIRECTION;
00154
                      eventRelation.direction =
00155
EventRelation.RELATION UNKNOWN;
                      eventRelation.confidence =
00156
EventRelation.STRONG_RELATION;
00157
00158
                      eventRelations.add(eventRelation);
                   }
00159
00160
              }
              catch (DataManagerException ex) {
00161
00162
                  throw new AnalyzeEventException(ex);
              }
00163
          }
00164
00165 }
```

The following describes the implementation of this method:

- 1. The "correlateEvents" method implements the second phase of the correlation algorithm, which is to validate the event pair matches by performing further checks on the pair of events. This method accepts input parameters of a database connection, the current event id, the event id to match against and an output parameter of containing a list of event relationships.
- 2. Line 123 retrieves a reference to the correlation data cache. This cache keeps event data fields available in memory for access using the event id. Line 125 then gets the correlation data from the cache using the event id. Similarly, line 136 gets data from the cache for the event to be matched against.
- 3. If the data is not present in the memory cache, it needs to be retrieved from the database. The method "getCorrelationData" on lines 128 and 139 does that. If the

data is retrieved from the database, it is inserted into the cache on lines 131 and 142.

- 4. Line 150 performs a check that two events of the same API are not correlated together.
- 5. Lines 150-158 involve creating an EventRelation object and adding it to the method output list "eventRelations". When this list is returned to the Analyzer, the Analyzer creates database entries relating these two events.

Note that each class and method referenced above is described further in Section 3.5.3.

While implementing this bean, care needs to be taken that the custom lookup key created is unique and not repeated between transactions. Not doing so will result in multiple, unrelated transactions correlated into one business transaction. This could also slow down the Analyzer performance.

3.7.3.4. Step 4: Enable the Analyzer to invoke the correlation bean

This involves editing the Beans.xml file to add the correlation bean. The following line in **old** needs to be added in the Beans.xml file, with the line in bold to be replaced with the correlation bean class being implemented:

3.7.3.5. Step 5: Test the correlation bean

The correlation bean can be verified by checking the transaction path in the transaction analysis view. A completely correlated path will have message path flows between local transactions.

4. Reference - Extending the Analyzer

4.1. Using the Beans.xml File

The file Beans.xml located in the <TVISION_HOME>/config/services directory controls the beans loaded by the Analyzer framework for event processing.

IMPORTANT: This file is used by the Analyzer internally. Modifying sections that are not documented here could break the correct functioning of the Analyzer.

Each module listed in the Beans.xml file has a type and a name. The type can be a "Context", which can hold other modules or a "Bean" type, which is loaded by a "Context". A module of type "Bean" contains the class that implements an interface which is used by its context. Each context defines a known interface for the beans it contains, loads the bean and calls into the interface implemented by the bean to perform its function. In the example segment below, the EventModifierCtx is a bean context which holds the DefaultModifierBean bean.

```
<Module type="Context" name=" EventModifierCtx">
```

```
<Module type="Bean" class="com.bristol.tvision.services.analysis.
eventmodifier.DefaultModifierBean"/>
```

</Module>

Each context uses its own rules to determine how its beans are invoked. The following contexts can be modified or added to:

- EventModifierCtx
- DBWriteExitCtx
- CorrelationTechHelperCtx

The following sections will document how each of the above contexts can be modified.

4.2. Unmarshalling Message Data

Typically, binary message data has a proprietary, user-defined format. The EventModifierCtx context allows a user to add a bean to "unmarshal" this binary data; that is, convert the binary data to XML for later use by TransactionVision in reports, for analysis or querying. To help converting binary data to XML, TransactionVision provides a set of utility classes.

4.2.1. The Default Modifier Bean

The TransactionVision installation comes with a default event modifier bean, the "com.bristol.tvision.services.analysis.eventmodifier.DefaultModifierBean". This bean scans the user data for any XML data and, if found, simply adds it to the Event XML document at the position "/Event/UserData/Chunk[@seqNo='n']" wher 'n' is the number of the data range (defined in the data collection filter).

4.2.2. Adding a Message Data Unmarshal Bean

Adding a custom message or user data unmarshal bean involves modifying the Beans.xml file to replace the default class with one or more custom written classes.

```
<Module type="Context" name="EventModifierCtx">
```

```
<Module type="Bean"
class="com.bristol.tvision.demo.stock.StockTradePayloadProc
essingBean"/>
```

```
</Module>
```

For example, in the code snippet above, a bean

com.bristol.tvision.demo.stock.StockTradePayloadProcessingBean processes any "stock trade" related custom data. If no event modifier bean is plugged in, the binary data will be saved into tables as a BLOB. The bean invoked by the EventModifierCtx context needs to implement the IEventModifier interface.

4.2.3. IEventModifier Interface

The method modify() of the IEventModifier interface is invoked by the EventModifierCtx context when it receives an event. This interface contains one method modify() defined as below.

public void modify(XMLEvent event)
 throws EventModifyException

Description:

The method modify() is called to convert the BLOB set stored in the XMLEvent object into the user-data section of the XML tree or modify the event's XML data. The BLOB set contains the event's binary message data.

IMPORTANT: data should typically be added in the XML event tree. Removing certain nodes from the tree could break the analysis and database write operations in later contexts. Parameters:

event - The XML event to which the XML format of the message data is appended to. The XMLEvent class is documented in detail in Section 4.2.4.

Throws:

EventModifyException – This exception represents a failure in the bean performing the XMLEvent modification.

4.2.4. XML Related Classes

This section documents the relevant public methods of the classes XMLEvent, XPathSearch and XMLParser. Class XMLEvent contains the incoming event converted to an XML DOM tree. Class XPathSearch is a utility class to search a DOM tree using XPath queries. Class XMLParser is a wrapper class around the Apache DOM parser, with better error handling facilities. TransactionVision event and event collection filter information is saved in XML document format. To retrieve values of different fields, an XPath expression is used to specify the location of the field. TransactionVision provides the file XPathConstants.java, which contains XPath expression constants used to locate different fields in the event. This file is useful for writing plug-in beans and reports and can be found at <TVISION_HOME>/java/src.

4.2.4.1. Class XMLEvent

package com.bristol.tvision.services.analysis
public class XMLEvent
extends com.bristol.tvision.util.xml.XMLDocument
implements java.io.Serializable

The class XMLEvent contains event data in XML DOM representation. It also holds a set of cached properties to carry inter-module communication information, and a list of BLOBs to hold application data which cannot be placed in the XML DOM tree. Note, that all the public methods of the class org.w3c.dom.Document are available to users of XMLEvent. The following methods are defined in the XMLEvent class.

Methods:

- getAttribute public java.lang.Object getAttribute(java.lang.String key)
- setAttribute

removeAttribute

public java.lang.Object **removeAttribute**(java.lang.String key) The above three methods allow the user to set a cached value at one stage of event processing, which can be used at another point of event processing without parsing the XML document. For example during the unmarshal message data phase values can be stored which may later be used during analysis. Typically, the key would be an XPath into the XML document and the value would be the XML element value. The user of the above APIs must ensure that TransactionVision internal values are not overwritten or deleted. This can be done by using unique XPaths to message data as the key.

getBlobCount

public int getBlobCount()

Returns the number of BLOBs available, using the blobIterator() method.

• blobIterator

public java.util.Iterator blobIterator()

Typically, event message data is stored into one BLOB field in the XMLEvent object. However, if data ranges are used in the data collection filter an array of BLOBs is created, one BLOB for each data range. This method returns an Iterator for instances of type XMLEvent.Blob.

deleteBlob

This method is used to delete the binary message data from XMLEvent. This method should typically be called if an EventModifier plugin bean converts binary data to XML. In that case, the binary data may no longer be required to be stored in the database and should be deleted using this method. If the message data is unmarshalled into the technology tree under, for example, the /Event/Technology/MQSeries/MQPUTEntry/Buffer subtree, the deleteUserDataRef and deleteDataChunk flags should be set to true. If the message data is unmarshalled into /Event/Data/Chunk, then both flags should be set to false. Also, if you want to replace a chunk with a different BLOB, call this method with both flags set to false and then call addBlob() to add a new BLOL into the XMLEvent.

Parameters:

seqNo - 0-based BLOB index

```
deleteUserDataRef - true if /Event//UserDataRef[@chunk=n] should be removed
deleteDataChunk - true if /Event/Data/Chunk[@seqNo=n] should be removed
```

getPiild
public int getPiiId()

getEventSeqNo

public int getEventSeqNo()

The PiiId (Program Instance Id) and the SeqNo (Sequence Number) together form a unique identifier to an event. They may be used to access event data from database tables.

Inner Class XMLEvent.Blob

Instances of this class are returned by the method 'blobIterator()' and represent the data ranges for the message data:

```
public static class Blob {
      public int id;
                            // id of the blob, starting with 0
      public int from;
                            // data range start
                            // data range end
      public int to;
     public int type;
                             // type of BLOB data
                             // (Binary, String, or XML,
                             // defined in TVisionCommon.java)
      public int ccsid;
                            // the character set id
      public byte[] blob;
                             // the data
      public Blob(int ID, int from, int to, int type, int ccsid,
           byte[] blob);
}
```

4.2.4.2. Class XPathSearch

```
package com.bristol.tvision.util.xml
    public class XPathSearch
    extends XPathSearchBase
The helper class XPathSearch allows access to elements of an XML document using the
XPath syntax.
```

Constructor:

• XPathSearch

XPathSearch(org.w3c.dom.Document doc) Creates an XPathSearch object from a DOM document or derived class like XMLEvent. XPathSearch

XpathSearch(java.io.InputStream stream) throws XMLException Creates an XPathSearch object from an InputStream. The InputStream is parsed into a DOM document without validation

• XPathSearch

XPathSearch(java.io.InputStream stream, boolean validate)
throws XMLException

Creates an XPathSearch object from an InputStream. The InputStream is parsed into a DOM document. Parameters: stream - The InputStream conatining the XML data

validate - Use parser validation

XPathSearch

XPathSearch(java.io.InputStream stream, String schemaFile) throws XMLException

Creates an XPathSearchobject from an InputStream.

The InputStream is parsed into a DOM document with validation turned on.

Parameters:

stream - The <CODE>InputStream</CODE> conatining the XML data

schemaFile - The xml schema file to use for validation

Methods:

• getNodes

This method returns a list of all nodes in the XML document matching the XPath query. The elements in the array are ordered according to the order of the elements in the DOM tree.

Overrides:

getNodes in class XPathSearchBase

Parameters:

 ${\tt xpath}$ - The XPath expression for the query

Returns:

A list of all nodes matching the query

Throws:

XMLException - Signals error during retrieving the values from the document

getValues

This method returns the value of all text elements in the XML document matching the XPath query. The elements in the array are ordered according to the order of the elements in the DOM tree.

Overrides:

getValues in class XPathSearchBase

Parameters:

 ${\tt xpath}$ - The XPath expression for the query

Returns:

The value of all text elements matching the query

Throws:

 $\tt XMLException$ - Signals error during retrieving the values from the document

getValue

public java.lang.String getValue(java.lang.String xpath)

throws XMLException

This method returns the value of the first text element in the XML document matching the XPath query.

Overrides:

getValue in class XPathSearchBase

Parameters:

xpath - The XPath expression for the query

Returns:

The value of the first matching text element

Throws:

XMLException - Signals error during retrieving the values from the document

4.2.4.3. Class XMLParser

package com.bristol.tvision.util.xml public class XMLParser

implements org.xml.sax.ErrorHandler

This class is a wrapper around the Apache DOM parser and is a utility useful to parse XML files or convert binary streams containing XML data into a DOM tree.

Constructor:

XMLParser

```
XMLParser(boolean validation)
Creates a parser instance
Parameters:
```

validation - whether to create a validating parser or not

Methods:

• parse

public org.w3c.dom.Document parse(java.lang.String systemId)

throws XMLException

Parses a XML file Parameters: systemId - The system id for the XML source Returns: The parsed document as a DOM tree Throws: XMLException - Signals errors during parsing

parse public org.w3c.dom.Document parse(java.lang.String systemId, java.lang.String schema) throws XMLException Parses a XML file and uses the specified XML schema rather than a schema reference in the document itself for schema validation Parameters: systemId - The system id for the XML source schema - The schema to use for validation Returns: The parsed document as a DOM tree Throws: XMLException - Signals errors during parsing parse public org.w3c.dom.Document **parse**(java.io.InputStream stream) throws XMLException

Parses a XML document from an input stream

Parameters:

stream - The input stream for the document Returns: The parsed document as a DOM tree Throws: XMLException - Signals an error during parsing

parse

public org.w3c.dom.Document **parse**(java.io.InputStream stream, java.lang.String schema) throws XMLException

Parses a XML document from an input stream and uses the specified XML schema rather than a schema reference in the document itself for schema validation Parameters: stream - The input stream for the document schema - The schema to use for validation Returns: The parsed document as a DOM tree Throws:

XMLException - Signals an error during parsing

4.2.4.4. Other Utility Classes

Often, binary structures embedded in the message data will need to be converted to XML. This can be accomplished with a two step process, first extract the binary data into Java data types and then convert these data types to appropriate XML elements. The Java class java.io.DataInputStream could be used to walk through a binary stream, extract and convert data into Java basic types. Also, the class "Translator" can be used to convert raw binary data into a Java UTF String with code page conversion: package com.bristol.tvision.util.charmapper

public class Translator {

```
public static Translator instance(int srcCcsid);
public String translate(byte[] rawData);
```

Once Java basic types have been extracted from the binary stream these values need to be converted to XML data. This can be done using the utility "XML builder" classes in the package com.bristol.tvision.util.xml. These classes allow a user to set values of native Java types, a element name and get the XML tag output appended to a DOM tree using the toDOM() method. These classes implement the DOMElement interface.

Interface DOMElement

public interface DOMElement

This class defines a common interface for classes which output XML into a DOM tree.

Methods:

}

• toDOM

This method appends nodes to the DOM tree doc at node location root.

Class EventElement

public abstract class EventElement
implements DOMElement

This class is the super class of all XML builder classes that output XML elements into a DOM tree.

Methods:

Constructor

public EventElement(java.lang.String name)

The constructor of the EventElement class takes in the element name as a parameter. The element name is used by the toDOM method to output the node of element name to the XML DOM tree.

• toDOM

This is the same method as in the interface DOMElement.

• Class TextElement

public abstract class TextElement
extends EventElement

This class is a super class for those XML element classes which have only one text node as a child. This class allows adding attributes to the XML element.

Methods:

Constructor

public TextElement(java.lang.String elementName)
The constructor takes in the element name of the node to be inserted into the XML DOM tree.

toDOM
 public void toDOM(org.w3c.dom.Document doc,

org.w3c.dom.Node root) Overrides:

toDOM in class EventElement

This method allows adding a name-value pair of attributes to the XML element.

hasNonNullValue

public abstract boolean **hasNonNullValue**() This method returns true if this element has a non-null value and false otherwise.

Class ByteElement

public class ByteElement
extends TextElement

Fields:

• value

public byte **value**

This field holds the byte value to be converted to an XML DOM tree node by the toDOM method.

Methods:

Constructor

public ByteElement(java.lang.String elementName)
The constructor takes in the element name of the tag to be output in the XML DOM tree node.

• toDOM

public void toDOM(org.w3c.dom.Document doc,

org.w3c.dom.Node root)

This method appends a node containing the byte value held by the field value to the DOM tree doc at node location root with the element name elementName specified in the constructor of this object.

• toString

public java.lang.String toString()

Overrides:

toString in class java.lang.Object

This method converts the byte held in the field value to a string representation.

hasNonNullValue

public boolean **hasNonNullValue**() Overrides: hasNonNullValue in class TextElement This method returns true if this element has a non-null value and false otherwise.

Class ByteStringElement

public class ByteStringElement
extends TextElement

Fields:

value

```
public byte[] value
```

This field holds the byte array value to be converted to an XML DOM tree node by the toDOM method.

Methods:

Constructor

public ByteStringElement(java.lang.String elementName)
The constructor takes in the element name of the tag to be output into the XML DOM tree node.

• toDOM

public void toDOM(org.w3c.dom.Document doc,

org.w3c.dom.Node root)

This method appends a node containing the byte array value held by value to the DOM tree doc at node location root with the element name elementName specified in the constructor of this object.

• toString

public java.lang.String toString()

Overrides:

toString in class java.lang.Object This method converts a byte array held in the value field to a string representation.

hasNonNullValue

```
public boolean hasNonNullValue()
```

Overrides:

hasNonNullValue in class TextElement

This method returns true if this element has a non-null value and false otherwise.

Class IntElement

public class IntElement
extends TextElement

Fields:

```
    value
```

```
public int value
```

This field holds the integer value to be converted to an XML DOM tree node by the toDOM method.

Methods:

Constructor

public IntElement(java.lang.String elementName)
The constructor takes in the element name of the tag to be output into the XML DOM tree node.

• toDOM

public void toDOM(org.w3c.dom.Document doc,

```
org.w3c.dom.Node root)
```

This method appends a node containing the integer value held by field value to the DOM tree doc at node location root with the element name elementName specified in the constructor of this object.

• toString

```
public java.lang.String toString()
```

Overrides:

toString in class java.lang.Object This method converts an integer to a string representation.

hasNonNullValue

```
public boolean hasNonNullValue()
```

Overrides:

hasNonNullValue in class TextElement

This method returns true if this element has a non-null value and false otherwise.

Class IntHexElement

public class IntHexElement

extends IntElement

This class's toDOM method outputs an integer value to a XML DOM node element as a hexadecimal string.

```
    Class LongElement
```

```
public class LongElement
extends TextElement
```

Fields:

value

public long **value**

This field holds the integer long value to be converted to an XML DOM tree node by the toDOM method.

Methods:

Constructor

public LongElement(java.lang.String elementName)
The constructor takes in the element name of the tag to be output into the XML DOM tree node.

• toDOM

public void toDOM(org.w3c.dom.Document doc,

org.w3c.dom.Node root)

This method appends a node containing the integer long value held by the field value to the DOM tree doc at node location root with the element name elementName specified in the constructor of this object.

```
• toString
```

```
public java.lang.String toString()
```

Overrides:

toString in class java.lang.Object This method converts the integer long held in the field value to a string representation.

hasNonNullValue

public boolean hasNonNullValue()

Overrides:

hasNonNullValue in class TextElement

This method returns true if this element has a non-null value and false otherwise.

• Class LongHexElement

public class LongHexElement

extends LongElement

This class's toDOM method outputs an integer long value to a XML DOM node element as a hexadecimal string.

Class StringElement

public class StringElement
extends TextElement

Fields:

value

public String value

This field holds the String value to be converted to an XML DOM tree node by the toDOM method.

Methods:

Constructor

public StringElement(java.lang.String elementName)

The constructor takes in the element name of the tag to be output into the XML DOM tree node.

toDOM

public void toDOM(org.w3c.dom.Document doc,

org.w3c.dom.Node root)

This method appends a node containing the String value held by the field value to the DOM tree doc at node location root with the element name elementName specified in the constructor of this object.

• toString

```
public java.lang.String toString()
```

Overrides:

toString in class java.lang.Object This method converts the String held in the field value to a string representation.

hasNonNullValue

public boolean hasNonNullValue()

Overrides:

hasNonNullValue in class TextElement This method returns true if this element has a non-null value and false otherwise.

Class RawStringElement

public class RawStringElement

extends TextElement

This class's toDOM method outputs a String value to a XML DOM node element as a string whose non-ASCII characters are converted to hexadecimal values.

4.2.5. Sample Usage of the IEventModifier Interface

Refer to the code in the directory <TVISION_HOME>/samples/stock/beans to see the usage of the IEventModifier interface.

Two Java beans have been developed for processing stock trade simulation under TransactionVision 3.0 - StockTradePayloadProcessingBean and StockTradeAnalysisBean. StockTradePayloadProcesingBean is a message data processing bean. It looks for the MQPUT/MQPUT1 and MQGET events from a StockTrade program (which initiates a trade). MQPUT/MQPUT1 and MQGET events mark the beginning and end of a stock trade transaction. For the MQPUT/MQPUT1 calls, the bean retrieves the message data blob, passes it to the XML parser, and attaches the resultant DOM tree (/Order) to the event document under /Event/Data. For the MQGET calls, the bean performs the same XML parsing on the message data blob, creates a new XML document (Event/Data/Result) reflecting the trade result.

The following code fragment is the change to the Beans.xml file. It tells the Analyzer framework to load the StockTradePayloadProcessingBean bean as a part of the EventModifierCtx context. The payload processing bean's IEventModifierCtx interface's modify() method is invoked by the context.

```
<Module type="Context" name="EventModifierCtx">
<!--
This context contains beans that modify XML event, which is
unmarshalled from raw event stream.
-->
```

```
<Module type="Bean"
class="com.bristol.tvision.demo.stock.StockTradePayloadProcessingBean"/
>
```

```
</Module>
```

The following code fragment is the implementation of the IEventModifierCtx interface. The class StockTradePayloadProcessingBean is derived from EventModifierBean and needs to implement the method modify(). The method modify() receives an XMLEvent as its parameter. This sample looks for a particular type of event namely MQPUTs, MQPUT1s and MQGETs APIs from certain programs.

```
public class StockTradePayloadProcessingBean extends
EventModifierBean {
/** Creates new StockTradePayloadProcessingBean */
    public StockTradePayloadProcessingBean() {
    }
    /**
     * Processing the event data and generate stock trade payload
document
     * @param event completed event document for the current event
     * @throws EventModifyException when processing failed
     */
     public void modify(XMLEvent event) throws EventModifyException
{
      try {
            XPathSearch lookup = new XPathSearch(event);
            int type = StockTradeHelper.getEventType(lookup);
            switch (type) {
                    case StockTradeHelper.MQSERIES_REQUEST_EVENT:
                    case StockTradeHelper.MQSERIES_REPLY_EVENT:
                            processMQSeriesEvent(event, lookup,
type);
                            break;
                    case StockTradeHelper.DONT CARE EVENT:
                    default:
                            break;
                    }
            }
            catch (XMLException e) {
                    if (Logging.debug)
EventReader.log.debug("StockTradePayloadProcessingBean-process: " +
                            "XML exception encountered");
            }
    }
}
```

The method getEventType() in the file StockTradeHelper.java does a lookup on certain parts of the XML event document using the class XPathSearch. For example, in the segment below the value of the event technology name and the program name is being accessed from the DOM tree using the class XPathSearch's getValue method. XPathConstants.TECH_NAME and XPathConstants.PROGRAM_NAME map to the XPath expressions

"/Event/StdHeader/TechName" and "/Event/StdHeader/ProgramName" respectively. Refer to the event XML schema at <TVISION_HOME>/config/xmlschema/Event.xsd for the schema layout of the XML event packet.

```
XPathSearch lookup = new XPathSearch(event);
String techName = lookup.getValue(XPathConstants.TECH_NAME);
if (techName.equalsIgnoreCase(TVisionCommon.TECH_NAME_MQSERIES)) {
    /* we are only interested in the initiating program events */
```

Once the right event has been identified, its message data is converted to XML by the getUserDataXML() method called from the processMQSeriesEvent method. The BLOB list is first obtained from the XMLEvent object using the blobIterator() method. The obtained BLOB is converted to an XML DOM tree using the XMLParser class. The method getUserDataXML is as below.

```
/**
 * Return the XML document for event user data blob
 * @param event TransactionVision event document
 * @return the byte array representing the event user data
 */
 public static Document getUserDataXML(XMLEvent event) throws
XMLException {
```

```
byte[] blob;
Iterator blobs = event.blobIterator();
if (blobs.hasNext())
            blob = ((XMLEvent.Blob) blobs.next()).blob;
else
            return null;
XMLParser parser = new XMLParser(false);
return parser.parse(new ByteArrayInputStream(blob));
```

Note that the parse() method of the XMLParser class will throw an exception if the BLOB is not a XML document. Once the XML document is obtained, the document tree is inserted under the node /Event/Data of the XMLEvent DOM tree. The method getDataNode() as below returns the location of the message data node in the event DOM tree.

```
/**
    * Return the /Event/Data node in the event document
    * @param lookup The XPathSearch lookup object of the
corresponding event document.
    * @return the /Event/Data node in the event document
    */
    public static Node getDataNode(XPathSearch lookup) throws
XMLException {
        NodeList nodes =
        NodeList nodes
        NodeList nodes
        NodeList nodes
        NodeList nodes
        NodeList nodeList nodes
        NodeList no
```

Once the binary data has been converted to an XML tree and the message data node has been identified, the next step is inserting the message data XML tree into the XML event. The code below from file StockTradePayloadProcessingBean.java shows how to append data into the XMLEvent DOM tree. Here payloadDoc is of type org.w3c.dom.Document. A call into getDocumentElement returns an object of type

}

org.w3c.dom.Element which is copied into the XMLEvent object event. The call appendChild() to attaches the copied nodes to the location for the message data, namely under "/Event/Data".

```
Document payloadDoc = StockTradeHelper.getUserDataXML(event);
    /* append the Order document to /Event/Data */
    dataNode.appendChild(event.importNode(payloadDoc.getDocumentEleme
nt(), true));
```

The above sample code is useful when the message data is already in an XML format. The following sample from file StockTradePayloadProcessingBean. java shows how to create nodes from part of the message data tree and append it to the XMLEvent tree.

```
/* create /Event/Data/Result/ID */
String orderid = payloadLookup.getValue("/Order/ID");
Element eltOrderID = event.createElement("ID");
eltOrderID.appendChild(event.createTextNode(orderid));
/* create /Event/Data/Result/Type */
String orderType = payloadLookup.getValue("/Order/Type");
Element eltType = event.createElement("Type");
eltType.appendChild(event.createTextNode(orderType));
/* create /Event/Data/Result/Status */
String orderStatus = payloadLookup.getValue("/Order/Status");
Element eltStatus = event.createElement("Status");
eltStatus.appendChild(event.createTextNode(orderStatus));
/* create /Event/Data/Result */
Element eltRes = event.createElement("Result");
/* attach ID, Type, Status to /Event/Data/Result */
eltRes.appendChild(eltOrderID);
eltRes.appendChild(eltType);
eltRes.appendChild(eltStatus);
/* attach new tree to /Event/Data */
dataNode.appendChild(eltRes);
```

4.3. Trimming Data From an Event

The DBWriteCtx context is invoked by the Analyzer framework before the database write operation. It gives a user defined bean an opportunity to trim out data from the XML event packet. Beans loaded by this context need to implement the IDBWriteExit interface.

4.3.1. Interface IDBWriteExit

public interface IDBWriteExit

Methods

```
    modify
```

public XMLEvent modify(XMLEvent event)

throws DBWriteExitException

This method trims data off the XML event. The bean has to make a copy of the XML event and return the trimmed copy.

Parameters:

event - The XML event to trim.
Returns:
The return value is the trimmed XML event
Throws:
TrimEventDataException - Trimming of the event failed
The sample code under <TVISION_HOME>/samples/dbwritexit shows how to write a
bean to plug into the database write exit context.

4.4. XML-Database mapping Using XDM Files

The TransactionVision database schema is made extensible through the XML to Database Mapping (XDM) files. As new technologies or message data specific information is added, new XDM files can be written to describe the lookup tables for the technology and message-specific data in those events. Hence the purpose of the XML to Database mapping is twofold:

- To describe which fields are to be extracted from the XML event data and stored in lookup tables for fast searching and retrieval.
- To make the database schema partially data-driven.

The definitions contained in the XML Database Mapping (XDM) file are used as input not only to the TransactionVision Data Manager (including the query services), but also to a program that generates the commands necessary to create the lookup tables. The XDM mappings can be technology or platform specific. The common mapping defined

in the file <TVISION_HOME>/config/xdm/Event.xdm (data in the standard event header) will be written for every event, but the mappings defined in the other XDM files will only be applied if the current event matches the mapping's "category" (technology or platform) definition. The XML schema format of XDM files is defined in

```
<TVISION_HOME>/config/xmlschema/XDM.xsd. The following code is an extract from the file Event.xdm.
```

```
<?xml version="1.0"?>
<Mapping documentTable="event" documentColumn="event_data">
        <Key name="proginst_id" type="INTEGER"
description="ProgramInstanceId">
                <Path>/Event/EventID/@programInstID</Path>
        </Kev>
        <Key name="sequence no" type="INTEGER"
description="SequenceNumber">
                <Path>/Event/EventID/@sequenceNum</Path>
        </Key>
        <Table name="EVENT_LOOKUP" category="COMMON">
            <Column name="host_id" type="INTEGER" description="Host"
isObject="true">
                <Path>/Event/StdHeader/Host/@objectId</Path>
            </Column>
            <Column name="program_id" type="INTEGER"
description="Program" isObject="true">
                <Path>/Event/StdHeader/ProgramName/@objectId</Path>
            </Column>
         . . .
        </Table>
</Mapping>
```

The above snippet from Event.xdm defines a table EVENT containing the XML document and a table EVENT_LOOKUP, containing various indexed columns of data from the XML document. The key columns proginst_id and sequence_no are integer types and mapped to XPath expressions /Event/EventID/@programInstID and

/Event/EventID/@sequenceNum. These key columns are primary keys common to the EVENT and EVENT_LOOKUP tables. Similarly columns host_id and program_id are mapped to XPath expressions /Event/StdHeader/Host/@objectId and

/Event/StdHeader/ProgramName/@objectId respectively.

The above XDM file specifies that when an XML event is written to the database by the DBWrite module in the Analyzer, these fields are extracted and written into the database columns mapped to in the XDM file. Similarly, when the database is queried using the QueryServices XML interface, these XDM files are used to construct the corresponding SQL query.

The isObject attribute for a Column tag in the above XDM file refers to that column being an identifier for an object in the system model table. The documentTable and documentColumn tags are the table and column where the actual XML document is stored. The key is the primary key and is common to the document table and the lookup tables. Each lookup column is indexed. The queryOnly attribute for a Column tag indicates that the value is not written by the Analyzer in the DBWrite module, but maybe written in the analysis phase of the Analyzer or by some other application. Hence, this field is for queries only.

```
</Column>
```

The generated attribute for a Column tag means that column is a database generated id.

```
<Column name="sequential_id" type="INTEGER" generated="true" des
cription="SequentialId">
```

```
<Path>/Event/SequentialId</Path>
```

</Column>

The conversionType attribute for a Column tag means that field requires a formatting conversion before writing to the database. The TypeConvService is called into before writing that field into the database. This is typically used for writing date/time or enumeration fields.

```
<Column name="entrytime" type="CHAR" size="20" description="Entr
yTime" conversionType="Date">
```

```
<Path>/Event/StdHeader/EntryTime</Path> </Column>
```

The category attribute on the Table tag contains either COMMON or the technology string or the platform string for the event data that should be written into this table. The string COMMON indicates that this table contains data common to every event and should be written for every event going through the Analyzer. A technology or platform name like "MQSERIES" or "OS390_BATCH" used in the category field indicates that this table should only be filled for events of that technology or platform.

```
<Table name="EVENT_LOOKUP" category="COMMON">
...
</Table>
<Table name="OS390_LOOKUP"
category="OS390_BATCH,OS390_CICS,OS390_IMS">
...
```

</Table>

A column can map to multiple XPath expressions as in the sample code below. This assumes that only one of the XPaths will exist in a given event document.

<Column name="datasize" type="INTEGER" description="DataSize">

<Path>/Event/Technology/MQSeries/MQGET/MQGETExit/DataLength</Path >

```
<Path>/Event/Technology/MQSeries/MQPUT/MQPUTExit/BufferLength</Path>
```

<Path>/Event/Technology/MQSeries/MQPUT1/MQPUT1Exit/BufferLength</ Path>

</Column>

The XDM files are also used by the CreateSqlScript utility to create the TransactionVision tables are setup time.

4.5. Performing Event Analysis

There are five categories of event analysis activities defined in TransactionVision:

- **Event Correlation:** Establishing relation(s) between any two events. Examples include message path relation representing a message flow from one event to another, and transaction path relation representing a control flow between the two events.
- Local Transaction Analysis: Grouping events of the same technology that participate in the same unit of work in the same thread of execution into one local transaction object.
- **Business Transaction Analysis**: Grouping local transaction objects participating in the processing of the same business activity instance into one business transaction object. This is achieved by establishing relation between any two local transaction objects through the corresponding message path or transaction path relation of respective events in the local transaction objects.
- Statistics Analysis: Calculating event statistics for the Static Topology View
- User Analysis: This can be any customized infrastructure or business level analysis. Each event analysis task is implemented in an event analysis bean. The class AnalyzeEventBean defines the base class for these beans:

The individual beans are managed under a multi-level analyze event context framework. The class AnalyzeEventCtx defines the top level context. The set of beans to be managed under this context are specified in the Beans.xml file. Each registered bean is executed following the order defined in the file. The following is an example of the event analysis context setup for the stock trade simulation example:

```
<Module type="Context" name="AnalyzeEventCtx">
```

```
TransactionVision Programmer's Guide
```

```
class="com.bristol.tvision.services.analysis.eventanalysis.Local
    TransactionAnalysisBean"/>
  <!-- TransactionVision Default Business Transaction Analysis bean
-->
  <Module type="Bean"
  class="com.bristol.tvision.services.analysis.eventanalysis.Busine
  ssTransactionAnalysisBean">
  <!-- TransactionVision Statistics beans -->
  <Module type="Context" name="StatisticsCtx"
  class="com.bristol.tvision.services.analysis.statistics.Statistic
  sCtx">
     <Module type="Bean"
  class="com.bristol.tvision.services.analysis.statistics.MQStatist
  icsBean"/>
     <Module type="Bean"
  class="com.bristol.tvision.services.analysis.statistics.JMSStatis
  ticsBean"/>
     <Module type="Bean"
  class="com.bristol.tvision.services.analysis.statistics.ServletSt
  atisticsBean"/>
     <Module type="Bean"
  class="com.bristol.tvision.services.analysis.statistics.EJBStatis
  ticsBean"/>
  </Module>
  <Module type="Bean"
    class="com.bristol.tvision.demo.stock.StockTradeAnalysisBean"/>
</Module>
```

4.5.1. Event Analysis Utility Classes and Interface

The following utility classes are extensively used in implementing various types of event analysis beans.

4.5.1.1. Interface Cache

```
package com.bristol.tvision.util.cache
```

public interface Cache

TransactionVision maintains various in-memory caches for miscellaneous objects. These caches are implemented as LRU caches, meaning that always the most recent processed data is available. For example, a local transaction cache is maintained to store a mapping from event ID to local transaction data. This interface defines the methods for manipulating the cache.

Methods:

```
• insert
```

```
public void insert(java.lang.Object key, java.lang.Object value)
```

Insert a new key-value pair into the cache. Parameters:

key – new cache object key field value – new cache object value field

• get

public Object get(java.lang.Object key)

This method returns the value field of the cache entry with the matching key. Parameters:

```
key – key field of the cache entry to be matchedReturns:The value field of the cache entry if a matching object is found.
```

remove
 public void remove(java.lang.Object key)

Remove the cache entry with the matching key.

Parameters:

 $\operatorname{key}-\operatorname{key}$ field of the cache entry to be matched

removeAll
 public void removeAll()

Remove all cache entries.

```
• getSize public int getSize()
```

Return the defined cache size specified in the CacheProperty file. Returns: The defined cache size

```
    resize
        public void resize(int size)
Resizes (and clears) the cache.
Parameters:
size – new cache size
```

getElementCount

public int getElementCount()
Return the current number of cache entries.
Returns:
The current number of cache entries.

getCacheName

public java.lang.String getCacheName()
Return the name of this cache.
Returns:
The name of this cache

4.5.1.2. Class ConnectionInfo

package com.bristol.tvision.datamgr

public class ConnectionInfo

This class is a simple structure for holding the TransactionVision database connection and schema name within an object which can be passed through the event analysis service framework.

Fields:

```
con
public java.sql.Connection con;
```

A TransactionVision Connection object to the database. This connection object implements the Java SQL Connection object interface.

schema
public java.lang.String schema;

String for the current project database schema.

4.5.1.3. Class EventID

package com.bristol.tvision.datamgr.dbtypes

public class EventID

Each event is uniquely identified by a pair of integer ID: a program instance (PII) ID and a sequence number. The program instance ID points to the program instance (threads, tasks, etc.) the event occurs within. This class defines a wrapper around these two identifiers for an event.

Constructor:

• EventID

```
EventID(int piiId, int seqNo)
Creates an event ID object for an event with the program instance ID piiId and sequence
number seqNo.
```

Fields:

• public int piild

The program instance id for this event

• public int seqNo

The sequence number of this event

Methods:

• equals public boolean equals(EventID eventId)

Determine if the input event is the same as this event. Parameters: eventId – eventId to be matched Returns: true if the event ID matches, false otherwise.

hashCode
 public int hashCode()

Return a unique integer has code for this event ID object. Returns: The integer hash code for this event ID object

• toString public java.lang.String toString ()

Return a string describing this event ID object. Returns: A string describing this event ID object

4.5.1.4. Class TechEventID

package com.bristol.tvision.datamgr.dbtypes public class TechEventID This class extends class EventID and additionally holds the technology ID of the event.

Constructor:

TechEventID

TechEventID(int piiId, int seqNo, int techId) Creates an event ID object for an event with the program instance ID piiId, sequence number seqNo., and techology ID techId

Fields:

• public int techId

The techology ID for this event.

4.5.2. Event Analysis Classes

4.5.2.1. Interface IAnalyze

package com.bristol.tvision.services.analysis.eventanalysis public interface IAnalyze This defines the interface for general-purpose event analysis beans.

Methods:

```
    analyze
```

This method implements a specific event analysis task on the given event. Parameters:

 $\verb|conInfo-database| connection info object for the current project|$

event-completed XML document for the current event

```
Throws:
```

AnalyzeEventException - Signals errors during the event correlation analysis

4.5.2.2. Class AnalyzeEventCtx

package com.bristol.tvision.services.analysis.eventanalysis public class AnalyzeEventCtx extends ChainManagerCtx implements IAnalyze This is the top level event analysis context class and holds all analysis beans for the event analysis. During analysis, the **analyze()** interface will be called for all beans contained in this context (in sequential order).

4.5.2.3. Class AnalyzeEventBean

package com.bristol.tvision.services.analysis.eventanalysis public abstract class AnalyzeEventBean extends ChainManagedBean implements IAnalyze This is the abstract base class for all event analyze bean. Any custom event analysis bean should derive directly or indirectly from this class, and implement the IAnalyze interface methods.

Fields:

```
• Analysis Type
```

```
public static final int EVENT_CORRELATION = 1;
public static final int LOCAL_TRANSACTION_ANALYSIS = 2;
public static final int BUSINESS_TRANSACTION_ANALYSIS = 3;
public static final int BUSINESS_PROCESS_ANALYSIS = 4;
public static final int USER_ANALYSIS = 5;
The type of analysis implemented by the event analysis bean
instance.
```

Methods:

• getAnalysisType public int getAnalysisType()

Return the analysis type of the event analysis bean.

4.5.3. Adding Custom Correlation Analysis Beans

For event correlation, the class CorrelationTechHelperCtx defines the top-level context for managing all event correlation beans. These beans are managed into different groups according to the technology categories the beans are associated with. Each category is managed by a technology specific event correlation context. Each context is designated to

handle a particular type of technology (e.g.: WebSphere MQ). That is, all the events being passed to the context belong to the same technology. The technology specific context itself holds a set of correlation beans which implements the Interface IEventCorrelation, each is responsible for correlating the current technology to one particular other technology. In Addition to these technology specific contexts it is possible to plug in a custom 'UserCorrelationBean', which will be invoked for every event processed by the event analysis service, irrespectively of the technology.

The following is an example of event correlation context definition in the Beans.xml file:

```
<Module type="Context" name="CorrelationTechHelperCtx">
```

```
<!-- This context contains beans that perform event correlation.
-->
<!-- For each event the correlation context that matches the
event's technology will be called. -->
<!-- This context contains beans that perform MQSeries event
correlation -->
<Module type="Context" name="CorrelationMQHelperCtx"
class="com.bristol.tvision.services.analysis.eventanalysis.Correl
ationMQHelperCtx">
 <!-- This bean is provided by TransactionVision for establishing
 default intra MQSeries event correlation such as MQPUT - MQGET
 message path relations -->
 <Module type="Bean"
 class="com.bristol.tvision.services.analysis.eventanalysis.MQTOM
 QRelationshipBean"/>
 <!-- This bean is provided by TransactionVision for establishing
 MQSeries - IMSBridge message path relations -->
 <Module type="Bean"
 class="com.bristol.tvision.services.analysis.eventanalysis.MQToB
 ridgeRelationshipBean"/>
 <!-- This bean is developed specifically the stock trade
 simulation for establishing a custom transaction path relation
 between a failed MQGET call and the MQPUT call issued by the
 stock trade initiating program -->
 <Module type="Bean" class="com.bristol.tvision.demo.
 stock.StockTradeRelationshipBean"/>
 <!-The CorrelationTechHelperCtx provides a hook for the user to
 plug in a technology independent custom correlation bean:
 <!-- UserCorrelationBean :
 1.) the 'createLookupKeys()' method of the user bean is called
 after the default lookup key generation for events of all
 technologies and can add additional lookup keys
```

```
2.) the `correlateEvents()' method of the user bean is called
after the default correlation for events of all technologies and
can generate additional event relations -->
<Attribute name="UserCorrelationBean"
value="com.bristol.tvision.services.analysis.eventanalysis.UserC
orrelationBean"/ -->
```

</Module>

</Module>

For WebSphere MQ, TransactionVision provides a bean MQToMQRelationshipBean that handles all WebSphere MQ correlation tasks. This includes matching MQPUT or MQPUT1 calls to MQGET calls that handle the same message. The resultant relation is known as the message path relation, indicating a data flow between the two corresponding applications. It is possible to add additional correlation logic in several ways:

- A new correlation bean can be developed and added to the correlation processing chain. In the above example, the StockTradeRelationshipBean bean is invoked in the MQSeries event context along with the MQToMQRelationshipBean.
- The default correlation bean can be replaced by a user bean through subclassing or aggregation. This allows modifications to the default correlation behavior. For example, a bean can be developed that invokes the MQToMQRelationshipBean correlation interfaces, examines the correlation results, and makes modifications to the results if necessary.
- Provide an implementation for the UserCorrelationBean.

An event correlation bean should implement the interface **IEventCorrelation**. The IEventCorrelation interface defines two methods **createLookupKeys** and correlate for the two phases discussed before. The class **CorrelationTechHelperBean** serves as the base class for all event correlation beans

In TransactionVision, event correlation is performed on a per event, per technology basis. The correlation task is divided into two phases.

The first phase involves generating lookup keys based on the characteristics of the current event. The purpose of setting up these keys is to identify the set of events bearing the same lookup key as the potential candidates for correlation in the second phase. For example, in the case of MQPUT(1) – MQGET message path relation generation, for each MQPUT(1) and MQGET event, a key composed of the message ID (MQMD.MsgId), correlation ID (MQMD.CorrelId), message put data and time is generated.

For any event, the createLookupKeys() method of each bean contained in the technology specific context will be called. In the above example, for a MQ event the MQToMQRelationshipBean as well as the MSToBridge RelationshipBean will both generate a lookup key for the current event.

The second phase involves relation generation. Specifically, a set of events is passed as potential candidate for matching with the current event. This set is composed of the events that have the same lookup key as the current event. For example, for a MQGET event, all the MQPUT(1) /MQGET events having the same key (message Id + correlation ID + message put data + message put time) are passed as potential match candidates. Further tests can now be conducted on individual candidate event to see if it is truly related to the current

event. For example, events with the same method/API name (MQPUT-MQPUT, MQGET-MQGET) should not result in a message path relation.

For a certain set of candidates with matching lookup keys, the type of the correlation (e.g., MQ-MQ or MQ-IMS) determines which beans correlateEvents() method is called. In the above example, a set of events with matching lookup key of type MQ-MQ will be passed on to the MQToMQRelationshipBean, a set of events with type MQ-IMS will be passed on to the MQToBridgeRelationshipBean. Currently the following correlation types are defined for TransactionVision as constants in class EventCorrelationBean:

public class EventCorrelationBean extends AnalyzeEventBean {

```
public static final int MQ_PUT_GET_TYPE = 1;
public static final int MQ_IMSBRIDGE_TYPE = 2;
public static final int IMSBRIDGE_ENTRY_EXIT_TYPE = 3;
public static final int JMS_SEND_RCV_TYPE = 4;
public static final int PROXY_TYPE = 5;
public static final int PUBSUB_TYPE = 6;
public static final int CICS_TRANS_TYPE = 7;
public static final int MQ_CICS_TYPE = 8;
```

}

•••

The correlation type for a correlation bean has to provided in the constructor call. For user defined correlation beans, new correlation types should be ≥ 100 .

4.5.3.1. Interface IEventCorrelation

package com.bristol.tvision.util.services.analysis.eventanalysis
public interface IEventCorrelation

The IEventCorrelation interface defines the methods to be implemented by any event correlation bean.

Methods:

createLookupKeys

```
public void createLookupKeys(ConnectionInfo conInfo, XMLEvent event,
java.awt.List lookupKeys) throws AnalyzeEventException
```

Generate one or more lookup keys for correlation purpose for the given event.

Parameters:

conInfo – database connection info object for the current project event – completed XML document for the current event lookupKeys – list of lookup keys to be added

Throws:

AnalyzeEventException - Signals errors during the event correlation analysis

correlateEvents

```
public void correlateEvents (ConnectionInfo conInfo, TechEventID id,
TechEventID idToMatch, List eventRelations)
throws AnalyzeEventException
```

Decide whether a relation should be established between the two events passed. If the conclusion is affirmative, generate new relation objects and add them to the given list. **Parameters:**

conInfo - database connection info object for the current project id - event ID object for the current event to be matched idToMatch - event ID object for the potential matching event candidate eventRelations - list of event relations generated

Throws:

AnalyzeEventException - Signals errors during the event correlation analysis

4.5.3.2. Class CorrelationTechHelperBean

package com.bristol.tvision.util.services.analysis.eventanalysis public abstract class CorrelationTechHelperBean extends ChainManagedBean implements IEventCorrelation

This is the abstract base class for all event correlation beans.

Constructor:

• CorrelationTechHelperBean

CorrelationTechHelperBean(java.lang.String technology, int correlationType) throws AnalyzeEventException

Creates an instance of this event correlation bean for the given technology and correlation type. The correlation type is a unique integer and should be ≥ 100 for new user-defined correlation types.

Methods:

createLookupKeys

Refer to the definition of IEventCorrelation.

correlateEvents

Refer to the definition of IEventCorrelation.

• getCorrelationType public java.lang.String getCorrelationType()

Return the correlation type string.

4.5.3.3. Class MQCorrelationData

package com.bristol.tvision.datamgr.dbtypes
public class MQCorrelationData

This class defines a collection of event attributes relevant to the event correlation process. For example, in the IEventCorrelation::correlateEvents method, event attributes for the two events to be matched can be retrieved through a correlation data cache. The attributes are returned in an object instance of this class.

Constructor:

MQCorrelationData

MQCorrelationData(int apiCode, java.lang.String putApplName, java.lang.String putApplType,String userId, int qmgrId, int mqObjId, java.lang.String eventTime, int programId)

Creates an instance of a WebSphere MQ correlation event attribute data collection object based on the given event attributes.

Fields:

- int apiCode
- String putApplName
- String putApplType
- String userId
- Int qmgrId
- Int mqObjId
- String eventTime
- Int programId

4.5.3.4. Class JMSCorrelationData

package com.bristol.tvision.datamgr.dbtypes public class JMSCorrelationData Similar to the class MQCorrelationData, this class defines a collection of event attributes relevant to the event correlation process of JMS events.

Constructor:

• JMSCorrelationData

JMSCorrelationData(**int** methodCode, String appId, String userId, String destination, String eventTime, **int** programId, String putApplType, **int** qmgrId, **int** mqObjId)

Creates an instance of a JMS correlation event attribute data collection object based on the given event attributes.

Fields:

- int methodCode
- String appId
- String userId
- String destination
- String eventTime
- int programId

- String putApplType
- int qmgrId
- int mqObjid

4.5.3.5. Class LookupKey

```
package com.bristol.tvision.datamgr.dbtypes
public class LookupKey
```

This class defines the lookup key object to be used in identifying potential events for correlation purpose.

Constructor:

LookupKey

LookupKey(java.lang.String keyValue, int typeId) Creates a new lookup key instance with the given key and the correlation type id.

Fields:

- String keyValue
- int typeId

Methods:

```
• equals public boolean equals (LookupKey lookupKey)
```

Decide whether the given lookupKey is equal to this key object. The two objects are equal if the corresponding key, correlation type string, and type ID are the same.

Parameters:

lookupKey – lookup key object to be compared

Returns:

true if the two keys are equal, false otherwise

4.5.3.6. Class EventRelation

```
package com.bristol.tvision.datamgr.dbtypes
public class EventRelation
```

This class defines an event relation object between any two events.

Fields:

```
• Relation Type
```

```
public static final int UNKNOWN_PATH = 0;
public static final int MESSAGE_PATH = 1;
public static final int TRANSACTION_PATH = 2;
public static final int BIDIRECTION = 16
```

Type of the event relation:

• MESSAGE_PATH indicates a direct message flow between the two events. That means the two events are associated with the same message data. For example, a MQPUT and MQGET call dealing with the same message bears a message path relation.

- TRANSACTION_PATH indicates a control flow between two events.
- BIDIRECTION is a type mask that indicates the bi-direction nature of the relation between the two events.

Relation Direction

```
public static final int RELATION_PATH_IN = 1;
public static final int RELATION_PATH_OUT = 2;
public static final int RELATION_UNKNOWN = 0;
```

Direction of the event relation. Note that the event object is created in conjunction with an event pair (event1, event2). This indicates the direction from event1 to event2.

Confidence Factor

```
public static final int WEAK_RELATION = 0;
public static final int STRONG_RELATION = 1;
```

This factor is assigned by the event correlation module. There are cases where the correlation module may not have perfect data for a deterministic decision on the event relation generated. In such case, the relation created can carry a WEAK_RELATION confidence factor indicating the uncertainty in the decision.

- *int relation* Bitfield indicating the relation type, e.g. MESSAGE_PATH | BIDIRECTION
- int direction

Bitfield indicating the relation direction, e.g. RELATION_PATH_IN | RELATION_PATH_OUT

• int confidence

Confidence factor, either WAEK_RELATION or STRONG_RELATION

int latency

The latency between the two events in milliseconds

Constructor:

EventRelation

```
EventRelation(int relation, int direction, int confidence, int
latency)
```

Creates a relation object with the given relation type, direction, confidence factor, and latency.

4.5.3.7. Class MQRelationDBService

package com.bristol.tvision.datamgr.dbservices
public class MQRelationDBService

This class defines an internal database service for accessing MQSeries correlation related information. For example, this service works in conjunction with the caching mechanism and stores MQSeries event correlation attributes. The following describes the public interfaces of interest to the custom event analysis beans developers.
Methods:

instance

public static MQRelationDBService instance (java.lang.String schema)

Return the singleton instance of the MQRelationDBServices.

Parameters:

schema - Database schema for the current project

Returns:

Singleton instance of the MQRelationDBService.

getCorrelationData

```
public MQCorrelationData getCorrelationData(java.lang.Connection
con, EventID eventID) throws DataManagerException
```

Return the MQSeries correlation event data for the given event.

Parameters:

con - Java SQL database connection handle, probably from the ConnectionInfo object. eventID - EventID object for the interested event

Returns:

A MQCorrelationData object for the given event.

Throws:

DataManagerException - Signals errors during internal database operations.

4.5.3.8. Class JMSRelationDBService

package com.bristol.tvision.datamgr.dbservices
public class JMSRelationDBService

This class defines an internal database service for accessing JMS correlation related information.

Methods:

instance

public static JMSRelationDBService instance (java.lang.String schema)

Return the singleton instance of the JMSRelationDBServices.

Parameters:

schema – Database schema for the current project

Returns:

Singleton instance of the JMSRelationDBService.

getCorrelationData

```
public JMSCorrelationData getCorrelationData(java.lang.Connection
con, EventID eventID) throws DataManagerException
```

Return the MQSeries correlation event data for the given event. **Parameters:**

con – Java SQL database connection handle, probably from the ConnectionInfo object. eventID – EventID object for the interested event Returns: A JMSCorrelationData object for the given event. Throws: DataManagerException - Signals errors during internal database operations.

4.5.3.9. Sample Custom Event Correlation Bean

Refer to the code in the directory

<TVISION_HOME>/samples/stock/beans/correlation to see a sample implementation of custom event correlation bean (StockTradeRelationshipBean.java).

StockTradeRelationshipBean implements the IEventCorrelation interface and is derived from the class CorrelationTechHelperBean. It builds a custom message path relation between a failed MQGET event (CompCode equals to MQCC_FAILED) and the MQPUT event that participates in the same trade request processing. The stock trade example follows a request-reply messaging model. The StockTrade program records the message ID field of the initial request message, and uses this value as the correlation ID value to be matched when it reads the reply message through the MQGET call. In other words, for a particular transaction, the message ID field in the MQMD object of the StockTrade – MQPUT(1) event should be the equal to the correlation ID field in the MQGET event. The following is the code fragment for the StockTradeRelationshipBean constructor. It specifies that the bean handles MQSeries events and generates custom event relation of type "REQUEST_REPLY_TYPE" correlation as described above:

```
public static final String REQUEST_REPLY_TYPE = 100;
public StockTradeRelationshipBean() throws AnalyzeEventException {
  super(TVisionCommon.TECH_NAME_MQSERIES, REQUEST_REPLY_TYPE);
  }
```

The next code fragement contains the implementation of the createLookupKeys method. As discussed before, the message ID or correlation ID value in the message descriptor record is used as the lookup key for MQPUT(1) and MQGET respectively.

```
public void createLookupKeys(ConnectionInfo conInfo, XMLEvent event,
                             List lookupKeys) throws
AnalyzeEventException {
try {
XPathSearch lookup = new XPathSearch(event);
String correlId;
/* for StockTrade->MQPUT call (request event), use MQMD.MsgID as */
/* lookup key, for StockTrade->MQGET call (reply event), use */
/* MQMD.CorrelId as the lookup key */
switch (StockTradeHelper.getEventType(lookup)) {
case StockTradeHelper.MQSERIES_REQUEST_EVENT:
correlId = lookup.getValue(XPathConstants.MSGID);
if (correlId == null)
    return;
break;
case StockTradeHelper.MQSERIES_REPLY_EVENT:
if (Integer.parseInt(lookup.getValue(XPathConstants.COMPCODE)) !=
   MQDefs.MQCC_FAILED)
return;
  correlId = lookup.getValue(XPathConstants.CORRELID);
    if (correlId == null)
```

```
return;
    break;
  default:
  return;
    }
   /* create a new lookup key and add it to the list */
  LookupKey key = new LookupKey(correlId, REQUEST_REPLY_TYPE);
  lookupKeys.add(key);
     }
     catch (XMLException ex) {
   throw new AnalyzeEventException(ex);
     }
   }
The next code fragment contains the implementation of the correlateEvents method:
  public void correlateEvents(ConnectionInfo conInfo, TechEventID id,
   TechEventID idToMatch, List eventRelations) throws
  AnalyzeEventException {
   try {
   /* Retrieve data relevant for event correlation from cache. */
  Cache cache = AnalysisCacheManager.instance().getCorrelationCache
   (conInfo.schema);
  MQCorrelationData data = (MQCorrelationData) cache.get(id);
  if (data == null) {
     data =
  MQRelationDBService.instance(conInfo.schema).getCorrelationData(
             conInfo.con, id);
     if (data != null)
      cache.insert(id, data);
    else
      return;
   }
  MQCorrelationData dataToMatch = (MQCorrelationData)
  cache.get(idToMatch);
   if (dataToMatch == null) {
     dataToMatch =
       MQRelationDBService.instance(conInfo.schema).getCorrelationData(
         conInfo.con, idToMatch);
     if (dataToMatch != null)
       cache.insert(idToMatch, dataToMatch);
     else
      return;
   }
   int apiId = data. apiCode;
   int apiIdToMatch = dataToMatch.apiCode;
   if (apiId != apiIdToMatch) {
    EventRelation eventRelation = new EventRelation();
     eventRelation.setRelation(EventRelation.MESSAGE_PATH |
                               EventRelation.BIDIRECTION);
     eventRelation.setDirection(EventRelation.RELATION_UNKNOWN);
     eventRelation.setConfidence(EventRelation.STRONG_RELATION);
     eventRelations.add(eventRelation);
```

}

Chapter 4 • Reference - Extending the Analyzer *Performing Event Analysis*

```
catch (DataManagerException ex) {
   throw new AnalyzeEventException(ex);
}
```

The AnalysisCacheManager object provides an internal memory cache for storing selected attributes of the events to be matched. Refer to the MQCorrelationData class definition for a list of attributes supported. This cache allows quick access to certain event attributes without executing an event data query, thus improving the correlation process performance. To decide whether the two events are indeed related, the API code of the two events are compared to ensure that one event is MQPUT(1) and the other one is MQGET. Since only MQPUT(1) and MQGET events can be potential candidates, it is enough to check whether the two event API codes are different or not.

Once it is decided that the two events are related, a new event relation object is created and inserted to the relation list. The relation is of type MESSAGE_PATH, has no direction attribute, and has a STRONG_RELATION confidence factor.

The following code fragment is the change to the Beans.xml file for including this custom event correlation bean. It tells the Analyzer framework to load and run the

StockTradeCorrelationBean bean as a part of the CorrelationMQHelperCtx context. This bean will be invoked after the default MQToMQRelationshipBean for every MQSeries event.

<Module type="Context" name="CorrelationTechHelperCtx">

```
<Module type="Context" name="CorrelationMQHelperCtx"
class="com.bristol.tvision.services.analysis.eventanalysis.CorrelationMQH
elperCtx">
        <Module type="Bean"
        class="com.bristol.tvision.services.analysis.eventanalysis.MQToMQRelati
        onshipBean"/>
        <Module type="Bean" class="com.bristol.tvision.demo.
        stock.StockTradeRelationshipBean"/>
        </Module>
```

</Module>

4.5.4. Custom Business Transaction Attributes and Classification

Business transaction attributes are stored in the table BUSINESS_TRANSACTION which is defined by an XDM file, and thus are easily extensible. Additional custom business transaction attributes can be simply added by modifying the corresponding Transaction.xdm file. The database schema which is defined by the standard XDM definition is as follows:

BUSINESS_TRANSACTION

business_trans_id: INTEGER
class_id: INTEGER
starttime: CHAR(20)
endtime: CHAR(20)
responsetime: BIGINT
state: INTEGER
result: INTEGER
label: VARCHAR(128)
sequential_id: INTEGER

- business_trans_id: a unique ID for the transaction generated by the database
- class_id: the ID of the transaction class (FK into table transaction_class)
- starttime: the start time of the transaction
- endtime: the end time of the transaction
- responsetime: the time difference between start and end time
- state: the current state of the transaction (UNKNOWN, IN_PROCESS, COMPLETED)
- result: the result of the transaction (SUCCESS, FAILED)
- label: a label for the transaction to display in the GUI
- sequential_id: a unique ID which gets incremented every time the transaction has been updated

When modifying the XDM definition to add custom business transaction attributes it is important not to alter or delete any of those predefined "standard" attributes. If no standard or custom transaction classification bean is plugged in into the Analyzer framework, the attributes will get populated with the following values during event transaction analysis:

business_trans_id	generated by the database
class_id	XMLTransaction.UNCLASSIFIED_ID (-1)
starttime	time of the earliest event in this business transaction
endtime	time of the latest event in this business transaction
state	XMLTransaction.Unknown (-1)
result	XMLTransaction.Unknown (-1)
label	null
responsetime	difference between starttime and endtime
sequential_id	generated by the database

There are two different ways to populate the values of custom transaction attributes or to modify the default values of the standard attributes:

• Use the StandardClassifyTransactionBean and define rules how to classify transactions and update attribute values. This approach does not require any additional coding, only the rule definition file has to be edited.

• Write a custom classification bean that implements the IClassifyTransaction interface. This approach is useful if more complex transaction classification is needed than the standard classification bean can provide

4.5.4.1. Transaction Classification

By default, TransactionVision does not classify the business transactions it processes; the class ID of each transaction will be 0, indicating that this transaction does not belong to any transaction class. To enable transaction classification, the following steps (which are explained in more detail in sections 3.5.4.2 - 3.5.4.6) are required:

- Enable classification in the Beans.xml file by removing the comment around the ClassifyTransactionCtx section and by placing the appropriate classification bean (standard or custom classification bean) into it.
- Define your classification rules in the file TransactionDefinition.xml (if using the standard classification bean).
- Insert each class with its attributes into the database table TRANSACTION_CLASS. The table must be populated before any transactions are processed by the Analyzer.

4.5.4.2. Transaction Classification with the Standard Classification Bean

The StandardClassifyTransactionBean is a default implementation of a classification bean and allows user customized transaction classification without the need to write a single line of code. Although the rule engine of this standard bean is simple and fairly limited, it may well be sufficient for a great amount of classification cases. It is well suited for transactions that can be classified based on the attributes of one event of the transaction. The classification logic is driven by rules in the configuration file

\$TVISION_HOME/config/services/TransactionDefinition.xml which define how and when transaction attributes are set or updated. These rules will get evaluated for each event being processed in the transaction analysis in the Analyzer. The main structure of this configuration file is:

<TransactionDefinition>

The transaction definition consists of one ore more <Class> definitions that contain rules that are applicable to events and transactions of that particular transaction class. The attribute @name has to be a valid transaction class name which has a corresponding entry in the transaction_class table. Each class definition can have an optional attribute @dbschema which restricts the definition to one or more (specified as a comma separated list) database schemas. If the database schema for the current event does not match the schema tag for the class definition, none of the rules for this class will get evaluated. If the schema attribute is missing, the definition is valid for all database schemas.

Each <Class> definition consists of one or more <Classify> sections that contain rules for identifying the transaction class, and a list of rules for setting and updating the transaction attributes. Each <Classify> tag needs an arbitrary, but unique @id attribute. The evaluation flow is as follows:

• If the current transaction has not been classified yet (class_id == XMLTransaction.UNCLASSIFIED_ID), then all <Classify> sections of all class definitions matching with the current event schema are evaluated. If a classification is successful, the transaction class ID of the transaction will get set and all attribute rules contained in the class definition will get evaluated as well. No further <Classify> section will be evaluated any more. If none of the classifications are successful, the union of all attribute rules (outside of <Classify> sections) of **all** class definitions for the current event schema are evaluated.

Note: This is necessary because the processing order of events in the analyzer can be different to the order the events really happened, and the classification algorithm needs to make sure that all rules for a certain class will get evaluated even if the event which will classify the transaction will be processed at a later time. As a consequence, rules outside of <Classify> sections should always be specific enough (by defining appropriate matching rules) to match only on events of the class they are meant for, because they will also get executed on events that might belong to another class for which the classifying event has not been processed yet.

• If the current transaction already has its class attribute set, only the attribute rules in the corresponding class definition outside of the <Classify> sections are evaluated. The conditions inside of the corresponding <Classify> section are **not** evaluated again.

Each <Classify> section contains one ore more <Match> conditions, e.g.:

```
<Class name="StockTrade" dbschema="Stock ">
<Classify id="1">
<Match xpath="/Event/Technology/JMS/Caller" operator="EQUAL"
value="StockTrade"/>
<Match xpath="/Event/Technology/JMS/MQObject/Queue"
operator="EQUAL" value="TRADE_REQUEST"/>
{...}
```

If the logial AND of these conditions results in true, the current transaction is considered to be 'classified', and the class_id attribute of the current transaction is set to the corresponding class ID of the definition class. In general, a match condition consist of a @xpath, @operator, and @value attribute. The @xpath attribute specifies a certain

value from either the current XML event or the transaction document. @operator can be either EQUAL or UNEQUAL, and @value can either contain a literal string value or an enumeration constant (if there is an enumeration defined for this XPath). The condition gets evaluated by string comparison of the document value with the specified value.

As mentioned before, the match conditions in one <Classify> section are logically AND-ed together. To specify an alternative set of conditions (logical OR), one or more additional <Classify> sections for the same class can be added.

In addition to the <Classify> section, each class definition can contain zero or more attribute rules to set or modify other transaction attributes. Here is an example of such an attribute rule:

```
<Attribute name="Declined">
<Path>/Transaction/Declined</Path>
<ValueRule name="SetDeclined1">
<Match xpath="/Event/StdHeader/ProgramName" operator="EQUAL"
value="DeclineTrade01"/>
   <Match xpath="/Event/Technology/MOSeries/MOObject/@objectName"
operator="EQUAL" value="TRADE REPLY"/>
    <Value type="Constant">true</Value>
</ValueRule>
<ValueRule name="SetDeclined2">
<Match xpath="/Event/StdHeader/ProgramName" operator="EQUAL"
value="DeclineTrade02"/>
    <Match xpath="/Event/Technology/MQSeries/MQObject/@objectName"
operator="EQUAL" value="TRADE REPLY"/>
   <Value type="Constant">true</Value>
</ValueRule>
</Attribute>
```

Each <Attribute> element defines rules for setting the value of a certain transaction attribute. An arbitrary but (for the class) unique @name attribute is required. The <Path> element specifies the Xpath for the transaction attribute. The possible values for this transaction attribute are specified in one or more <ValueRule> sections. Each <ValueRule> specifies a set of match conditions (logical AND) and the new value for the attribute if the match conditions 'fire'. Again, each <ValueRule> requires an arbitrary but unique @name attribute. The <ValueRule> definitions for an <Attribute> are evaluated in sequential order, and once a certain rule has 'fired', the transaction attribute will get updated with the value defined within this rule, and all following <ValueRule> sections will get skipped.. The new values for a transaction attribute are specified within the <Value> element and can have one of two possible types (specified with the @type attribute):

- "Constant" specifies a literal String value or an enumeration constant (if there is an enumeration defined for this XPath)
- "XPath" specifies that the new value should be retrieved dynamically at runtime from either the XML event or transaction document

It is possible to specify multiple <Value> element for one attribute, in which case the attribute value will be the concatenation of all evaluated <Value> definitions, like .e.g.:

```
<Attribute name="Label">
<Path>/Transaction/Label</Path>
<ValueRule name="SetLabel">
<Value type="XPath">/Event/Data/Order/Ticker</Value>
<Value type="Constant">_</Value>
```

```
<Value type="XPath">/Transaction/Account</Value>
<Value type="Constant">_</Value>
<Value type="XPath">/Transaction/OrderID</Value>
</ValueRule>
</Attribute>
```

Every time the transaction analysis calls into the standard classification bean for an event all <Attribute> definitions for the corresponding transaction class are getting evaluated in sequential order. But by default the <Attribute> rules are only evaluated if the corresponding transaction attribute has no value yet, the definition is considered to be "final". Once a final rule has set the value of the transaction attribute, it (and other final rules that refer to the same attribute) will not be evaluated again.

To allow transaction attributes to get set and updated more than once, the attribute rule can be declared with an attribute @final set to "false":

This forces an attribute rule to get evaluated every time, even when the transaction attribute is already set. An attribute rule without the @final attribute is equivalent to @final="true". Another rule attribute, @precedence, can be used to control the setting of new values for transaction attributes :

```
<Attribute name="State" precedence="true">
<Path>/Transaction/State</Path>
{...}
```

This attribute can only be set for rules referencing integer valued transaction attributes. If set to true then an existing attribute value only gets overwritten if the new value is greater than the old value. This mainly makes sense for 'state' and 'result' like attributes where all values can be ordered according to a priority (e.g. UNKNOWN->PROCESSING->COMPLETE), though in general it can be applied to any integer valued attribute. All @precedence rules are automatically considered to be non-final too. By default (if the @precedence attribute is not specified) the value is false.

Any transactions that have been successfully classified will show up with their respective class name in the reports that categorize by class, such as the Transaction Tracking Report. Also, any errors that are encountered during the classification process will get logged in the Analyzer.log file.

4.5.4.3. Classification Action Rules

In addition to setting values within a classification value rule, custom actions can be performed when the value rules are met. This is done by specifying a java class implementing com.bristol.tvision.services.analysis.eventanalysis.IAnalyzerAction with an <Action> element under the <ValueRule> element.

In this example, if both the StartTime and EndTime of the transaction documents have been set (not equal to empty string), set the transaction attribute named "State" to "Completed". When this occurs, the bean specified in the action tag is invoked. This allows invocation of response time computation logic only once per transaction. This logic can also log service level violations or other transactional information.

<Action> elements can be chained. If more than one <Action> element is specified within a classification attribute, they will be invoked in the order they appear as long as the actions return true. As soon as one action returns false, the invocation chain is stopped for that transaction.

Currently, the only Action type available is "JAVACLASS". Code and reason provide a means of passing an integer and/or string for use in the action method. They are not required.

The com.bristol.tvision.services.analysis.actions.LogSLAViolation class provided with TransactionVision logs service level agreement violations for a given transaction to AnalyzeActivityLog defined in the Analyzer.Logging.xml:

```
<category additivity="false"
   class="com.bristol.tvision.util.log.XCategory"
   name="AnalyzerActivityLog">
    <priority class="com.bristol.tvision.util.log.XPriority"
      value="info"/>
      <appender-ref ref="ANALYZER_ACTIVITY_LOGFILE"/>
      <appender-ref ref="AMIT_APPENDER"/>
</category>
```

If you write a custom action class, it must implement com.bristol.tvision.services.analysis.eventanalysis.IAnalyzerAction interface and must provide an action method to be invoked by the standard classification bean. The custom class is added to the Analyzer's CLASSPATH by setting service_additional_classpath in the Analyzer.properties file.

4.5.4.4. The ClassifyTransactionCtx and the IClassifyTransaction Interface

Transaction classification beans are plugged in into the Analyzer framework by placing them into the ClassifyTransactionCtx in the Beans.xml file; for example:

```
<Module type="Context" name="ClassifyTransactionCtx">
        <Module type="Bean"
class="com.bristol.tvision.services.analysis.eventanalysis.StandardCl
assifyTransactionBean"/>
        </Module>
```

The context can contain multiple beans, in which case the beans are processed in sequential order. Each classification bean has to implement the IClassifyTransaction interface:

public boolean

classify(com.bristol.tvision.services.analysis.XMLEvent event,

Performs transaction classification **Parameters:**

event - The current event txn - The transaction document for the current event correlatedEvents - The list of correlated events conInfo - The current database connection

Returns:

true if the transaction doc has been updated, false otherwise

Throws:

com.bristol.tvision.services.analysis.eventanalysis.AnalyzeEventExce
ption - The analysis process failed

For each event that gets processed during the event transaction analysis phase the classify method of each registered classification bean will be called, and the logical OR of all bean invocations will be returned back to the transaction analysis phase in the Analyzer. If the returned value is true (meaning one or more beans have modified the transaction document) the corresponding row values in the business_transaction table will get updated by the Analyzer framework.

By default, the ClassifyTransactionCtx is disabled in the Beans.xml file. To enable the standard classification, remove the XML comments around the following section :

4.5.4.5. Writing a Custom Classification Bean

A classification bean has to implement the classify interface described above and can trigger the update of business transaction attributes by modifying the XMLTransaction object (the business transaction for the current event), which gets passed into the call. The bean has access to all XMLDocument values in the current event and the

corresponding business transaction object by using the method getDocumentValue (String xpath); for example:

```
String progName =
event.getDocumentValue(XpathConstants.PROGRAM_NAME);
String oldLabel = txn.getDocumentValue(XMLTransaction.LABEL_XPATH);
```

The bean can set and modify all of the additional custom transaction attributes, and most of the standard ones. The only exception is business_trans_id; updating this value is not allowed and may lead to unexpected results in the Analyzer. The update of transaction attributes is done by using the method setDocumentValue(String xpath, String value); for example:

```
tnx.setDocumentValue(XMLTransaction.LABEL_XPATH, newLabel);
```

If the bean has modified any of the transaction attributes, it has to return a boolean true value from the classify call; otherwise, the new values will not be written to the database in the Analyzer framework.

If the transaction document remains unchanged, the bean should return false to avoid unnecessary database write overhead.

To classify a certain transaction, the bean has to update the class_id attribute of the transaction document (XMLTransaction.CLASS_ID_XPATH). This integer value is a foreign key into the transaction_class table and thus should only contain values that correspond to valid transaction class entries. The transaction class Ids can easily be accessed by using the utility class CachedTransactionClass :

```
int classId = CachedTransactionClass.getClassId(conInfo,
className);
```

As the transaction class table content is static, the utility class reads the transaction class data only once from the database and returns all Ids without any further database access.

4.5.4.6. The Transaction Class Table

As attributes for a certain transactions are stored in the business_transaction table, transaction classes and their attributes are stored in the table transaction_class. The standard database schema (without any user-defined class attributes added) and the relationship to the business_transaction table is as follows:



A row in this table makes a certain transaction class 'known' to the analysis system. The table contents is static, meaning that all transaction classes have to be defined (and the table populated) before the Analysis service is started. The Analyzer will read the transaction class definitions at startup and use this information to map class names (wherever specified) to the corresponding class Ids stored in the database.

The table schema is defined by the XDM mapping TransactionClass.xdm . The standard mapping is:

Custom transaction class attributes can be simply added by editing the XDM file and adding the appropriate <Column> definitions. Here is an example how to add a custom attribute 'SLA':

```
<Column name="SLA" type="INTEGER" description="SLA">
<Path>/TransactionClass/SLA</Path>
</Column>
```

Although there is no "TransactionClass" XML document that gets processed by the XMLDatabaseMapper (as mentioned above the contents is static), defining the transaction_class table through a XDM mapping has the advantage of allowing queries on transaction data that can include references to the transaction class attributes. The class_id column definition in the business transaction XDM mapping includes a JOIN reference to the transaction class table and thus makes it possible to create queries that use both document types.

There is currently no utility that can aid in populating this table, so all class definitions have to be inserted manually, e.g. by an SQL script. A sample script would look like this:

```
INSERT INTO schema.transaction_class (class_id, class_name, SLA)
VALUES(1, `StockTrade', 5)
INSERT INTO schema.transaction_class (class_id, class_name, SLA)
VALUES(2, `CashFlow', 0)
{...}
```

The class IDs and class names should be unique, and the class ID value 0 is reserved for the UNCLASSIFIED transaction class.

Important! It is important to distinguish the class_id for each Class definition from the classify id set in each Classify section within a Class definition. You should enter one unique class id value into the transaction_class table for each Class that you define. Do not add an entry for each Classify section that you define. The classify ids are only required for the standard classification bean to uniquely identify each classify section.

4.5.4.7. Business Groups

A Business Group is a group of one or more Transaction Classes or child business groups. There are two tables that must be populated in order to make use of reports such as the Business Impact Report which use and display information by Business Group.

Business Group Table:

The Business Group table defines the set of business groups to be used. Each table entry defines the text name of the group, assigns it a unique id, and the id of its parent group. If the group is a root level group and has no parent, the parent id should be specified as -1.

The schema definition is as follows: <u>BUSINESS_GROUP</u> BIZGRP_NAME: VARCHAR(64) - name of the business group PARENT_BIZGRP_ID: INTEGER - ID of the parent business group of this group (-1 if no parent) BIZ_GRP_ID: INTEGER - ID of this business group

Transaction Class to Business Group Table:

The Transaction Class to Business Group table assigns transaction classes to their corresponding business groups. Each table entry specifies the id of a transaction class and the id of its corresponding business group.

The schema definition is as follows: <u>TRANSCLASS_TO_BIZGRP</u> TRANSCLASS_ID: INTEGER - ID of the Transaction Class BIZGRP_ID: INTEGER - ID of the Business Group In order to use the Business Impact Report and other future reports that use Business Groups, these two tables must be populated. As with the Transaction Class table, there is currently no utility that can assist in populating these tables, and they must be populated manually through an SQL script.

A sample would look like the following: INSERT INTO TRADE.BUSINESS_GROUP(BIZGRP_NAME, PARENT_BIZGRP_ID, BIZGRP_ID) VALUES('Purchase', -1, 0); INSERT INTO TRADE.BUSINESS_GROUP(BIZGRP_NAME, PARENT_BIZGRP_ID, BIZGRP_ID) VALUES('Trade', -1, 1); INSERT INTO TRADE.TRANSCLASS_TO_BIZGRP(TRANSCLASS_ID, BIZGRP_ID) VALUES(1, 0); INSERT INTO TRADE.TRANSCLASS_TO_BIZGRP(TRANSCLASS_ID, BIZGRP_ID) VALUES(2, 0); INSERT INTO TRADE.TRANSCLASS_TO_BIZGRP(TRANSCLASS_ID, BIZGRP_ID) VALUES(2, 0); INSERT INTO TRADE.TRANSCLASS_TO_BIZGRP(TRANSCLASS_ID, BIZGRP_ID) VALUES(2, 1);

4.6. Extending the System Model

Use the <TVISION_HOME>/config/services/RemoteDefinition.xml file to define objects in your system that the sensor might otherwise not be able to fully resolve.

For example, suppose you have a remote queue on queue manager QM1 that points to some queue on queue manager QM2. A sensored application putting to the queue on QM1 does not connect to QM2 to fully discover what type of object the final destination queue is. The destination queue might be an alias queue or even another remote queue. If no sensored application on QM2 ever connects directly to the destnation of the QM1 remote queue, then the object will never be fully resolved, possibly resulting in a missing link in the correlation of events.

By manually defining objects in RemoteDefinition.xml, you can specify the details of objects that the sensor could not completely resolve otherwise.

Each <RemoteObject> tag defines an object. When the analyzer attempts to resolve the target of a remote queue, it checks whether an entry exists with the same object and queue manager name. If such a match is found, the MQObject definitions within the RemoteObject tag will replace the generic queue definition provided by the sensor. Embedding an additional MQObject tag within the first MQObject tag creates a "resolveto" relationship.

Therefore, the first RemoteObject tag in the following example can be interpreted as: If the destination of a remote queue has the name RALIAS2.QUEUE on queue manager host.tv2.manager, create for this object an alias queue RALIAS2.QUEUE that resolves to a local queue RRR.QUEUE.

Possible values for the objectType attribute include:

- Q_LOCAL
- Q_MODEL
- Q_ALIAS
- Q_REMOTE
- Q_CLUSTER

- Q_LOCAL_CLUSTER
- Q_ALIAS_CLUSTER
- Q_REMOTE_CLUSTER

Take care in creating and modifying these definitions as inserting objects that don't actually match the topology of your system could break the correlation of events.

Example RemoteDefinition.xml file:

```
<?xml version="1.0" encoding="UTF-8"?>
<RemoteDefinition>
       <RemoteObject objectName="RALIAS2.QUEUE" queueManager="perplex7.tv2.manager">
               <MQObject objectName="RALIAS2.QUEUE" objectType="Q_ALIAS"
queueManager="perplex7.tv2.manager">
                      <MQObject objectName="RRR.QUEUE" objectType="Q_LOCAL"
queueManager="perplex7.tv2.manager"/>
              </MQObject>
       </RemoteObject->
       <RemoteObject objectName="TEST.CLUSTER.QUEUE" queueManager="SECOND_CLUSTER">
       <MQObject objectName="TEST.CLUSTER.QUEUE" objectType="Q_REMOTE"
queueManager="SECOND_CLUSTER">
               <MQObject clusterName="SECOND_CLUSTER" objectName="TEST.CLUSTER.QUEUE"
objectType="Q_LOCAL_CLUSTER" queueManager="deepakelap.tv3.manager"/>
       </MQObject>
       </RemoteObject>
```

</RemoteDefinition>

4.7. Generating Application Events to Tivoli Enterprise Console (TEC)

TransactionVision allows plugging in custom code to generate TEC events when certain application events occur. These can either be plugin beans into the Analyzer or as scheduled jobs running in the application server hosting the UI. This custom code can use the log4j classes to generate log messages. The log4j appender, TECAppender, routes these log messages to Tivoli, when enabled. The MonitoringEvent class is provided to allow setting of parameters into the log4j message, which are then mapped to Tivoli slots by the TECAppender. The TECAppender uses the file <TVISION_HOME>/config/logging/tivoli/SlotMap.properties to map MonitoringEvent parameters to Tivoli slots.

4.7.1. Class MonitoringEvent

public class MonitoringEvent implements Cloneable

Constants: // severity levels. public static int FATAL = 1; public static int ERROR = 2; public static int WARN = 2; public static int INFO = 3;

public static final String TVISION_EVENT_APPLICATION =
"TVISION_EVENT_APPLICATION";

Constructor: public MonitoringEvent(String clas, int s) Takes in constant TVISION_EVENT_APPLICATION as the class name and severity levels FATAL, ERROR, WARN and INFO, as defined above.

Methods:

public void setMessage(String msg) public String getMessage() Sets and gets the message string to be logged.

public void setParameter(String name, String value) public String getParameter(String name) Sets and gets additional parameters that can be set into Tivoli slots.

4.7.2. SlotMap.properties

This file is used by the log4j TECAppender to allow mapping of parameters set into MonitoringEvent to Tivoli slots. The file format is:

<MonitoringEvent parameter> = <Tivoli slot>

Any parameter specified here is explicitly mapped to a Tivoli slot. Parameter names unspecified in this file are mapped to Tivoli slots $tv_attrib[1|2|3]$ and their values are mapped to slots $tv_value[1|2|3]$.

4.7.3. Example Usage:

The following sample code writes an ERROR log message of class BTV_app_red, with parameters "application", "transaction_class" set.

```
MonitoringEvent ev = new MonitoringEvent
(MonitoringEvent.TVISION_EVENT_APPLICATION, MonitoringEvent.ERROR);
ev.setParameter("application", "Trade");
ev.setParameter("transaction_class", "TRADE_CLASS");
ev.setParameter("message_id", "TransactionError");
ev.setMessage("Error fulfilling transaction xyz");
Logging.analyzerActivityLog.error(ev);
```

4.7.4. BTV Class Definitions and Rulebase

Class definitions supplied address the following events:

- Internal events events generated regarding the TransactionVision application itself
- Applications events events generated by entities that TransactionVision is monitoring
- Unknown events events that have not fit the criteria to be defined beyond coming from TransactionVision.

• Escalation events - events of either internal or application that have exceeded count thresholds The rules file creates the following classes related to TransactionVision:

BTV_app_black BTV_app_red BTV_app_yellow BTV_app_green BTV_int_black BTV_int_red BTV_int_yellow BTV_int_green BTV_unk

The classes BTV_int_[black|red|yellow|green] are used by TransactionVision internally while the classes BTV_app_[black|red|yellow|green] may be used by application plugin code. The color black, red, yellow, green indicates the severity level to be FATAL, ERROR, WARN and INFO respectively.

The following slots will be created:

message_id tv_component tv_attrib1 tv_attrib2 tv_value1 tv_value2 tv_value3 err_code application event_time transact_class transact_id

All slots may not be filled by TransactionVision internal messages.

Rulesets have supplied rules for the following:

- First instance rule which takes action upon an event the first time it arrives, or if there are no other like events in either OPEN or ACK status
- Duplicate rule which identifies an event as a duplicate to a previous event in either OPEN or ACK status, increments the repeat count on the original event, and drops the new event
- Escalation rule which takes action when an event has been received in succession for a defined count and status is of OPEN or ACK
- internal events, which are focused on the TV application itself.

5. Using the Query Services

The Query Services interfaces provide a means to retrieve XDM mapped data from the database using an XML based query document. The Query Services consist of the following interfaces and classes: QueryServices interface is the top-level interface to create and run queries. The methods in this class return an object that implements the Query interface, which can be used to execute the query. Many of the methods in this class take an object implementing the QueryDoc interface as a parameter. The QueryDoc object describes the query to be obtained in the form of an XML document. A WhereClause can be set into the QueryDoc, which describes what matching criteria should be used, and a SelectClause, which describes what fields should be retrieved. A Cursor object is returned from several of the QueryServices methods, which allows a user to iterate over the results. The QueryServices implementation converts the input XML query into an SQL statement and executes it. The Cursor class is a wrapper around the JDBC cursor classes.

The following sections will describe each of these objects and interfaces and show sample code to document their usage.

5.1. Sample Usage

The following sample code shows how to create a query document, populate the document with a query description, use the QueryServices interface to get a Query object back and then execute the query. The sample counts the number of events for each MQPUT, MQPUT1 and MQGET.

```
qdoc.updateWhereClause(clause);
```

```
// select the fields to be retrieved in this case the program id.
String[] selects = { XPathConstants.PROGRAM_ID };
qdoc.insertSelect(selects);
// gets and execute the query.
Cursor queryCursor = customReportBean.getQueryResults(qdoc);
// map of API name versus event count for that API.
HashMap nameToCount = new HashMap();
int maxValue = 0;
// iterate through the query fetching the results from the database.
while (queryCursor.next()){
    String objValue = queryCursor.getValue(1,true);
    Integer count = (Integer)nameToCount.get(objValue);
    if (count == null)
    {
        count = new Integer(1);
        nameToCount.put(objValue,count);
    } else {
        int newValue = count.intValue() + 1;
        nameToCount.put(objValue,new Integer(newValue));
        if (newValue > maxValue)
            maxValue = newValue;
    }
}
```

The method getQueryResults used in the above code snippet is as follows. This method gets the QueryService instance (QueryService is a singleton object per schema), gets an event list query object and executes the query, returning the result set cursor.

5.2. Class QueryServices

public class com.bristol.tvision.datamgr.query.QueryServices extends java.lang.Object

QueryServices is the main interface to query the XDM tables. It is a singleton object that has methods that take a XML query document as the query definition and returns a query object. This query object can then be executed to obtain a cursor, which is then used in consecutive calls to retrieve data. All the methods in this interface that get a cursor or data from the database require a valid JDBC SQL connection handle. The methods throw a DataManagerException on an error condition occurring.

This interface defines the following methods.

5.2.1. Methods:

instance

This method returns the singleton instance for the specified schema

Parameters:

schema - The schema for which the instance should be returned. The schema name used by the current project can be obtained by using the TVisionServlet class documented in Chapter 4 of this manual.

Returns:

The return value is a reference to the singleton instance.

Example:

A servlet requiring access to a QueryServices instance could use the code below to get the schema name using the getSessionBeanFromSession() method and then use the instance method of the QueryServices singleton using the schema name.

```
String schemaName =
TVisionServlet.getSessionBeanFromSession(session).getSchemaName();
QueryServices queryServ = QueryServices.instance(schemaName);
```

getEventDetail

This method returns the event XML document for a given event.

Parameters:

con	The database connection to use
eventId	The specified event
convSvr	The TypeConvService allows fields like date and time formatting, time-zone and other conversions to be applied to the retrieved data. A value of null implies that no conversions are applied. Refer to the section on TypeConvService for more information on the supported conversions.

Returns:

The return value is an XML document containing event data.

Throws:

DataManagerException - if retrieving of the XML document fails

getUserDataLength

This method returns the length of a given message data segment for a given event. Typically, message data is segmented when a data collection filter using data ranges is used to collect data. In that case, this method allows you to get the size of a particular data segment.

Parameters:

con -	he database connection to use.
eventId	The event id the event that the message data belongs to.
dataNum	The segment number of the message data, where the first segment has index 0.

Returns:

The return value is the length of the message data segment.

Throws:

DataManagerException – occurs if the database operation fails.

getUserData

This method returns a segment of a message data segment. This segment is specified by a starting offset (offset) and the length (length) to return.

Parameters:

con	The database connection to use.
eventId	The event id the user data belongs to.
dataNum	The segment number of the user data.
offset	The starting offset of the segment to retrieve.
length	The number of bytes to return.

Returns:

The return value is the message data part of the event of id eventId.

Throws:

DataManagerException - if database operation fails.

Example:

The following code retrieves the first (index 0) segment of the message data buffer into a byte array.

```
QueryServices queryService =
    QueryServices.instance(TVisionServlet.getSchemaNameFromSess
ion(session));
int dataLength =
    (int)queryService.getUserDataLength(con, eventId, 0);
byte[] rawData =
    queryService.getUserData(con,eventId,0,0,dataLength);
```

getEventListQuery

This method creates a Query object for the given event list query document. The Cursor obtained from executing the query can be used in the following calls to getNextEventListDocument to get a specific part of the result returned as an XML document.

Parameters:

con	The connection to use for executing the query
queryDoc	The XML query document specifying the event list query

Returns:

A Query object ready for execution

Throws:

DataManagerException - if parsing the query doc or executing the query fails.

Example:

The following sample gets and executes a query for a given query document. It then creates a document containing an event list of all the events in the database. Note, that this event list does not contain all the event data, only those that have been indexed in lookup tables.

getNextEventListDocument

This method returns the event list XML document for a given query cursor, start index, and number of next rows to return. This event list document does not contain the complete event, but only the data in database lookup tables.

Parameters:

cursor	The query cursor on the events
startInd ex	The index of first row to include in the document
nrOfRows	The number of rows following the stating position to include in the document
convSvr	The TypeConvService allows fields like date and time formatting, time-zone and other conversions to be applied to the retrieved data. A value of null implies that no conversions are applied. Refer to the section on TypeConvService for more information on the supported conversions.

Returns:

The return value is the XML event document for the event list.

Throws:

DataManagerException - if retrieving of the data or assembly of the XML document fails

The format of the returned XML documents is:

```
<?xml version="1.0" encoding="UTF-8"?>
<EventList>
        <EventListItem Program="Trade" APICode="..." ... />
        <EventListItem ... />
        ...
<//EventList>
```

Each <EventListItem> contains the data for one row of the result.

getPreviousEventListDocument

This method returns the event list XML document for a given query cursor, start index, and number of previous rows to return.

Parameters:

cursor	The query cursor on the events
startIndex	The index of first row to include in the document

nrOfRows	The number of rows preceeding the stating position to include in the document
convSvr	The TypeConvService allows fields like date and time formatting, time-zone and other conversions to be applied to the retrieved data. A value of null implies that no conversions are applied. Refer to the section on TypeConvService for more information on the supported conversions.

Returns:

The XML event document for the event list

Throws:

DataManagerException - if retrieving of the data or assembly of the XML document fails

insert

public void insert(java.sql.Connection con,

com.bristol.tvision.services.analysis.XMLDocument doc)
 throws com.bristol.tvision.datamgr.DataManagerException

Inserts all values from the given XML document for which a XDM mapping is defined into the corresponding lookup tables.

Parameters:

con	The database connection to use
doc	The XML document

Throws:

com.bristol.tvision.datamgr.DataManagerException - If the database insert fails

update

Updates all rows in the lookup table which get selected by the query document. The columns to update and the new values are passed in as a Map: The key is the XPath specifying the column, the value is the new value. Note that the WHERE conditions in the query document are only allowed to reference one lookup table.

Parameters:

con	The database connection to use
values	The map containing column xpaths and new values
queryDo c	The query document specifying which rows to update

Throws:

com.bristol.tvision.datamgr.DataManagerException - If the database update fails

delete

Deletes all rows in the lookup table which get selected by the query document. Note that the WHERE conditions in the query document are only allowed to reference one lookup table.

Parameters:

con	The database connection to use
queryDo c	The query document specifying which rows to delete

Throws:

com.bristol.tvision.datamgr.DataManagerException - If the database delete fails

getLocalTransactionQuery

This method creates a Query object to query the database for all events contained in the same local transaction as eventId. Execution of the query returns a Cursor which can be used in following calls to getEventListDocument to get a specific parts of the result returned as an XML document. The SELECT clauses in the query document define which rows to include in the result, the WHERE clauses are ignored.

Parameters:

con	The connection to use for executing the query
queryDoc	The XML query document specifying the rows to include in the result

Returns:

A Query object ready for execution

Throws:

DataManagerException - if parsing the query doc or executing the query fails

get Business Transaction Query

This method creates a Query object to query the database for all events contained in the same business transaction as eventId. Execution of the query returns a Cursor which can be used in following calls to getEventListDocument to get a specific part of the result returned as an XML document. The SELECT clauses in the query document define which rows to include in the result, the WHERE clauses are ignored.

Parameters:

con	The connection to use for executing the query
eventId	The event ID
queryDoc	The XML query document specifying the rows to include in the result

Returns:

A Query object ready for execution

Throws:

DataManagerException - if parsing the query doc or executing the query fails

getBusinessTransactionQuery

public	Query	<pre>getBusinessTransactionQuery(java.sql.Connection con,</pre>
		int businessTxnId,
		org.w3c.dom.Document queryDoc)
		throws DataManagerException

This method creates a Query object to query the database for all events contained in the business transaction denoted by businessTxnId. Execution of the query returns a Cursor that allows access to all columns specified in queryDoc

Parameters:

con	The connection to use for executing the query
businessTxn Id	The business transcation ID
queryDoc	The XML query document specifying the rows to include in the result. The WHERE clauses are ignored.

Returns:

A Query object ready for execution.

Throws:

DataManagerException - if parsing the query doc or executing the query fails

getCorrelatedEventsQuery

This method creates a Query object to query the database for all events correlated to the event denoted by eventId. Execution of the query returns a Cursor that allows access to all columns specified in select section of the query, as well as to the column "confidence", "direction", and "relation_type" of table event_relation.

Parameters:

con	The connection to use for executing the query
eventId	The eventID
queryDoc	The XML query document specifying the rows to include in the result. The WHERE clauses are ignored.

Returns:

A Query object ready for execution

Throws:

DataManagerException - if parsing the query doc or executing the query fails

$update {\it Business Transaction Label}$

public	void	updateBusinessTransaction	Label	(java.sql.Connection con,
			-	int businessTransactionId,
				java.lang.String label)
		t	hrows	DataManagerException
			_	

This method updates the label for a business transaction. This label is displayed on the left-side panel of the transaction analysis view.

Parameters:

con	The database connection to use
businessTransactionId	The ID of the business transaction.
label	The new label.

Throws:

DataManagerException - If database operation fails

getBusinessTransactionId

This method returns the business transaction id for an event

Parameters:

con	The database connection to use
eventId	Event ID of the event

Returns:

int - The business transaction id for the event, or -1 if no business transaction exists

Throws:

DataManagerException - If database operation fails

getEventCount

This method returns the number of events in the database for a particular schema.

Parameters:

con	The database connection to use.
schema	The schema for which to retrieve the event count

Returns:

The event count

Throws:

DataManagerException - If database operation fails

5.3. Class QueryDoc

public class com.bristol.tvision.projectmgr.QueryDoc extends com.bristol.tvision.util.xml.XMLDocument

The QueryDoc class is used as the input query definition into the QueryServices object. The schema the XML document is defined in the file <TVISION_HOME>/config/xmlschema/Query.xsd.

A sample query document is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
  <Query desc="" name="default" timeZone="America/New_York">
        <Group name="MQSERIES">
        <Where name="apicode" negated="false" translateValue="false">
            <XPath>/Event/Technology/MQSeries/@apiCode</XPath>
            <Operator>equal</Operator>
            <Value>8</Value>
            <Value>11</Value>
            <Value>12</Value>
            </Where>
            </Group>
            </Group>
            </Query>
```

The above query searches for events on the XPath

"/Event/Technology/MQSeries/@apiCode", that is the lookup column corresponding to the MQSeries API code for values 8 (MQGET), 11(MQPUT) and 12 (MQPUT1). Note that there is a separate <Group> section for each technology include in the query, and the conditions of all <Group> sections are ORed together for the final query.

```
<?xml version="1.0" encoding="UTF-8"?>
<Query desc="" name="default" timeZone="America/New_York">
<Group name="MQSERIES">
```

The above query searches for WebSphere MQ API "MQGET" events from program name of "amqsput".

An "AND" operation is performed on the two "Where" clauses in the above query, while an "OR" operation is performed on values within the same "Where" clause. To use actual values instead of object ids the attribute "translateValue" has to be set to true in the "Where" clause. The static inner class QueryDoc.WhereClause can be used to construct the QueryDoc document instead of providing an XML document.

5.3.1. Constructors

QueryDoc

```
public QueryDoc()
```

This constructor creates new QueryDoc. The root element 'Query' is created automatically.

QueryDoc

public QueryDoc(java.lang.String name)

This constructor creates new a QueryDoc of a given name. The root element 'Query' is created automatically.

Parameters:

name - The session unique name of this QueryDoc.

QueryDoc

The constructor creates a new QueryDoc from the input queryDoc bytes.

Parameters:

queryName	Name of the query
queryDoc	Document for the query.
modified	Modification status for the query document.

Throws:

ProjectManagerException - Error generating XML document for the query

QueryDoc

public QueryDoc(QueryDoc other)

This copy constructor creates a new QueryDoc from the given input QueryDoc.

Parameters:

other - QueryDoc instance used to create a new QueryDoc from.

QueryDoc

public QueryDoc(boolean transDoc)

This constructor creates a new QueryDoc, if it is passed 'true' it will initialize the query document to use the Transaction XDM Document type. This is required if you wish to query on your transaction classes. By default the QueryDoc uses the Event document type.

5.3.2. Methods

getDocName

public java.lang.String getDocName()

This method gets the name of the query document.

Returns:

Name of query document

setDocName

public void setDocName(java.lang.String str)
This method sets the query name.

Parameters:

str - Name of query document.

setDocDescription

public void **setDocDescription**(java.lang.String desc) This method sets the description string for the query document.

Parameters:

desc - Description string for the query document.

getDocDescription

public java.lang.String getDocDescription()
This method retrieves the description string for the query document.

Returns:

Description string for the query document.

setTimeZone

public void setTimeZone(java.lang.String tzId)
This method sets the time-zone string for the query document.

Parameters:

desc - timezone string for the query document.

getTimeZone

public java.lang.String getTimeZone()

This method retrieves the time-zone string for the query document.

Returns:

The return value is the time-zone string for the query document.

toByteArray

public byte[] toByteArray()

throws ProjectManagerException

This method returns the XML query document as a byte array.

Returns:

The return value is a byte array of the XML document. Returns null on failure.

insertSelect

public boolean **insertSelect**(java.lang.String[] xpaths) This method sets an array of XPath expressions, which form the "SELECT" part of the query.

isModified

public boolean isModified()

Check if query document is changed since the last calling of setClean()

Return:

true, if document is modified.

setClean

public void setClean()

Reset the modified flag.

updateWhereClause

public boolean updateWhereClause(QueryDoc.WhereClause clause)

This method sets the "WHERE" part of the query. It will check the current query group setting. If current query group id is QueryDoc.TECH_ALL, it updates where clauses for all existing technology, otherwise just update the where clause of current selected query group.

Returns:

Parameters:

clause - If there is no "Where" clause of same name in the query document, this clause is added into the document, else the existing "Where" clause is updated.

updateWhereClause

public boolean updateWhereClause(WhereClause clause, int groupId)

Update where clause under given query group. The groupId can be any integer. The following values are reserved for TransactionVision technology; if groupId is QueryDoc.TECH_ALL, it updates where clauses for all existing query groups.

TVisionCommon.TECH_ID_MQSERIES, TVisionCommon.TECH_ID_BTTRACE, TVisionCommon.TECH_ID_SERVLET, TVisionCommon.TECH_ID_JMS, TVisionCommon.TECH_ID_MQIMSBRIDGE, TVisionCommon.TECH_ID_EJB, TVisionCommon.TECH_ID_CICS

Parameters:

clause	where clause
groupId	group id

Return:

true if operation succeeds.

updateBufferClause

public boolean updateBufferClause(BufferClause clause)

Update buffer clause under current query group. If current query group id is QueryDoc.TECH_ALL, it applies on all existing query groups.

Parameters:

clause - buffer clause

Return:

true if operation succeeds.

updateBufferClause

public boolean updateBufferClause(BufferClause clause, int groupId)

Update buffer clause under given query group. The groupID can be any integer. The following values are reserved for TransactionVision technology. If groupId is QueryDoc.TECH_ALL, it updates all buffer clauses under existing query groups.

Parameters:

clause	buffer clause
groupId	group id

Return:

true if operation succeeds.

deleteWhereClauseByName

public void deleteWhereClauseByName(String name)

Delete where clause under current query group. If the current query group is QueryDoc.TECH_ALL, it deletes where clauses from all existing query groups.

Parameters:

name - where clause name

deleteWhereClauseByName

public void **deleteWhereClauseByName**(String name, int groupId) Update where clause under given query group. GroupId can be any integer. The following values are reserved for TransactionVision technology; if groupId is QueryDoc.TECH_ALL, it deletes where clauses from all existing query groups:

Parameters:

name	where clause name
groupId	group id

deleteBufferClause

public void deleteBufferClause()

Delete buffer clause under current query group. If the current query group id is QueryDoc.TECH_ALL, it deletes all buffer clauses from existing query group.

deleteBufferClause

public void deleteBufferClause(int groupId)

Delete buffer clause under give technology. The groupID can be any integer . the following values are reserved for TransactionVision Technology. If groupId is QueryDoc.TECH_ALL, it deletes all buffer clauses from existing query group.

Parameters:

groupId - group id

findWhereClauseByName

public WhereClause **findWhereClauseByName**(String name) Retrieve the where clause of given name under current query group

Parameter:

name - where clause name

Return:

WhereClause instance

findWhereClauseByName

public WhereClause findWhereClauseByName(String name, int groupId)

Retrieve the where clause of given name under given query group. GroupId should be the ID of an existing query group.

Parameters:

name	where clause name
groupId	query group id

Return:

WhereClause instance.

getBufferClause

public BufferClause getBufferClause()

Get buffer clause under current query group

Return:

BufferClause instance

getBufferClause

public BufferClause getBufferClause(int groupId)

Get buffer clause under given query group.

Return:

BufferClause instance

getWhereClauseNames

public String[] getWhereClauseNames()

Get all where clause names under current query group.

Return:

An array of where clause names.

getWhereClauseNames

public String[] getWhereClauseNames(int groupId)

Get all where clause names under given query group.

Return:

An array of where clause names.

isLinearSearch

public boolean isLinearSearch()

Check if query document contains linear search clause.

Return:

true if query document is linear searching.

isBufferSearch

public boolean isBufferSearch()

Check if query document contains buffer clause

Return:

true if there's at least one buffer clause

equals

public boolean **equals**(QueryDoc d) Check if two queries equal or not.

groupCompare

public boolean groupCompare(QueryDoc d, int groupId1, int groupId2) Compare group of different query doc.

printQueryDoc

public void printQueryDoc(OutputStream out)

Dump query document to given output stream.

Parameter:

out - output stream instance.

getCurGroup

public int getCurGroup()
Get current query group id

Return:

Query group ID

setCurGroup

public void setCurGroup(int groupId)

Set current query group id. GroupID can be any integer. The following values are reserved for TransactionVision technologies:

TVisionCommon.TECH_ID_MQSERIES, TVisionCommon.TECH_ID_BTTRACE, TVisionCommon.TECH_ID_SERVLET, TVisionCommon.TECH_ID_JMS, TVisionCommon.TECH_ID_MQIMSBRIDGE, TVisionCommon.TECH_ID_EJB, TVisionCommon.TECH_ID_CICS

Parameter:

groupId – The query group Id.

setCurGroup

public void setCurGroup(String name)
Set current query group by given technology name. QueryDoc will map the name to TransactionVision technology ID.

Parameter:

name - technology name, must be one of the following values:

TVisionCommon.TECH_ID_MQSERIES, TVisionCommon.TECH_ID_BTTRACE, TVisionCommon.TECH_ID_SERVLET, TVisionCommon.TECH_ID_JMS, TVisionCommon.TECH_ID_MQIMSBRIDGE, TVisionCommon.TECH_ID_EJB, TVisionCommon.TECH_ID_CICS

setTechnologyOn

public void setTechnologyOn(boolean on, int techId)

Turn on/off the query for given technology. TechID must be one of the following values:

TVisionCommon.TECH_ID_MQSERIES, TVisionCommon.TECH_ID_BTTRACE, TVisionCommon.TECH_ID_SERVLET, TVisionCommon.TECH_ID_JMS, TVisionCommon.TECH_ID_MQIMSBRIDGE, TVisionCommon.TECH_ID_EJB, TVisionCommon.TECH_ID_CICS

Parameters:

on	flag turn on/off
techId	technology id

isTechnologyOn

public boolean **isTechnologyOn**(int techId) Check if query on given technology is on or off.

getTechnologyNameFromID

public static String **getTechnologyNameFromID**(int id) Get technology name for given tech ID.

getTechnologyDescFromID

public static String **getTechnologyDescFromID**(int id) Get technology display string for given tech ID.

getTechnologyIDFromName

public static int **getTechnologyIDFromName**(String name) Get technology id from given technology name.

5.4. Class QueryDoc.WhereClause

This is an inner static class in the class QueryDoc. It is a utility class that helps to define the where condition of the query. This condition is the matching criteria for which events should be retrieved from the database.

5.4.1. Fields

- name public java.lang.String name Name of the where clause
- negated public boolean negated Whether the where clause has "not" condition
- *xpath* public java.lang.String **xpath** XPath for the where clause
- operator public java.lang.String operator Operator for the where clause
- values public java.lang.String[] values Values for the where clause
- isLinearCond
 public boolean isLinearCond
 Specifies whether the "Where" clause is a linear search condition.
- valueType public java.lang.String valueType
- TYPE_BIN public static final java.lang.String TYPE_BIN
- TYPE_TEXT public static final java.lang.String TYPE_TEXT
- codePage public java.lang.String codePage

5.4.2. Constructors

QueryDoc.WhereClause

public QueryDoc.WhereClause()

This constructor creates an empty object.

QueryDoc.WhereClause

This constructor creates a "WhereClause" object using the given data.

Parameters:

Name - Name of the "Where" clause. Negated - Whether the where clause has "not" condition. XPath - XPath of "Where" clause. Operator - Operator of "Where" clause. Values - Values of "Where" clause. IsLinearCond - True if "Where" clause is a linear search condition.

QueryDoc.WhereClause

java.lang.String Operator, java.lang.String[] Values, boolean IsLinearCond, java.lang.String valueType, java.lang.String codePage)

This constructor creates a "WhereClause" object using the given data.

Parameters:

Name - Name of the where clause Negated - Whether the "Where" clause has "not" condition XPath - XPath of "Where" clause Operator - Operator of the "Where" clause Values - Values of the "Where" clause IsLinerCond - True if "Where" clause is a linear condition valueType - Either of the following values, it's just for query edit page display QueryDoc.WhereClause.TYPE_BIN QueryDoc.WhereClause.TYPE_TEXT codePage - The character code page

$\label{eq:QueryDoc.WhereClause} QueryDoc.WhereClause$

This constructor creates a "WhereClause" object using the given data.

Parameters:

Name - Name of the where clause
Negated - Whether the "Where" clause has "not" condition
XPath - XPath of "Where" clause
Operator - Operator of the "Where" clause
Values - Values of the "Where" clause
IsLinerCond - True if "Where" clause is a linear condition
valueType - Either of the following values. For display purpose only.
QueryDoc.WhereClause.TYPE_BIN
QueryDoc.WhereClause.TYPE_TEXT
codePage - The character set code page, that is used when converting the hexidecimal
value string into text
techID - Technology ID

5.4.3. Methods

equals

public boolean equals(QueryDoc.WhereClause c)

This method compares two WhereClause objects.

Parameters:

c - another instance of WhereClause

Returns:

true if two are considered be equal

5.4.4. Example

The sample code below creates a query document with a "WhereClause" and a "SelectClause" using the methods updateWhereClause and insertSelect. The query condition is named "mqputget" and specifies to match all MQPUT, MQPUT1 and MQGET APIs. The data fetched out of the database is specified by the selects String array and contains the XPath expressions for the fields entry time, exit time, API code, host id, program id, program instance id and sequence number.

XPathConstants.PROGINST_ID,
XPathConstants.SEQUENCE_N0 };

```
qdoc.updateWhereClause(clause);
qdoc.insertSelect(selects);
```

5.5. Interface Query

public interface com.bristol.tvision.datamgr.query.Query This interface provides the functionality to run a query. This object is obtained from methods in the QueryServices class.

5.5.1. Methods

execute

public Cursor **execute**()

throws DataManagerException

This method executes the query and returns a Cursor object to be iterated over.

Throws:

DataManagerException - If executing the query fails

close

public void close()
 throug DataManagerException

throws DataManagerException

This method closes the query and releases the database resources. The query can not be executed again once close has been called.

Throws:

DataManagerException - If release of the database resources fails

cancel

public void **cancel**()

throws DataManagerException

This method can be called from a different thread to cancel the current query execution.

Throws:

DataManagerException - If the cancel fails

5.6. Interface Cursor

public interface com.bristol.tvision.datamgr.query.Cursor

The cursor interface is used to iterate over data returned by a query.

5.6.1. Methods

getRowCount

public int getRowCount()

This method returns the number of table rows in the query result, or -1 if this feature is not supported

Returns:

The number of rows

getColumnCount

public int getColumnCount()

This method returns the number of columns in the query result

Returns:

The number of columns

getColumnDescription

public java.lang.String getColumnDescription(int index)

This method returns the column description for the specified column. The index of the first column is 1.

Parameters:

index - The index of the column

Returns:

The column name

getColumnName

public java.lang.String getColumnName(int index)

This method returns the database column name for the specified column. The index of the first column is 1.

Parameters:

index - The index of the column

Returns:

The column name

getRow

public int getRow()

throws DataManagerException

This method returns the current row for this cursor

Returns:

The current row, or 0 if there is no current row

getValue

This method returns the value of the column as a String value. The index of the first column is 1.

Parameters:

index - The index of the column

Returns:

The value of the column, converted into a String

Throws:

DataManagerException - If getting the value from the underlying ResultSet fails

getValue

This method returns the value of the column as a String value (converted by the type conversion service). The index of the first column is 1.

Parameters:

index	The index of the column
convSvr	The type conversion service to use.

Returns:

The value of the column, converted into a String

Throws:

DataManagerException - If getting the value from the underlying ResultSet fails

getIntValue

This method returns the value of the column as an integer value. The index of the first column is 1.

Parameters:

index - The index of the column

Returns:

The value of the column, converted into a integer

Throws:

DataManagerException - if getting the value from the underlying ResultSet fails

getValue

This method returns the value of the column as a String value. The column is identified by a key (XPath for XDM columns).

Parameters:

key - The key for the column

Returns:

The value of the column, converted into a String

Throws:

DataManagerException - If getting the value from the underlying ResultSet fails

getValue

This method returns the value of the column as a String value (possibly converted by the type conversion service). The column is identified by a key (XPath for XDM columns).

Parameters:

key	The key for the column
convSvr	The type conversion service to use

Returns:

The value of the column, converted into a String

Throws:

DataManagerException - If getting the value from the underlying ResultSet fails

getIntValue

This method returns the value of the column as an integer value. The column is identified by a key (XPath for XDM columns).

Parameters:

key - The key for the column

Returns:

The value of the column, converted into a integer

Throws:

DataManagerException - if getting the value from the underlying ResultSet fails

getValueMap

This method returns a Map object which contains a mapping from XPath to current column value, or null if this feature is not supported.

Parameters:

convSvr - The type conversion service to use.

Returns:

A Map object containing the values of the current row

Throws:

DataManagerException - if getting the values from the underlying ResultSet fails

wasNull

public boolean **wasNull**()

throws DataManagerException

This method reports whether the last column read with getValue() or getIntValue had a value of SQL NULL

Returns:

true if the last column value read was SQL NULL and false otherwise

Throws:

DataManagerException - if accessing the ResultSet fails

next

public boolean next()

throws DataManagerException

This method moves the cursor forward one row from its current position. A Cursor is initially positioned before the first row, calls to next() advance the cursor to the next row.

Returns:

true if the new current row is valid; false if there are no more rows

Throws:

DataManagerException - if moving the cursor in the underlying ResultSet fails

previous

public boolean **previous**()

throws DataManagerException

This method moves the cursor backwards one row from its current position. A Cursor is initially positioned before the first row, calls to previous() advance the cursor to the previous row.

Returns:

true if the new current row is valid; false if there are no more rows

Throws:

DataManagerException - if moving the cursor in the underlying ResultSet fails

absolute

public boolean absolute(int row)

throws DataManagerException

This method moves the cursor to an absolute row position.

Parameters:

row - The row to position on

Returns:

true if the new current row is valid; false if cursor is not positioned on valid row

Throws:

DataManagerException - if positioning the cursor in the underlying ResultSet fails

close

public void close()

throws DataManagerException

This method closes the cursor and all with the cursor associated database resources

Throws:

DataManagerException - if closing the underlying JDBC resources fails

5.7. Class DataManagerException

public class DataManagerException extends TVisionException

This exception class contains errors from the DataManager package.

5.7.1. Constructors

DataManagerException

public DataManagerException()

This constructor creates new DataManagerException without a detail message string.

DataManagerException

public DataManagerException(java.lang.Throwable t)

This method constructs a DataManagerException with the specified embedded Throwable.

DataManagerException

public DataManagerException(java.lang.Object[] args)

This method constructs a DataManagerException with the specified logging arguments.

Parameters:

args - the logging arguments

DataManagerException

This method constructs a DataManagerException with the specified embedded Throwable and the specified logging arguments.

Parameters:

t - the exception to chain args - the logging arguments

5.7.2. Methods

getSQLException

public java.sql.SQLException getSQLException()

This method returns the embedded exception as a SQLException if it is an instance of SQLException, null otherwise.

Returns:

The SQLException, or null if the embedded exception is not an instance of SQLException

is Unique Violation Exception

public boolean isUniqueViolationException()

Returns true if the embedded exception is a SQLException indicating a violation of an unique constraint, false otherwise.

Returns:

true if exception is a unique constraint violation

isStringTruncationException

public boolean isStringTruncationException()

This method returns true if the embedded exception is a SQLException indicating that a string has been truncated because it is too long for the column, false otherwise.

Returns:

true if exception is a truncation exception

isComplexityException

public boolean isComplexityException()

This method returns true if the embedded exception is an SQLException indicating that the executed SQL statement was too complex, false otherwise.

Returns:

true if exception is a SQL complexity violation

$is Operation Canceled {\sf Exception}$

public boolean isOperationCanceledException()

This method returns true if the embedded exception is an SQLException indicating that the executed SQL query has been canceled, false otherwise.

Returns:

true if exception is a SQL cancellation violation

6. Extending the User Interface

6.1. Writing TransactionVision Reports

TransactionVision reports are essentially JSPs and servlets which make queries into TransactionVision project tables to extract, analyze and present data collected. These reports may either use the QueryServices classes or make direct JDBC SQL calls to perform queries. The presentation of the reports may be in any browser support technology such as HTML, SVG or Java applets.

The TransactionVision report framework is based on the Model-View-Controller (MVC) design pattern. When creating new reports, you must code the "View" and "Model" parts of the framework, then hook them into the report framework. To facilitate report development, the TransactionVision report framework provides the following:

- A library of custom JSP tags
- Interfaces for handling report generation and report parameters

The TransactionVision report framework also provides a default implementation of the interfaces. You are encouraged to use the default implementation and override only those aspects that are unique to your report. The TransactionVision installation provides a set of sample reports; use them as a reference when creating your own reports.

To creat a new TransactionVision report, you must do the following:

- 1. Identify report parameters.
- 2. Create a new implementation of the **IReportData** interface, or derive a class from the **BaseReportBean**. Create get/set methods for each parameter that has to be updated by the framwork, with values from either the HTTP request or from a saved database record.
- 3. Create a new implementation of the **IReportAction** interface to generate the report. If additional actions are defined for the report, then provide an implementation of **IReportAction** for each of these as well. The DefaultReportActionImpl Java class handles most of the operations known to the framework; you are expected to override only the **CreateReport** action.
- 4. Write up a new JSP to display the report. The JSP custom tag library assists in JSP creation.
- 5. Add the report to the <TVISION_HOME>/config/ui/reports.xml file.

Note: trace and debug messages from the report framework and tag classes are written to the UI_TRACELOG under the category "ReportTrace."

The following diagram shows an overview of how the report framework handles an incoming report request:



6.1.1. Report Interfaces

The **IReportData** and **IReportAction** interfaces enable the report framwork to handle reports in a uniform way. Each report can have one implementation of **IReportData** and one or more **IReportAction** implementations. All report interfaces may be implemented in either one Java class or a separate class for each interface.

6.1.1.1. IReportData

}

This interface encompasses the most common operations executed on a specific report's parameters. The following operations are currently identified:

- Creating an XML document holding a list of parameters and their current values for easy storage in a database.
- Extracting the parameters from either a submitted HTTP request object or an XML document and updating the report bean properties.

The com.bristol.tvision.ui.report.DefaultReportDataImpl class provides a default implementation of this interface. You may either derive from this class or provide your own implementation of this interface.

Note: The framework assumes and depends on the report bean to provide the get/set function for each individual parameter. The prototypes of these functions are as follows:

void setXXXXX(String value); String getXXXXX();

If you do not provide a get/set method for any parameter, the default implementation of the framework is unable to handle that parameter.

For an example, see the SLA Analysis Report. This report contains two parameters: ReportDate and SelectedTxnClasses. The bean

com.bristol.tvision.report.samples.performance.SLAAnalysisReportBean derives from the DefaultReportDataImpl class and also provides set/get methods for the two parameters.

6.1.1.2. IReportAction

```
public interface IReportAction {
    public void perform(HttpServletRequest request)
        throws UIException;
}
```

This interface represents a single atomic operation that can be operated on a report. You must provide at least one implementation of this interface to handle creation of the report itself. The following other operations are also recognized by the framework:

- Saving the report parameters to the database. This helps end-users generate reports with just a click of a URL. Multiple configurations may be saved for each report.
- Retrieving the previously saved report parameters and generating the report.
- Deleting obsolete saved report parameter record(s) from the database.

The operations of the framework are CreateReport, SaveParameters,

DeleteParameters, and GetParameters. The

com.bristol.tvision.ui.report.DefaultActionImpl class provides a default implementation for these framework actions as follows:

Action	Default Implementation
CreateReport	Does nothing.
SaveParameters	Calls an IReportData function to wrap all the report parameters in an XML document, then saves the document in the REPORT_PARAMETERS table.
GetParameters	This action is executed when a saved report configuration link is clicked from the list of reports page. It extracts the saved record from the databse and passes it onto an IReportData function to update the report bean's properties then generate the report using these parameters.
DeleteParameters	Deletes an instance of the saved report parameters record and redisplays the list of available reports.

A report can override any or all these actions by providing a different IReportAction implementation.

You may also define new actions for your reports. The new actions must be handled by the report defining them. One implementation of this interface can handle one action, or it can be made to handle multiple actions by identifying the action by anme and handling it accordingly.

For examples, see the following reports installed with TransactionVision:

- The SLA Analysis Report bean implements the IReportAction interface for CreateReport.
- The Dashboard Report bean derives from the BaseReportBean class and implements the IReportAction::perform() method for generating the report.
- The com.bristol.tvision.ui.report.framework.ReportListBean report defines a new action—SelectReport—and the bean provides a common implementation for the CreateReport and SelectReport actions.

6.1.1.3. BaseReportBean

This is an abstract class extending DefaultReportDataImpl. It includes some additional methods that are common to most reports, such as handling the reporting time period parameter (From and To dates) and obtaining a database connection handle.

You may create new report beans by deriving from this class. At a minimum, the new bean must provide set/get methods for each report parameter and provide an implementation of the IReportAction::perform() method to generate the report data.

6.1.2. TransactionClass

The class com.bristol.tvision.datamgr.dbtypes.TransactionClass is provided to to get access to the transaction classification definitions. Its interface is as follows; an example of its use can be found in the report samples:

public Integer[] getClassIds()

Returns an Integer Array of ClassIds

public Map getClassAttributes(int classId)

Parameters: classId - The class id to retrieve the attributes for

Returns a map of all this classes attributes

public String getClassAttributeValue(int classId, String xpath)

Parameters: classId - The class id to retrieve the attributes for

xpath - The specific xpath of the attribute you want to lookup

Returns the attribute value, or if it doesn't exist, null.

public int getCount()

Returns the numbers of definitions for the given project.

6.1.3. JSP Custom Tag Library

TransactionVision provides a JSP custom tag library that you may use while writing the report JSP. Most of the tags are for creating the HTML form that obtains report parameters from the user generating the report.

Two basic report tags are required for every report JSP within the report framework: the report tag and the form tag. The following example shows these tags:

```
<tvreport:report>
<tvreport:form name = "reportForm" >
[...] Put all your form elements here.
[...] Put all your button definitions here - See ActionBuildTag.
</tvreport:form>
[...] Put all display related JSP code here.
</tvreport:report>
```

6.1.3.1. The Report Tag

The report tag (<tvreport:report>) frames the entire contents of your report. All contents of your JSP should occur within the contents of this tag. This tag sets up the basic infrastructure that the report needs. Most importantly, it sets a number of page

Variable	Description
reportRequest	This object is an instance of the ReportRequestParms (see javadoc) object for this report.
reportData	This object is the instance of your data bean for this report.
categoryName	category name of current report
subCategoryName	subcategory name of current report
reportName	name of current report

context attributes that contian relevant information about the currently running report. The following table describes these variables, which are accessible from the JSP:

Important:

Only access these variables **within** the report tag. For example, trying to access them before the <tvreport:report> start tag results in an error.

6.1.3.2. The Form Tag

The other required tag is the form tag. The body of this tag contains all the HTML you want displayed to set the parameters for configuring what the report will display and buttons for deciding what action the report should perform. This tag creates an HTML form tag and inserts a number of hidden form elements that contain information for the framework on how to handle this report when the form is submitted.

Anything occuring after the end of the form tag (after </tvreport:form>) is part of display portion of the report. These contents will only be displayed once the report has been run; it will not show on the initial entry to the report. Place any representation of your data in this section.

6.1.3.3. Tag Reference

This section provides reference information for TransactionVision report tags.

Form

This tag provides an HTML form for holding all the report parameters. The framework itself uses hidden HTTP form fields for identifying the report in the current request, user name, etc. This tag automatically generates the HTML for these hidden variables. When this form is submitted to the report framework, all fields submitted from the form are checked against the reports data bean. If a get/setter method matching the name of the form field is found, the value will be read and saved into the bean.

Important:

This tag has changed from the previous TransactionVision version; be sure to read the 'Migrating Form Tag' section to learn about these differences.

Attributes of this tag are:

name	The name corresponds to the name of the form. You can then use that name to access the form via JavaScript.
onValidate	(optional) Name of a JavaScript function to perform client

side validation of data entered into a reportform. The javscript function specified in onValidate will be called whenever the user submits the form. If the specified function returns false, the form will not be submitted.

Example:

Button

By default a report form will contain no buttons. Use the button tag to add a desired button.

This tag has four attributes: type, label, callback and action.

Type

The type attribute can be used to create the standard report buttons, and will override the other attributes. The three standard buttons are the Generate Report, Print Preview, and Save Settings buttons. They can be created as shown in the following example. These buttons can then be added, or removed as desired from your report.

```
<tvreport:button type='<%=ActionButtonTag.GENERATE_REPORT%>'/>
<tvreport:button type='<%=ActionButtonTag.PRINT_PREVIEW%>'/>
<tvreport:button type='<%=ActionButtonTag.SAVE_REPORT%>'/>
```

If you wish to further customize your buttons, the other three attributes allow you to control this.

Label

The text on the button. The label attribute specifies what text will appear in the button.

callback

JavaScript callback to call when button is pressed. The callback names an optional JavaScript function that you want called to do some processing before the form is submitted. The default callbacks for the standard buttons are generateReport, printPreviewReport, and saveReport respectively. These can be specified in most cases, but if you have further special tasks you want done you can write your own callback.

action

Report Action to initiate. The action is the report framework action this button initiates. If you have extended your report bean to support additional action types, you can use this field to create a button for doing this action.

The below shows an example of creating a button called 'Replay Events' that when pressed will generate the report.

<tvreport:button label="Replay Events" callback="generateReport" action="CreateReport"/>

Multiselect

This is a helper tag that provides some useful features. This tag creates a list of checkboxes that are all associated with a single parameter as given through the name attribute. This can be useful if you have a number of dynamically generated checkboxes, and thus your bean will not have set/get methods predefined. In this case all the option tags associated with this multiselect concatenate their values, delimeted by a ';' to the named value. The report bean can then extract these values from a single parameter. This is used commonly in the shipped TransactionVision reports to list all the transaction classes in a report. (see the following example). For lists of items that can have a large set of unique values, this tag also creates this checkboxes within a scrollable area.

name	Required, is the name of the variable in the bean where the data will be sent.
initValue	(Optional) The checkbox matching this name will be selected when this tag is initialized. If its not set, all checkboxes will be initialiy selected.
Listsize	(Optional) Controls the minimum size of the list of checkboxes to display before a scrollbar appears.

This tag has two attributes:

Example:

```
<tvreport:multiselect name="selectedTxnClasses" listsize="10">
    <% for (int i = 0; i < numTxnClasses; i++) { %>
        <tvreport:option value='<%=txnClassNames[i]%>'/>
        <% } %>
    </tvreport:multiselect>
```

6.1.3.4. Migrating the Form Tag from TransactionVision 4.0

A new tag called form replaces the parameterform of TransactionVision 4.0. This tag has some changes from the old tag. The goal of the new tag is to make the use of reports more flexible and in the hands of the user. The tag no longer makes any assumptions on how you want your report form to look. The old form tag forced the report form to use a table layout with a certain look. The new form tag no longer does this; it leaves the layout and look of the fields up to the web page developer. This allows you to use standard html form tags to create your form instead of needing to learn how to use cutom TransactionVision tags. One effect of this is that the form tag will no longer creates an overall table that all the controls appear in. In migrating your reports, if you want them to maintain the similar table format, you need to add the appropriate table tags to the HTML (see the following example). Instead of a label attribute (the label is now set explicitly by you in your report HTML), the form has a name attribute.

New jsp code:

```
<tvreport:form name="reportForm">
```

```
Report Parameters

A Checkbox

<input type="checkbox" name="mychkbox" checked/>

<t/tr>

</tvreport:tvform>
Old jsp code:
<tvreport:parameterform label="Report Parameters">
<tvreport:parameterform label="Report Parameters">
<tvreport:checkboxparm label="Report Parameters">
<tvreport:checkboxparm</tvreport:checkboxparm>
<tvreport:checkboxparm>
</tvreport:checkboxparm>
</tvreport:checkboxparm>
</tvreport:parameterform>
</tvreport:parameter
```

dateParm

This tag displays the reporting time period. The time period is relative and has values like "Today", "Yesterday", "This Year", "Last Year" etc. There is also a "DateRange..." value which when selected displays two input fields for providing the From and To dates of the range.

Attributes of this tag are:

name	Hardcoded to "ReportDate".
Label	<i>Required.</i> This is the label or static text that appears to the left of this select field.
htmlAttrs	<i>Optional.</i> All the HTML <select> tag attributes for which there is no corresponding attribute in this tag goes here.</select>

Example:

6.1.3.5. Deprecated Tags

The tags in this section were available in TransactionVision 4.0; however, they may not work in future releases.

ReportParmInputTag

This tag creates an HTML <input> element. Along with the normal HTML it generates additional code to place this element within a table with the "label" on the left and the element itself on the left. Apart from these, it handles additional things specific to the report framework functionality. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

Туре	Required. Can be any valid value of <input/> HTML tag like "text", "button", etc.
Name	<i>Required.</i> Identifier for this UI element. Note: If this is a report parameter then provide the same name as the bean's corresponding property.
Label	<i>Optional.</i> This is the label or static text that appears to the left of this text field, button etc. Default is NULL.
Value	Optional. Default value for this field. Default is NULL.
Mandatory	<i>Optional.</i> If set to "true", this makes sure that the text field is always filled up before submitting the form. Makes sense for text fields only. Default is "no".
Serialize	<i>Optional.</i> If set to "true", this form field is saved to the database when the "Save Parameters" button is clicked. Default is "yes".
	Note: This requires the 'name' of this field to be same as the report bean's property or the very least the report bean should provide a get/setXXXX() method where XXXX is the name of this field.
htmlAttrs	<i>Optional.</i> All the HTML <input/> tag attributes for which there is no corresponding attribute in this tag goes here.
Example:	
[]	
<tvrej< td=""><td>type="text"</td></tvrej<>	type="text"
	name="QueueManager"

[...]

ReportParmCheckboxTag

This tag creates an HTML <input type="checkbox"> element. Along with the normal HTML it generates additional code to place this element within a table with the "label" on the left and the element itself on the left. Apart from these it handles additional things specific to the report framework functionality. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

name

Required. Identifier for this UI element.

value="<%=queueManager%>"/>

	Note: If this is a report parameter then provide the same name as the bean's corresponding property.
Label	<i>Optional.</i> This is the label or static text that appears to the left of this checkbox. Default is NULL.
Value	<i>Optional.</i> Default value for this field. This is the value returned when the checkbox is checked. Default is NULL.
Serialize	<i>Optional.</i> If set to "true", this form field is saved to the database when the "Save Parameters" button is clicked. Default is "yes".
	Note: This requires the 'name' of this field to be same as the report bean's property or the very least the report bean should provide a get/setXXXX() method where XXXX is the name of this field.
htmlAttrs	<i>Optional.</i> All the HTML <input/> tag attributes for which there is no corresponding attribute in this tag goes here.
Checked	<i>Optional.</i> If "true" the checkbox is checked. Default value is "false".
Example:	
[]	
<tvreport:checkb< td=""><td>poxparm</td></tvreport:checkb<>	poxparm
	label="Include local transactions:"
	name="wantLocalTxns">
<tvreport:ch< td=""><td>necked></td></tvreport:ch<>	necked>
<%= (rec	<pre>quest.getParameter("wantLocalTxns") != null) ?</pre>
	"true" : "Ialse" %>
<td>boxparma</td>	boxparma
<pre>>/ cvreborc.cliecy</pre>	

ReportParmSelectTag

This tag creates an HTML <select> element. Along with the normal HTML it generates additional code to place this element within a table with the "label" on the left and the element itself on the left. Apart from these it handles additional things specific to the report framework functionality. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

Name	<i>Required.</i> Identifier for this UI element. Note: If this is a report parameter then provide the same name as the bean's corresponding property.
Label	<i>Optional.</i> This is the label or static text that appears to the left of this select field. Default is NULL.
Serialize	<i>Optional.</i> If set to "true", this form field is saved to the database when the "Save Parameters" button is clicked. Default is "yes".

	Note: This requires the 'name' of this field to be same as the report bean's property or the very least the report bean should provide a get/setXXXX() method where XXXX is the name of this field.
htmlAttrs	<i>Optional.</i> All the HTML <select> tag attributes for which there is no corresponding attribute in this tag goes here.</select>
Example:	
[]	
<tvreport:se< td=""><td>lectparm label="Sort By:" name="sortByCol"></td></tvreport:se<>	lectparm label="Sort By:" name="sortByCol">
<tvreport:< td=""><td>option value='Name'></td></tvreport:<>	option value='Name'>
	Component Name
<td>t:option></td>	t:option>
<tvreport:< td=""><td>option value='Min'></td></tvreport:<>	option value='Min'>
	Minimum Latency Time
<td>t:option></td>	t:option>
<td>selectparm></td>	selectparm>
[]	

CheckboxTag

This tag creates an HTML <input type="checkbox"> element. Along with the normal HTML it handles additional things that are specific to the report framework functionality. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

Name	<i>Required.</i> Identifier for this UI element. Note: If this is a report parameter then provide the same name as the bean's corresponding property.
Value	<i>Optional.</i> Default value for this field. This is the value returned when the checkbox is checked. Default is NULL.
htmlAttrs	<i>Optional.</i> All the HTML <input/> tag attributes for which there is no corresponding attribute in this tag goes here.
Checked	<i>Optional.</i> If "true" the checkbox is checked. Default value is "false".

Example:

```
[...]
<tvreport:checkbox name="txnClass"
htmlAttrs=" onclick='updateSelectedTxnClasses(false)'"
/>
[...]
```

CheckedTag

This tag works only with a CheckBoxTag or a ReportParmCheckboxTag. It either checks or unchecks the checkbox. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

None

FormTag

This tag provides an HTML form. It provides some basic validation of the form fields. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

Name	Required. Name of the HTML form.
Action	Required. Action to be performed when this form is submitted.
Method	Required. GET or POST http method.

Example:

```
[...]
    <%@ taglib uri="/tvReportTags" prefix="tvreport" %>
    <tvreport:form name="myForm" action="foo.jsp" method="GET">
        [...]
</tvreport:form>
[...]
```

InputTag

This tag creates an HTML <input> element. When this tag is used with FormTag or ReportParmFormTag then some basic data validation is done. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

type	<i>Required.</i> Can be any valid value of <input/> HTML tag like "text", "button", etc.
name	<i>Required.</i> Identifier for this UI element. Note: If this is a report parameter then provide the same name as the bean's corresponding property.
Value	Optional. Default value for this field. Default is NULL.
Mandatory	<i>Optional.</i> If set to "true", this makes sure that the text field is always filled up before submitting the form. Makes sense for text fields only. Default is "no".
Serialize	Optional. If set to "true", this form field is saved to the

	database when the "Save Parameters" button is clicked. Default is "yes".
	Note: This requires the 'name' of this field to be same as the report bean's property or the very least the report bean should provide a get/setXXXX() method where XXXX is the name of this field.
htmlAttrs	<i>Optional.</i> All the HTML <input/> tag attributes for which there is no corresponding attribute in this tag goes here.
Example:	
[]	
	<tvreport:input <="" td="" type="text"></tvreport:input>

[...]

OptionTag

This tag works only with a SelectTag or a ReportParmSelectTag. It identifies one of the available options in the select list. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

Value

This is the value sent to the report framework which this option is selected.

Example:

SelectTag

This tag creates an HTML <select> element. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

Name	<i>Required.</i> Identifier for this UI element. Note: If this is a report parameter then provide the same name as the bean's corresponding property.
Serialize	<i>Optional.</i> If set to "true", this form field is saved to the database when the "Save Parameters" button is clicked. Default is "yes".

as the

	Note: This requires the 'name' of this field to be same as t report bean's property or the very least the report bean should provide a get/setXXXX() method where XXXX is the name of this field.
htmlAttrs	<i>Optional.</i> All the HTML <select> tag attributes for which there is no corresponding attribute in this tag goes here.</select>
Example:	
- []	
<tvreport:se< td=""><td>elect name="sortByCol"></td></tvreport:se<>	elect name="sortByCol">
<tvreport:< td=""><td>option value='Name'></td></tvreport:<>	option value='Name'>
-	Component Name
<td>rt:option></td>	rt:option>
<tvreport:< td=""><td>option value='Min'></td></tvreport:<>	option value='Min'>
_	Minimum Latency Time
<td>rt:option></td>	rt:option>
<td>:select></td>	:select>
[]	

SelectedTag

This tag works only with a SelectTag or a ReportParmSelectTag. It identifies the selected option. This tag is deprecated; it may not work in future releases.

Attributes of this tag are:

None

Example:

```
[...]
<tvreport:selectparm label="Sort By:" name="sortByCol">
  <tvreport:option value='Name'>
                     Component Name
   </tvreport:option>
  <tvreport:option value='Min'>
                     Minimum Latency Time
   </tvreport:option>
   <tvreport:selected>
       <%=request.getParameter("sortByCol")%>
    </tvreport:selected>
 </tvreport:selectparm>
 [...]
```

6.1.3.6. Report Example

This section provides an overview of the steps that can be followed to get a dashboard style report with controls that update in real-time, and how to accomplish this effect using the Report Framework.

The Transaction Scorecard Dashboard report is an extension of the 'How are my transactions performing' report. Instead of specifying a date range, the dashboard version of the report will automatically refresh itself by the given interval and show the latest data for the past N minutes. The steps taken to convert from the fixed time report to an updating report should provide insight into how this process could be applied to

other reports to accomplish a similar effect. This mechanism can be used primarily if you have an SVG control that can be independently updated.

The basic mechanism for auto-updating the reports SVG control exists within the following JSP code:

```
window.setTimeout("Update()",60000);
</script>
```

In a nutshell, the way this javascript function works is that it will periodically cause the object referenced, in this case 'gauge0' to reload its SVG content.

The next step becomes, how do we fit this within the report framework? What our report needs to do is be able to handle a special request, and return data in SVG content when this request is made. Thus the first step in this process is to extend our report's definition to configure this action.

In Reports.xml, we define a new action called 'Dashboard'. The important part about this definition is the 'redirectUrl' attribute we give this action. What this does is tells the report framework that after this bean has done it processing (querying the database, constructing the neccassary prerequisites for this SVG's construction), it is to forward this request to another servlet which will convert that data into an SVG control/graph. This special redirection is needed because the default behavior of the report framework is to send the browser back to the main jsp defined for this report.

Report.xml example xml:

```
<Report name="myreport.report" title="A Dashboard report"
layoutfile="/reports/dashboard.jsp"
dataclass="my.dashboard.reportbeanclass">
<ShortDesc>"A Dashboard report</ShortDesc>
<Description> ... </Description>
<Action name="CreateReport"
class="my.dashboard.reportbeanclass" />
<Action name="Dashboard"
class="my.dashboard.reportbeanclass"
redirectUrl="/SvgGeneratorServlet"/>
```

</Report>

Now, in the ReportBean itself we need to do a couple things differently than the standard simple report. First, you must distinguish between whether this is a dashboard action or the standard report generation action. Because the dashboard generation request and the report creation request are now split into separate calls, there must be some way of communicating the options selected for this report to the dashboard request. For example, perhaps what transaction class to track, or any other criteria selected from the form that the bean needs to generate all its data. These configuration settings can be saved into the session so that when a subsequent dashboard request comes in the configuration can be retrieved from there.

Additionally, before completing the dashboard action should save the appropriate results so that out SvgGeneratorServlet knows how to generate its image. The below pseudo code shows how this might be done.

ReportBean example code:

```
ReportRequestParms reportRequest = (ReportRequestParms)
request.getAttribute(ReportConstants.TVREPORT_REQUEST);

if (reportRequest.getActionName().equals("Dashboard")) {
    Retrieve report configuration parameters
    Generate my data here
    Save graph definition where SvgGeneratorServlet can
retrieve it.
    } else { // Handle standard report actions
    Save report configuration parameters into session
}
```

The data from the report bean could be passed via a request object and the SvgGeneratorServlet in this case might do something like the below to generate the svg contents.

Note that in the case of a PopChart SVG control, such as those used by the current TransactionVision report samples, TransactionVision already provides a servlet that can do exactly this. By placing the appropriately configued PCISLibEmbedder object into the PopChartServlet.OBJECT_NAME request attribute, and setting your redirectUrl to "/PopChartServlet" you will generate a popchart SVG control.

Now go back to the initial JSP code sample above to tie everything together. You've now defined all the report framework steps needed for this new action, and only need to

complete the hookup in the JSP. For this a helper function of the BaseReportBean will be used - getRequestUrl(). This method will generate the appropriate base url string that will initiate the passed in action, in this case 'Dashboard'. Now everything should be ready to go.

Several implementation details of the report itself have been omitted so that this example can be explained succinctly. To see a complete sample that uses the steps above, take a look at the Transaction Scorecard Dashboard.

6.1.4. Adding a Report to the Framework

All reports are configured and defined in the file

<TVISION_HOME>/config/ui/Reports.xml. When a report is added to this file, a link to the report becomes available in the "Reports" page of the TransactionVision user interface.

The Reports.xml file defines one or more reports. Each report definition includes information such as the report title, description, a URL to the report, and the report categories the report belongs to.

A report category is a logical grouping of reports. TransactionVision groups together links to reports that belong to a category in a table in the Reports page. A ReportCatagory also defines the group of reports that are linked to a particular user right. A ReportCategory's name corresponds to the name specified in a user's LDAP settings. If this name appears in the user's report LDAP settings then that user will be able to view this group of reports. A user whose list of report groups contains the value of "allGroups" can see all available reports.

The ReportCategory and ReportSubcategory require both a name and a title.

The following is a sample from the Reports.xml file. A category called "Performance Analysis & Problem Resolution." This category contains a "Reports" subcategory, which contains the SLA Analysis Report. The link text for this report is "Amy I meeting my Service Level Agreement (SLA) availability and response requirements?" The URL for the report is /reports/performance/SLAAnalysis.jsp.

```
time and availability service level analysis for the
current project. You can specify a maximum
response time for the transactions and a minimum
transaction volume for the time period interval.
</Description>
</Report>
</ReportSubCategory>
</ReportCategory>
```

6.1.4.1. Required Configuration Information

You must provide the following configuration information for each report in Reports.xml:

name The name of the report. A report name must be unique within the subcategory in which it is defined.

layoutfile The location of the JSP for displaying the report.

Note: To add reports created for a release of TransactionVision prior to release 4.0, specify the required configuration information for the report as follows:

<Report name="myOld.report" layoutfile="old.jsp" />

You may also add optional information such as a title and description to provide more information about these reports.

6.1.4.2. Optional Configuration Information

In addition to the name layoutfile, you may also specify a number of optional attributes for each report.

title

The title is displayed at the top of the report (above the parameters form), if it is provided. The default value is NULL, menaing that no title wil be displayed.

dataclass

Specifies the Java class providing an implementation of the IReportData interface. For convenience, the framework provides a default implementation of IReportData in BaseReportBean. The default implementation takes care of assigning values to the parameters from the data in the HTTP request or the retrieved database record. IT can also generate and XML document encompassing all the report parameters so that they can be saved to the database easily.

If this attirbute is not provided, the framwork renders the provided JSP and assumes that the JSP knows how to extract the parameters and generate the report, similar to the way TransactionVision reports were written prior to release 4.0.

Action

Defines the various actions that can be performed on the report. Each operation one the report is identified by a <Action> tag. Most common action is "CreateReport".

Each action is an implementation instance of IReportAction. One java class can implement more than one action. Each action is uniquely identified by the **name** attribute. Each name should be accompanied by an IReportAction **class**.

Some of the actions may also need to perform other actions. For example, after getting a saved parameter record from the database, the report has to be created with these retrieved parameters and hence we need to call "CreateAction". This redirection is achieved by adding **redirect** attribute to the <Action> tag. For example:

```
<Action name="GetParameters"
    class=" report.framework.DefaultReportActionImpl"
    redirect="CreateReport" />
```

You can also redirect to a different report, subcategory or category using additional attributes as follows:

```
<Action name="DeleteParameters"
    class=" report.framework.DefaultReportActionImpl"
    redirect="CreateReport"
    redirectCategory="main"
    redirectSubCategory="main"
    redirectReport="ReportList.report" />
```

If redirectCategory is not provided, then the current category is assumed and so on.

ShortDesc

Use this tag to provide a very short concise description of the purpose of the report, usually in the form of a question. This is the description that will be displayed for this report in the main list of reports page. If this tag is not provided, then the report title is shown in the list of reports. If the title is not given, then the name itself is displayed.

Description

Use this tag to explain the report's purpose in more detail along with a brief explanation of how the report data is calculated. This helps the user's perception of the report and also helps them to understand the generated report better. This data is displayed when a report is selected from the list of reports page. No description is displayed if none is given.

redirectURL

If set, the report framework will forward the request to this URL upon completion of the report action.

6.1.4.3. Adding Actuate Reports

In addition to writing reports with the TransactionVision reporting framework, you may also use the Actuate Formula One eReport Designer application. This application facilitates the development of TransactionVision reports and is provided on the TransactionVision installation CD. This application is available for Windows platforms only.

Installing and Using the Actuate Formula One Reort Designer

To install the application, double-click the **FormulaOneERD.exe** icon on the TransactionVision installation CD.

Once you install the application, perform the following steps to invoke the designer and connect to a TransactionVision database:

- 1. Using the Windows Control Panel, create an ODBC data source.
- 2. Invoike the Actuate designer application.
- Choose the Data > JDBC Data Source menu item to connect to the data source using a JDBC connection. For DB2, specify the driver as COM.ibm.db2.jdbc.app.DB2Driver and the database URL as jdbc:db2:TVISION, where TVISION is the name of the TransactionVision database.

For instructions on writing reports with the Actuate designer application, see the Actuate designer documentation.

Adding Actuate Reports to Reports.xml

After authoring your report, add it to the Reports.xml file for inclusion in the Reports page, as in the following example:

```
<Report name="Account Summary.report"

title="Account Summary"

layoutfile="/reports/Actuate.jsp"

templatefile="Account Summary.jod"

dataclass="com.bristol.tvision.ui.report.framework.ActuateReport

Bean">

<Action name="CreateReport"

class="com.bristol.tvision.ui.report.framework.ActuateReportBean" />

<ShortDesc>Account Summary Report</ShortDesc>

<Description>
```

```
A sample Actuate report that displays
transactions per account along with
summary information on response time by
transaction type.
</Description>
```

</Report>

The differences between specifying an Actuate report as opposed to a non-Actuate report are as follows:

- For the layoutfile attribute, all Actuates reports should specify a **layoutfile** value of "/reports/Actuate.jsp", which is the default JSP for Actuate reports, unless customization is required. This JSP automatically parses the report and generates the report parameter form. This page can be copied and customized as needed if different behavior and/or formatting is required.
- Acutate reports require the **templatefile** attribute, which specifies the name/location of the Actuate template file to use in generating the report. The path specified is relative to the TVISION_HOME/config/actuate/jodfiles directory. In the example above, the file "Account Summary.jod" is assumed to reside in TVISION_HOME/config/actuate/jodfiles. You may organize Actuate templates by creating subdirectories within the jodfiles folder. Actuate template fiels that the extension .jod.

- For the **dataclass** attribute, all Actuate reports should use the class "com.bristol.tvision.ui.report.framework.ActuateReportBean". This class handles saving and loading of report parameters.
- For the Action attribute, Actuate reports really only need a single CreateReport action, which should use the class "com.bristol.tvision.ui.report.framework.ActuateReportBean", by default. This class interacts with the ActuateReportServlet to generate the report.

Note: The Actuate report engine requires an X connection to run. Your web server should be started in an environment where DISPLAY is set to a valid X server. By default, it will try to connect to the default display, "0:0". If X is not available or permission does not exist for the user to access X on the server machine, then you must set the DISPLAY environment variable to a valid X server.

Setting CLASSPATH for Formula One

To enable a Formula One report to connect to a DB2 JDBC datasource so it can access your data, you must modify the Formula One CLASSPATH to reference the system CLASSPATH as follows:

- 1. Open the <FormulaOne_Insallation>/bin/env.bat script for editing.
- 2. To the CP= line, add the path to db2java.zip (for example, C:\Program Files\IBM\SQLLIB\java\db2java.zip).
- 3. Close and save env.bat.

6.2. Adding Query Pages

Data indexed into lookup tables using the XDM files can be queried using the TransactionVision query page. The query page consists of two parts: the left-hand side query navigator pane and the right-hand side value input pane. The <TVISION_HOME>/config/ui/PresentationQuery.xml file needs to be modified to add a new entry to the query navigator pane and the value input pane the new entry uses.

```
The following is a sample entry in the PresentationQuery.xml file:
<Group name="Stock">
    <Category name="OrderID" desc="Order ID" type="default
jsp="guerySimpleString.jsp">
            <Path>/Event/Data/Order/ID</Path>
    </Category>
    <Category name-"Account" desc="Account Number" type="default"
jps="querySimpleString.jsp">
            <Path>/Event/Data/Order/Account</Path>
    </Category>
</Group>
<Group name="General">
    <Category name="entrytime" desc="Event Entry Time" type="time"
jsp="queryTime.jsp">
            <Path>/Event/StdHeader/EntryTime</Path>
    </Category>
    <Category name="host" desc="Host" type="object" objectType="1"
jsp="queryObject.jsp">
            <Path>/Event/StdHeader/Host/@objectId</Path>
```

</Category>

</Group>

This sample creates two query groups. These groups create a criteria grouping the query navigator pane list box. Each category in the group adds a criteria entry in the navigator pane list box. Each category corresponds to a possible WHERE clause in the generated SQL query. The description text is used as an entry in the navigator list box. When the criteria entry is clicked on, the JSP specified in the jsp attribute is invokded in the right-hand side panel, which brings up the user interface to input the value to query upon. The type attribute describes the type of the field being queried. the Path element specifies the XPath in the XML document of the field being queried. The name attribute of Category gives a unique name to the category. This is used by the query navigator pane to select which value input pane to use. For example,

ViewServlet?viewSelect=query&queryPage=program, where *program* corresponds to the name attribute value.

In the above sample, the Stock group containing categories OrderID and Account are created. The OrderID is a simple string type with and XPath of /Event/Data/Oder/ID.

The type attribute of Category identifies the data type of the WHERE clause. It can only be one of the following values:

time	This is a raw 20-character strin of format <i>yyyymmddbhmmssnnnnn</i> . The query page will use queryTime.jsp to display the edit page. Supported operations include <, <=, ==, >, >=, and query for a time frame.
simpleInt	This is an integer field, such as DataLength. The query page will use querySimpleInt.jsp to display the edit page.
multipleInt	This is a set of integer fields. However, only an enueration of integers is valid for this field, and each value has the corresponding user-friendly description. A typical example is fields of WebSphere MQ completion code. The GUI engine uses queryMultipleInt.jsp or queryMultipleIntExt.jsp to display the edit page, the latter one displayes the value along with the description, while the former one does not. The == operation is supported.
simpleString	This is a formatted string in hexadecimal character code format. If a user wants to query on 'ABC', but the real data in the dabase is stored as '65 66 67', which is the character code of 'ABC' in code page 1252. The query page uses queryByteArray.jsp to facilitate the conversion between 'ABC' to '65 66 67' (CP 1252). It also allows the user to query the hexadecimal string directly.
default	This is a plain text string field. The query page uses querySimpleString.jsp to display the edit page. The only operation supported is 'LIKE'.
object	This query field is an integer of object id. The value of @objectType must be a valid object type id in the system

model tables. The query edit page will retrieve all the object
ids currently in the system model tables for the user to
choose. The value of @jsp determines which JSP is used in
the query page. The value 'queryOjbect.jsp' simply displays a
pane that lets the user multi-select all possible values.
'queryGroupedOject.jsp' displays a multi-select list box of all
possible values and a list of object types to help group-
selecting objects of that type.

groupedObject This is only used for WebSphere MQ objects. This type differs from 'object' in that the query page uses queryGroupedObject.jsp to display all MQ objects, which are further grouped under they MQ queue manager names.

</Category>

In this example, the object id for the XPath to BrokerName is searched in XML documents linearly. Note that this kind of query may be significantly slow on large databases and the query should typically be narrowed down to smaller results using other query conditions like time.

6.3. User Interface Utility Classes

The following utility classes may be useful while developing custom reports and user interface enhancements for TransactionVision.

6.3.1. Class TVisionServlet

public abstract class TVisionServlet extends javax.servlet.http.HttpServlet

6.3.1.1. Methods

getUserBean

getProjectBeanFromSession

getSchemaNameFromSession
getQueryBeanFromSession

getFilterBeanFromSession

getQueryDocFromSession

getSessionBeanFromSession

getCommLinkTmplMgrBeanFromSession

```
public static SessionCommLinkTmplMgrBean
getCommLinkTmplMgrBeanFromSession(javax.servlet.http.HttpSession ses
sion)
```

throws UIException

$get {\it ProjCommLink} Mgr Bean From Session$

```
public static SessionProjCommLinkMgrBean
getProjCommLinkMgrBeanFromSession(javax.servlet.http.HttpSession ses
sion)
throws UIException
```

6.3.2. Class TypeConvService

public class com.bristol.tvision.util.typeconv.TypeConvService extends java.lang.Object

This class contains convenient utility methods to format strings for user interface presentation. The QueryService calls into this object's convert method to perform conversions based on the user interface settings. This class provides date, time and enumeration formatting capabilities.

6.3.2.1. Methods

convert

```
public java.lang.String convert(TypeConvService.Type convType,
```

java.lang.String xpath, java.lang.String value) throws TVisionException

This method converts from a raw value string into user-friendly named description

Parameters:

convType - the conversion type

- TypeConvService.Type.DATE: converts a 20-character time string "*yyyymmddhhmmssnnnnn*" to a user-friendly string. The exact output can be configured through the properties: timeFormat, timeZoneID.
- TypeConvService.Type.DATEONLY: converts an 8-character date yyyymmdd string to user-friendly string.
- TypeConvService.Type.TIMEONLY: converts an 8-character time string hhmmss to user-friendly string.
- TypeConvService.Type.ENUM: give XPath and the raw value in XMLEvent, return the user-friendly description. For example 0 for XPath field "..CompletionCode.." would return MQCC_OK.
- TypeConvService.Type.MSUNIT: appends milli-seconds to the value
- TypeConvService.Type.SECUNIT: append seconds to the value.
- TypeConvService.Type.TIMESKEW: converts from time-skew 20-character string to a user-displayable string.

xpath - The XPath to value's original data field

value - contains string format of the data

Returns:

The return value is a description. This may be null if no XPath is found in the technology constant files at <TVISION_HOME>/java/config/technologyconst.

Throws:

TVisionException – An error occurred during conversion.

retrieve

This method converts from a user-friendly named description to raw value string. This is used for transforming enumeration values to their enumeration values. For example, for a given XPath field, the value 0 can be looked up from the descriptive field say MQCC_OK.

Parameters:

convType - the conversion type

xpath - The XPath to value's original data field

desc - the user-friendly named description

Returns:

value, may be null, if no XPath is found in knowledge base

Throws:

TVisionException - An error occurred during conversion.

getTimeFormat

public int getTimeFormat()

setTimeFormat

public void setTimeFormat(int format)

This method allows setting the time format on all time fields in the result document returned by the QueryService. The valid values for this field are:

```
TVisionCommon.TIME_MILLISECOND_ONLY = 1;
TVisionCommon.TIME_MICROSECOND_ONLY = 2;
TVisionCommon.TIME_MILLISECOND_AND_DATE = 3;
TVisionCommon.TIME_MICROSECOND_AND_DATE = 4;
```

setTimeZoneID

public void setTimeZoneID(java.lang.String id)

This method sets the time zone to which all time fields in the result document returned by the QueryService are converted to. The id string should be a valid Java time zone ID.

getTimeZoneID

public java.lang.String getTimeZoneID()

convertDateStringToDateObj

public static java.util.Date
convertDateStringToDateObj(java.lang.String str)

This method converts a date-time string in the TransactionVision format to Java standard Date object. A date-time string looks like "yyyymmddhhmmssmmm" e.g. "20020123105501123". The last three numbers are millisecond. The micro-second part is ignored.

Parameters:

str - Date String

Returns:

Data object. If failed to convert due to incorrect input string return null.

6.4. Using Job Beans

A job is a taks that runs at a specified frequency. A job typically gathers statistics of recently arrived events and stores calculated results in a way that is easily accessible by a

report. By using a job, the reports themselves do not have to perform comlex, timeconsuming queries to present report data. Instead, they use already calculated data that is periodically updated by a job running in the background. A job is a bean that implements a particular task.

6.4.1. JobBean

```
package com.bristol.tvision.services.analysis.job;
public class JobBean extends implements IJob;
```

All job beans should extend the JobBean basic class and implement the IJob interface. This is required in order for the job to be managed by the job manager.

The JobBean class implements a method called

allowMultipleJobsPerSchema(), which returns true by default. If a job should only ever have one instance allowable per schema, override this method to return false:

Public Boolean allowMultipleJobsPerSchema()

6.4.2. IJob Interface

The IJob interface contains the following functions:

public void init(String startupparam) throws JobException;

The Init method is called once as the job transitions from a stopped to started state. Any one-time initilization for this job can be performed here. The JobManager will pass in this job's startup parameters (specified in the job definition).

public void exit() throws JobException;

The exit method is called when a job is stopped. Any cleanup can be done here. public void run(ConnectionInfo con) throws JobException;

This method will be called when this job is scheduled to perform its particular task. Note: Do **not** call con.close() to disconnect a connection, as it is used internally by TransactionVision.

Job beans do not need to implement database connect logic if the TransactionVision DataManager classes are used for database access. If you use TransactionVision DataManager classes, the Job scheduler will handle the database reconnection. The only requirement for this is to embed a DataManagerException in the JobException thrown from the run() method, as in the following example:

```
Public void run(ConnectionInfo con) throws JobException {
    try{
        [...]
        } catch (DataManagerException ex) {
        throw new JobExceptin(ex);
    }
}
```

public void cancel() throws JobException;

This function is called in order to interrupt a job that is currently processing (that is, while it is in its run method). If you want it to support the ability of the job to be

cancelled while it is running, code your run and cancel methods so that the job gracefully breaks out of whatever processing its in the middle of.

```
public void forceStop() throws JobException;
```

This function is similar to cancel, but indicates that the user wants the job to end immediately without wating for finishing up any cleanup. By default, it calls cancel. It is up to the bean to override and handle this correctly if it is to support a forced stop.

In addition the JobBean class provides the following helper functions public int getJobID()

This will return the current jobs id, which is a unique identifier for a job. public String getState(ConnectionInfo con) throws JobException public void setState(ConnectionInfo con,String state) throws JobException

Up to 128 bytes of customer user data can be stored and associated with a particular job. A Job Bean can optionally take advantage of this feature to store any state information it might want to maintain from one run to another. The getState method will return the current contents of this data. The setState method allows setting of this data.

6.4.3. Creating Jobs at Project Creation

When you create a new project, the project wizard will present you with a page that has a list of standard job templates. This page provides an easy way to include commonly used Jobs into new project when a project is created, without needed to manually create each job individually. The jobs available in this list are controlled by a configuration file (config/ui/Jobs.xml). This file has the following format:

Any definition included in this xml file is listed in the project creation wizard when you create a new project. The priority attribute is currently there for future support, and has no effect at the moment. Supported values for the 'startup' attribute are manual, automatic, and disabled. Supported values for the units attributes are minutes, hours, days, months.

7. Database Schema

7.1. System Object Model Tables

The System Object Model tables are used to store all the System Model objects and the relationships between them. System model objects include general resources as well as technology-specific resources.

Object Types

As such, different technologies will be assigned different ranges of object types. This is described in the table below.

Object Types

Value (range)	Description
0 – 999	Basic System Model Objects (hosts, technologies, Program Instances, etc.)
1	Host
2	Not used
3	Program
4	Program Instance
5	OS/390 Jobname
6	OS/390 Jobstep
7	OS/390 CICS Region
8	OS/390 CICS Transaction
9	OS/390 IMS ID
10	OS/390 IMS Region Type
11	OS/390 IMS Region ID
12	OS/390 IMS Transaction
13	OS/390 IMS PSB
14	OS400 Jobname
15	OS/390 CICS Task
16	User Name
17	Proxy

18	Statistics	
1000-2000	MQSeries Objects	
1000	Unknown type	
1001	None	
1002	Queue	
1003	Local Queue	
1004	Model Queue	
1005	Alias Queue	
1006	Remote Queue	
1007	Cluster Queue	
1008	Local Cluster Queue	
1009	Alias Cluster Queue	
1010	Remote Cluster Queue	
1011	Namelist	
1012	Process	
1013	Queue Manager	
1014	Distribution List	
1015	Cluster	
1016	MQSI Message Flow	
1017	MQSI Broker	
1018	Connection Name	
1019	Cluster Name	
1020	ReplyTo Queue	
1021	1 ReplyTo Queue Manager	
2000	Proxy Object	
3000-3100	Servlet Objects	
3000	Server	
3001	Web Application	
3002	Servlet	
3003	Internet	
3004	JSP	
3005	EJB	
3006	EJB Method	
3101-3199	JMS Objects	

Topic
Queue
CICS Objects
SYSID
APPLID
TREMID
File
TD Queue
TS Queue
TD Alias Queue

Signatures

Each System Model Object has a unique object id that is assigned when the object is inserted into the table. In addition to this unique identifier, each object can be considered to have a signature that identifies that object uniquely. The signature of the object can be generated from event data and looked up in the SYS_MDL_OBJECT table to find the corresponding unique object id. The signature can be uniquely generated from the attributes of the object in an event.

The general format for a signature is a list of all the successor objects from left (highest) to right (the final object), separated by forward slashes. In addition, the object type identifier (see table above) is a prefix to the signature since two objects of different types might otherwise have the same signature.

Object Type	Example Signature	
Host	1/macbeth	
	(Object type/hostname)	
Program Instance (Unix/NT)	4/U/2001080617592300000/132/1	
	(Object type/platform id/start time/process id/thread id)	
Program Instance	2/C/CICS/ABCD/A0F1	
(CICS – OS/390)	(Object type/platform id/CICS region/transaction id/task id)	
MQSeries Queue Manager	1001/qm1 (Object type/queue manager name)	
MQSeries Queue (local)	1002/qm1/LOCAL.QUEUE	
	(Object type/queue manager/queue)	
MQSeries Queue (alias)	1003/qm1/ALIAS.QUEUE	
	(Object type/queue manager/queue)	

Signature Examples

Logical Model



Figure 7-1: System Object Tables Logical Model





Figure 7-2: System Object Tables Physical Model

7.2. Event Tables

Data in the event tables is split up into three basic sections:

- The core event data
- The user data
- Lookup tables

The core event data contains a unique compound key identifying that event and an XML document, which contains the entire event data (minus user data which was not unmarshalled into XML.) The XML data gets stored in LOB columns. For performance reasons, the Analyzer can be configured to store the XML data into a VARCHAR column instead. Should the event XML data exceed the maximum size of this VARCHAR column, a separate row will be inserted into the EVENT_OVERFLOW table, which defines the event_data as LOB. To configure the Analyzer to use VARCHAR, edit the DatabaseDef.xml file in \$TVISION_HOME/config.datamgr and replace:

```
<Table name="EVENT" volatile="true">
<Column name="event data" type="VARCHAR" size="3960"/>
```

Please note that this change will only improve performance if most of the events will fit into the VARCHAR column (thus minimizing the need to use the overflow table). The

maximum size for the VARCHAR is dependent on the database tablespace page size and should be determined by a DBA.

The PARTIAL_EVENT table is a temporary container for Entry- or Exit only events. If the corresponding partial event arrives in the the Analyzer within a defined time interval, a matching thread running in the Analyzer will merge those events and store them in the EVENT table as usual.

User data that was not unmarshalled into XML is stored in the USER_DATA table in the raw format (no data conversion). As with the XML event data, the Analyzer can be configured to use VARCHAR instead of BLOB columns (edit the DatabaseDef.xml file in \$TVISION_HOME/config.datamgr and replace:

```
<Table name="USER_DATA" volatile="true">
<Column name="user_data" type="BLOB" size="10M"/>
```

with the following:

<Table name="USER_DATA" volatile="true"> <Column name="user_data" type="VARBINARY" size="3960"/>

The size of the user_data column may also be changed via the size attribute. However, note that if this value is changed or if the column type is changed, the USER_DATA table must be dropped and then re-created for the changes to take effect.

The lookup tables are used to store fields for quick searching; all columns in these tables are indexed. The XML to Database Mapping (XDM) file uses XPath statements to identify which data items are to be extracted from the XML event data and placed into the lookup tables. Lookup tables for the basic event data and the technology/platform specific MQSeries, OS390, OS400, JMS, Servlet, EJB, and BTTRACE event data are shown in the following figures.

Logical Model



Figure 7-3: Event Tables Logical Model



Figure 7-4: Event Tables Physical Model

7.3. Event Relationship Tables

EVENT_RELATION table stores the relationship between two events determined by technology specific event correlation logic. If the relationship type is defined as BIDIRECTION, there will be two entries in this table: event1 -> event 2 and event2 -> event1. If the logic determines the two events are correlated in certain way with 100% certainty, the confidence factor is set to STRONG_RELATION, otherwise WEAK_RELATION.

RELATION_LOOKUP table stores a correlation lookup id for each event. The logic to generate this lookup id is specific to the technology used by this event.

Logical Model



Figure 7-5: Event Relationship Tables Logical Model

Physical Model



Figure 7-6: Event Relationship Tables Physical Model

7.4. Transaction Tables

Local and Business Transactions are created and updated during the Event Analysis phase in the Analyzer. The local transaction analysis bean populates the LOCAL_TRANSACTION table and links the event data to the corresponding transaction through the column local_trans_id in the table EVENT_LOOKUP. The BUSINESS_TRANSACTION and LOCAL_TO_BUSINESS_TRANS tables are populated during business transaction analysis.

The TRANSACTION_CLASS table contains attributes of all transaction classes that will be made known to the Analyzer, it is static and has to get pre-populated by the user. The BUSINESS_TRANSACTION and TRANSACTION_CLASS tables are defined through an XDM file.

Logical Model



Figure 7-7: Transaction Tables Logical Model



Figure 7-8: Transaction Tables Physical Model

7.5. Statistics Tables

The statistics tables contain data used by various Reports in the TransactionVision web application. The data in the TOPOLOGY_STATS is collected by the Analyzer and used for the static Topology View and as a Datasource for event based reports. The data in the table TRANSACTION_STATS is generated by the TransactionStatisticsJobBean running in the web application and is used for transaction based reports.

Logical Model



Figure 7-9: Statistics Tables Logical Model

Physical model

TOPOLOGY_STATS start_time: TIMESTAMP end_time: TIMESTAMP source_objid: INTEGER dest_objid: INTEGER	TRANSACTION_STATS begin: TIMESTAMP end: TIMESTAMP class_id: INTEGER
link_objid: INTEGER tech_id: INTEGER msg_success: INTEGER msg_error: INTEGER putget_success: INTEGER putget_error: INTEGER byte_error: INTEGER byte_warn: INTEGER byte_error: INTEGER min_latency: INTEGER max_latency: INTEGER avg_latency: DOUBLE type: INTEGER	count: INTEGER success_count: INTEGER failure_count: INTEGER exceeding_sla_count: INTEGER min_response_time: BIGINT max_response_time: BIGINT avg_response_time: BIGINT max_response_txnid: INTEGER

Figure 7-10: Statistics Tables Physical Model

7.6. User Preference Tables

User preference tables stores the personal setting for each user. The setting includes, view options for the event list view, topology view, and time-zone information. When a user logs onto the TransactionVision, the application server will firstly check if there is a database record of the same user id, if not it will read the default setting from <TVISION_HOME>/config/usermgr/DefaultUserData.xml, and create a record for the user. The database record will be updated when the setting is changed by the user in any view. The Query table is used to store the queries for a ceratin project. The query document is saved as XML into the query_doc column. The Storage table is only used for internal purposes.

Logical model



Figure 7-11: User Preference Tables Logical Model

Physical model



Figure 7-12: User Preference Tables Physical Model

7.7. Administration (System) Tables

Logical model



Figure 7-13: Administration Tables Logical Model

Physical model



Figure 7-14: Administratin Tables Physical Model

All of the following tables are created in the "TVISION" schema, and not the userdefined schema to store events.

Analyzer Table: contains all Analyzers defined in the system. Only one Analyzer instance for a given schema is allowed and only one Analyzer instance is allowed to run on a given host. One Analyzer may process data from multiple schemas since the Analyzer is multi-threaded.



Figure 7-15: Analyzer Table

Analyzer_Schema Table: saves a map of Analyzers and the schemas they process data from. One Analyzer can process data for one or more schemas, though multiple Analyzers cannot process data for the same schema.



Figure 7-16: Analyzer_Schema Table

Schema Table: saves the list of database schemas available.

SCHEMA schema_id: INTEGER schema: CHAR(18) apply_timeskew: INTEGER

Figure 7-17: Schema Table

Schema_Version Table: This tables is used to perform a version check between the DataManager and the project schemas in the database.

```
SCHEMA_VERSION
schema: CHAR(18)
version: INTEGER
```

Figure 7-18: Schema_Version Table

Project Table: saves the list of project names and schemas their data is stored into.

PROJECT

```
project_id: INTEGER
project_name: VARCHAR(80)
schema_id: INTEGER
is_active: INTEGER
description: VARCHAR(128)
```

Figure 7-19: Project Table

Project_CommLink Table: One-to-many relation between a project and the communication links it contains. The communication links here are a copy from the global table of communication links. When they are copied over, the commlink name has the project name prepended to it. If the same communication link in global

communication link table is copied into two different projects, the two copies have different commlink_id.

PROJECT_COMMLINK commlink_id: INTEGER project_id: INTEGER commlink_name: VARCHAR(164) commlink_doc: VARCHAR(3072) commlink_secs: INTEGER commlink_msecs: INTEGER

Figure 7-20: Project_CommLink Table

Filter Table: This is a one to many relationship between a project and the data collection filters it contains. The same filter_name in different projects has different filter_id.

FILTER



Figure 7-21: Filter Table

Analyzer_Project_CommLink Table: This is to allow assigning particular communication link (in a particular project) to a particular Analyzer.

ANALYZER_PROJ_COMMLINK analyzer_id: INTEGER

commlink_id: INTEGER

Figure 7-22: Analyzer_Project_CommLink Table

CommLink_Filter Table: This is to allow assigning a particular data collection filter in to a particular communication link (in a particular project) on a particular Analyzer.

COMMLINK_FILTER commlink_id: INTEGER filter_id: INTEGER analyzer_id: INTEGER

Figure 7-23: CommLink_Filter Table

CommLink Table: This is the global communication link table. All created global communication link templates are saved here, and copied into the project table when loaded into a project by user.

COMMLINK

commlink_id: INTEGER

commlink_name: VARCHAR(80)
commlink_doc: VARCHAR(3072)

Figure 7-24: CommLink Table

8. Event XML Schema

This section describes the various XML documents stored in TransactionVision database tables. XML schemas are used to describe TransactionVision data.

8.1. Basic Types

Basic types are technology specific data types and are described using schema tags *xsd:simpleType* or *xsd:complexType*. For example, MQMD belonging to the MQSeries technology may be described in a schema as:

```
<xsd:complexType name="MQMD">
    <xsd:sequence>
        <xsd:element name="Strucld" type="MQCHAR4"/>
        <xsd:element name="Version" type="MQLONG"/>
        <xsd:element name="Report" type="MQLONG"/>
        <xsd:element name="MsgType" type="MQLONG"/>
        <rsd:element name="MsgType" type="MQLONG"/>
        </xsd:element name="MsgType" type="MQLONG"/>
        </xsd:sequence>
        </xsd:sequence>
        </xsd:complexType>
and the basic types MQCHAR4 and MQLONG are:
```

```
<xsd:simpleType name="MQCHAR4">
<xsd:restriction base="xsd:string">
<xsd:length value="4" fixed="true"/>
</xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="MQLONG">
<xsd:restriction base="xsd:long"/>
</xsd:simpleType>
```

Similarly, all datatypes in a particular technology need to be described as above.

Technology specific methods such as MQGET, MQPUT etc. extend the "API" base type.

```
<xsd:element name="MQPUT">

<xsd:complexType>

<xsd:sequence>

<xsd:element name="Hconn" type="MQHCONN"/>

<xsd:element name="Hobj" type="MQHOBJ"/>

<xsd:element name="pMsgDesc" type="PMQMD"/>

<xsd:element name="BufferLength" type="MQLONG"/>

<xsd:element name="pCompCode" type="pMQLONG"/>

<xsd:element name="pReasonCode" type="pMQLONG"/>

</xsd:sequence>

</xsd:complexType>

</xsd:element>
```

8.2. Event Schema Description

An event packet saved in the database would have the following layout: Detailed Schema definition can be found under <TVISION_HOME>/config/xmlschema/Event.xsd.

```
<?xml version="1.0" encoding="UTF-8"?>
<Event>
 <EventID programInstID="642" sequenceNum="7"/>
 <StdHeader minorVersion="1" uow="..." version="5">
   <HostArch>
     <OS>AIX</OS>
     <Vendor>IBM</Vendor>
     <HostArchValue>0xFFFFFF80030780</HostArchValue>
   </HostArch>
   <Encoding>273</Encoding>
 </StdHeader>
 <Technology>
   <MQSeries API="MQPUT" ... >
     <MQPUT>
       <MQPUTEntry>
        <HConn>0x5</HConn>
        <HObj>0x200EC268</HObj>
        <MQMD parameterName="MsgDesc" pointerValue="0x2FF22288">
          <StrucId>MQMD_STRUC_ID &quot;MD&quot;</StrucId>
          <Version>MQMD_VERSION_1 1</Version>
          <Report>MQRO_NONE 0</Report>
          <MsqType>MQMT_DATAGRAM 8</MsqType>
        </MOMD>
         <MQPMO parameterName="PutMsgOpts" pointerValue="0x2FF223F8">
          <StrucId>MQPMO_STRUC_ID &quot;PMO&quot;</StrucId>
          <Version>MQPMO_VERSION_1 1</Version>
          <Options>MQPMO_NONE 0x0</Options>
        </MOPMO>
         <BufferLength>25</BufferLength>
        <Buffer pointerValue="0x2FF2253C">
          <UserDataRef chunk="0"/>
         </Buffer>
        <CompCode pointerValue="0x2FF224FC">N/A</CompCode>
        <ReasonCode pointerValue="0x2FF22500">N/A</ReasonCode>
       </MQPUTEntry>
       <MQPUTExit>
        <HConn>0x5</HConn>
        <HObj>0x200EC268</HObj>
       </MQPUTExit>
     </MQPUT>
   </MQSeries>
 </Technology>
 <Data>
   <Chunk blobType="0" ccsid="0" from="0" seqNo="0" to="24"/>
 </Data>
</Event>
```

The diagram below shows the basic structure of the type hierarchy of objects used to describe an event.



9. The Data Manager

9.1. Using the DataManager to Access the Database

Custom beans and reports that need to access the database may use the service interface of the DataManager class to conveniently perform tasks which otherwise would have to be coded on the JDBC level.

A reference to the DataManager object can be obtained with the instance() method.

If the DataManager instance is used outside of the TransactionVision application context (for example, in a standalone Java application), the first call into the DataManager **must** be

DataManager.instance().init()

Beans and reports that run within the TransactionVision application are not required to do this; they can expect the instance to be successfully initialized.

Custom beans running within the TranactionVision Analyzer Framework will usually get the current database connection passed in as a parameter of class ConnectionInfo, which encapsulates the JDBC connection handle and the database schema name for the current processed event:

```
public class ConnectionInfo {
```

}

```
/** The database connection */
public Connection con;
/** The database schema */
public String schema;
```

```
public ConnectionInfo(Connection con, String schema) ;
```

In cases where the custom code needs to obtain its own database connection, the DataManager offers three different methods for this purpose:

getThreadConnection() will return a connection for the current thread. If this is the first time the thread calls into this method, a new connection to the database is returned. Every following call from the same thread will return the same connection, until it is getting released with releaseThreadConnection().

getSessionConnection(String sessionId) will return a new connection the first time this method is called for a certain session Id, and then return the cached connection for all further calls until the connection is relased with releaseSessionConnection(String sessionId) getConnection() will always create and return a new connection to the database. This connection will get released with a call to releaseConnection(Connection con).

Here is the complete list of the methods that make up the supported DataManager interface:

public static DataManager **instance**() Returns the DataManager Singleton instance

Returns:

The DataManager instance

public void init(java.lang.String propertyFile)

throws com.bristol.tvision.datamgr.DataManagerException Initializes the DataManager according to the settings in the specified properties file. NOTE : This method has to be called before any other method.

Parameters:

dbProperties - The Database.properties file containing the db settings

Throws:

com.bristol.tvision.datamgr.DataManagerException - If initialization fails

public void init()

throws com.bristol.tvision.datamgr.DataManagerException Initializes the DataManager with the default properties file (Database.properties)

Throws:

com.bristol.tvision.datamgr.DataManagerException - If initialization fails

```
public java.sql.Connection getThreadConnection()
```

throws

com.bristol.tvision.datamgr.DataManagerException

Returns the database connection for the current thread. If there is no connection stored in the connection map for this thread, a new connection is established by calling into the configured ConnectionSource, and this connection will be returned for all following calls.

Returns:

The database connection for the current thread

Throws:

com.bristol.tvision.datamgr.DataManagerException - if getting a new
connection from the ConnectionSource fails

```
public void releaseThreadConnection()
```

throws

com.bristol.tvision.datamgr.DataManagerException

Releases (closes) the connection for the current thread.

Throws:

com.bristol.tvision.datamgr.DataManagerException - if closing the connection fails

com.bristol.tvision.datamgr.DataManagerException

Returns the database connection for the specified web session. If there is no connection stored in the connection map for this session, a new connection is established by calling into the configured ConnectionSource, and this connection will be returned for all following calls.

Parameters:

sessionId - The session id

Returns:

The database connection for the session

Throws:

com.bristol.tvision.datamgr.DataManagerException - if getting a new connection from the ConnectionSource fails

```
public void releaseSessionConnection(java.lang.String sessionId)
```

throws

com.bristol.tvision.datamgr.DataManagerException

Releases (closes) the connection for the specified session.

Throws:

com.bristol.tvision.datamgr.DataManagerException - if closing the connection fails

public java.sql.Connection getConnection()

throws

com.bristol.tvision.datamgr.DataManagerException

Returns a new database connection which is not cached, which means every call into this method will obtain a new connection from the configured ConnectionSource.

Returns:

The database connection

Throws:

com.bristol.tvision.datamgr.DataManagerException - if getting a new connection from the ConnectionSource fails

Close the connection which has been obtained from a call to getConnection.

Throws:

 $\verb|com.bristol.tvision.datamgr.DataManagerException-if closing the connection fails||$

public void commitTransaction(java.sql.Connection con)

throws

 $\verb|com.bristol.tvision.datamgr.DataManagerException||$

Performs a commit on current the database transaction

Parameters:

con - The connection holding the transaction to commit

Throws:

com.bristol.tvision.datamgr.DataManagerException - if the commit fails

public void rollbackTransaction(java.sql.Connection con)

throws

com.bristol.tvision.datamgr.DataManagerException

Performs a rollback on the current database transaction

Parameters:

con - The connection holding the transaction to roll back

Throws:

com.bristol.tvision.datamgr.DataManagerException - if the rollback fails

9.2. XML-Database Mapping Using XDM Files

The TransactionVision database schema is made extensible through the XML to Database Mapping (XDM) files. XDM is a generic way to describe the mapping of values contained in XML documents onto table columns in the database and allows fast, indexed XML data retrieval by the database engine.

The XML mapping is implemented by the class XMLDatabaseMapper and is used in TransactionVision to store the event and transaction data into lookup tables for fast retrieval. This class is also accessible from custom beans and reports and allows user written code to map basically any XML data to the database.

XML mappings are grouped into different 'document types'. Each document type is defined by the root tag value for its documents and describes a mapping from XML to a set of database tables that logically belong together. These tables must share the same primary key, and the join across all these tables represents the mapped XML data for one XML document. In TransactionVision there are three predefined document types:

/Event

This document type consists of all event based XML mappings, including standard header event data, technology specific event data, and platform specific event data.

/Transaction

This document type maps data for the transaction analysis to the database tables.

/TransactionClass

This document type contains a single mapping for business class attributes.

9.3. The XDM Syntax

XML mappings are defined in XDM files in the <TVISION_HOME>/config/xdm directory. The XML schema format of XDM files is defined in

<TVISION_HOME>/config/xmlschema/XDM.xsd. Each XDM file defines a mapping of XML data to a particular database table. The syntax to describe this mapping is as follows:

<Mapping documentType="/Event">

Defines the document type for this mapping. This mapping is only valid for XML documents that have the same root tag as "documentType".

<DBSchema>Stock</DBSchema>

Defines that this mapping is only valid for the specified database schema. The insert() method of the XMLDatabaseMapper will not write any columns of this mapping if the supplied database schema parameter does not match. There are multiple <DBSchema> tags allowed. If this tag is missing, the mapping is valid for all schemas.

```
<Key name="proginst_id" type="INTEGER"
description="ProgramInstanceId">
<Path>/Event/EventID/@programInstID</Path>
</Key>
<Key name="sequence_no" type="INTEGER" description="SequenceNumber">
<Path>/Event/EventID/@sequenceNum</Path>
</Key>
```

Defines the primary key for the database table. All XDM mappings of the same document type must have the same key definition. There may be multiple key tags, in which case a compound primary key will get created. The structure of the key tag is similar to the Column tag and will be described there.

<Table name="EVENT_LOOKUP" category="COMMON">

Specifies the database table for the mapping. For mappings of the document type "/Event", the XDM mappings can be technology or platform specific. The category attribute on the Table tag contains either "COMMON" or the technology string or the platform string for the event data that should be written into this table. The string "COMMON" indicates that this table contains data common to every event and should be written for every event going through the Analyzer. A technology or platform name like "MQSERIES" or "OS390_BATCH" used in the category field indicates that this table should only be filled for events of that technology or platform. Examples:

```
<Table name="EVENT_LOOKUP" category="COMMON">
...
</Table>
```

```
//Table>
</Table name="OS390_LOOKUP"
category="OS390_BATCH,OS390_CICS,OS390_IMS">
...
</Table>
```

For other document types, the value of the category attribute should always contain the string "COMMON".

```
<Column name="host_id" type="INTEGER" description="Host" isObject="true">
```

```
<Path>/Event/StdHeader/Host/@objectId</Path>
```

</Column>

Each table mapping consists of several Column definitions that describe which XML value has to be mapped onto which database table column. The **name** attribute specifies the column name, and the **type** attribute specifies the column type, which can be one of the following:

- INTEGER
- FLOAT
- DOUBLE
- CHAR
- VARCHAR
- DATE
- TIMESTAMP

Both **name** and **type** are required. Types CHAR and VARCHAR require an additional attribute **size**.

The **description** attribute specifies the name of the tag containing the value for that column in the query result document returned by the QueryServices. Required.

The **isObject** attribute for a Column tag in the above XDM file refers to that column being an identifier for an object in the system model table. This allows to use the object name instead of the numerical, system generated object id in XDM based queries. Possible values: 'yes/no'. Default value if missing: 'No'.

The **queryOnly** attribute for a Column tag indicates that the value is not written by the XMLDatabaseMapper when invoking its write() method, but from some other code independent of the XDM mapping. Hence, the column value is only considered for queries.

The **generated** attribute for a Column tag means that column is a database generated id. Possible values: 'yes/no'. Default value if missing: 'No'.

The **conversionType** attribute for a Column tag means that field requires a formatting conversion after reading from the database. The TypeConvService is called into after reading that field from the database. This is typically used for writing enumeration fields
(conversionType='enum'). Refer to the TypeConversionService for more information on how values are converted.

The **indexed** attribute specifies if a database index should be created for this column for faster query access. Possible values: 'yes/no'. Default value if missing: 'Yes'.

<Path> contains the XPath of the document value to write into the table column. The XMLDatabaseMapper will extract the value form the XML document and insert it into the database. Note that only XPaths pointing to Text nodes and attribute values are valid. If a value specified by the XPath does not exist in the XML document, a NULL value is inserted to the database.

A column can map to multiple XPath expressions as in the sample code below. The XPath expressions are evaluated in a sequential order and the first value found will get inserted into the database.

```
<Column name="datasize" type="INTEGER" description="DataSize">
<Path>/Event/Technology/MQSeries/MQGET/MQGETExit/DataLength</Path>
<Path>/Event/Technology/MQSeries/MQPUT/MQPUTExit/BufferLength</Path>
<Path>/Event/Technology/MQSeries/MQPUT1/MQPUT1Exit/BufferLength</Pat
h>
```

```
</Column>
```

In addition to the <Path> element, a column definition can contain a <Join> definition like in the following example:

```
<Column name="class_id" type="INTEGER" description="ClassId">
<Path>/Transaction/ClassId</Path>
<Join documentType="/TransactionClass"</Join>
</Column>
```

Join definitions offer a way to link two different document types together in order to use column definitions of both document types in one query. Internally this will generate a database join between the column of the current table and the primary key of the other table.

9.3.1. Creating the XDM database tables

One important aspect of the XDM framework is that the creation of the underlying database tables is entirely data-driven. The definitions in the XDM files are not only being used for updating or querying the XML data, but also as an input to the TransactionVision Table Manager, which is responsible for creating and dropping the project tables as projects in the Analyzer GUI get created and deleted. Thus there is no need to issue any SQL DDL calls to the database. Once the XDM file is placed into the proper directory, and provided the document type is registered with the Table Manager, the new tables defined in the XDM mapping get automatically created for a new project. The same holds true if the project tables get created or dropped by using the command line tool CreateSqlScript.

The registration with the Table Manager is only needed if the XDM mapping uses a new user defined document type. The only thing to do is to add the new document type to the following section of the DatabaseDefinition.xml in the <TVISION_HOME>/config/datamgr directory:

<XDM>

```
<DocumentType>/Event</DocumentType>
<DocumentType>/Transaction</DocumentType>
```

```
<DocumentType>/TransactionClass</DocumentType>
<DocumentType>/MyNewDocType</DocumentType>
</XDM>
```

9.3.2. Properties of the TransactionVision Document Types

9.3.2.1. The /Event Document Type

Event-based XDM files specify that when an XML event is written to the database by the DBWrite module in the Analyzer, these fields are extracted and written into the database columns defined by the XDM mappings. Similarly, when the database is queried to retrieve event based data in the Analyzer GUI, these XDM files are used to construct the corresponding SQL query. The XML document for each event gets stored in the database table EVENT.

9.3.2.2. The /Transaction Document Type

This mapping is used to write business transaction attributes during the transaction analysis phase in the Analyzer. One noticeable difference to the event-based mappings is that there is no XML document inserted into the database, all document values are always mapped to the database tables.

9.3.2.3. The /TransactionClass Document Type

This document type contains only one XDM mapping file for the "transaction_class" table and describes the attributes of a transaction class. This document type is 'artificial' in the way that the underlying database table contents is static and must not be updated during the runtime of the application. Thus the XDM mapping solely serves a query purpose and allows, together with the document type join feature, to use transaction class attributes in business transaction queries.

9.4. The XMLDatabaseMapper Interface

The XMLDatabaseMapper can be used in 2 different ways: implicitly when writing custom bean code in the Analyzer bean framework or using the query facilities of the QueryServices, or explicitly by obtaining a reference to an XMLDatabaseMapper instance and calling into one of the available service methods.

To obtain a reference to an instance, the instance() method has to be called with the particular document type as an argument, e.g.:

XMLDatabaseMapper xdm = XMLDatabaseMapper.instance(myDocType);

The interface contains methods for reading, inserting, updating, and deleting XML values. All methods take a parameter of class XMLDocument, which denotes the XML document containing the data. The XMLDocument class implements the org.wc3.dom.Document Interface and can be constructed in several ways: from an existing document using the constructor XMLDocument (org.w3.dom.Document doc), or entirely bypassing the generation of any XML objects and creating a 'lightweight' XMLDocument instance by using the constructor XMLDocument (java.util.Map).

The class contains an internal HashMap for caching XPath expressions to the corresponding values in the XML document. The key of the map entry is an XPath

expression, the value of the map entry is the value in the XML document corresponding to that XPath. If an instance is created by using the latter constructor, then any value lookup on the document translates into a simple HashMap lookup, whereas a lookup on an instance created with the first constructor is performed by executing an XPath search on the XML document (unless the corresponding XPath is already in the cache). This is implemented transparently for the caller by the following method of XMLDocument:

public String getDocumentValue(String xpath) throws XMLException;

If there is a value for the given XPath in the HashMap, the stored value is returned. Otherwise an XPath search on the document is performed.

With these 'lightweight' XML documents it is possible to provide data to the XMLDatabaseMapper without having to make expensive XML operations. The XMLTransaction class used in the transaction analysis is one example of such a 'lightweight' XML object.

Following is the list of available XMLDatabaseMapper methods: public void **read**(com.bristol.tvision.datamgr.ConnectionInfo conInfo,

Reads all lookup table rows for the given key values and store the values in the attribute map of the XML document. The document passed in only needs to contain the key values.

Parameters:

con - The database connection to use

doc - The document containing the key values

Throws:

com.bristol.tvision.datamgr.DataManagerException - Error while accessing the document or reading from the database tables

```
public void
write(com.bristol.tvision.datamgr.ConnectionInfo conInfo,
```

com.bristol.tvision.services.analysis.XMLDocument doc)
 throws com.bristol.tvision.datamgr.DataManagerException

Writes the values of the mapped document elements to the lookup tables. For each mapped column defined in the xdm files, the value of the corresponding XPath expression is searched in the xml document and written to the table column defined in the mapping.

Parameters:

con - The database connection to use

doc - The document to search

Throws:

com.bristol.tvision.datamgr.DataManagerException - Error while accessing the document or writing to the database tables

```
public void
update(com.bristol.tvision.datamgr.ConnectionInfo conInfo,
```

com.bristol.tvision.services.analysis.XMLDocument doc)

throws com.bristol.tvision.datamgr.DataManagerException

Updates the values of the mapped document elements in the lookup tables. All columns that are defined by the document type will get updated. The rows to update are determined by the key values in the XML document.

Parameters:

con - The database connection to use

doc - The document containing the updated values

Throws:

com.bristol.tvision.datamgr.DataManagerException - Error while accessing the document or writing to the database tables

```
public void
```

delete(com.bristol.tvision.datamgr.ConnectionInfo conInfo,

Deletes rows in all lookup tables of the document type for the given key values in the XML document.

The document passed in only needs to contain the key values.

Parameters:

con - The database connection to use

doc - The document containing the key values

Throws:

com.bristol.tvision.datamgr.DataManagerException - Error while accessing the document or writing to the database tables

9.5. Extending the /Event Document Type

The XDM mappings of the /Event document type can be easily extended to map additional XML data to indexed database columns for faster retrieval. First, this can be done for XML values that are already present in the standard XML event data but which are not included in the default event based XDM mapping definitions. In this case the mapping for the desired values can be simply added (with its XPath and database column) to the corresponding XDM file (event.xdm. mqseries.xdm, etc.).

Second and more important, additional mappings can be defined for XML data that has been assembled from the contents of the user data buffer by an EventModifierBean (see chapter 3.2). Although this user defined XML data could also be mapped to the existing lookup tables (by simply modifying one of the existing XDM files), this is not advisable. For this purpose a new XDM file defining a mapping to a new table should be created. The mapping definition is required to have the document type /Event and the key columns proginst_id and sequence_no like all other event based XDM files. The column definitions should include all XDM values intended for display in the Analyzer GUI or queries through the query services. For steps to configure the Analyzer GUI to display these new columns see Chapter 3.

The TransactionVision DeleteEvents utility and job use an optimized fast deletion scheme based on timestamp columns if the -older option is used. To delete data in user-defined XDM tables, theis timestamp column must be present in any additional XDM mapping you define. Therefore, the following section is mandatory in the XDM file:

9.6. Extending the /Transaction and /TransactionClass Document Type

The /Transaction and /TransactionClass document types can be extended to add custom business transaction and transaction class attributes to the transactional data in TransactionVision. See chapter 3.5.4 for details.

The TransactionVision DeleteEvents utility and job use an optimized fast deletion scheme based on timestamp columns if the -older option is used. To delete data in user-defined XDM tables, theis timestamp column must be present in any additional XDM mapping you define. Therefore, the following section is mandatory in the transaction document type:

Adding New Document Types

It is possible to create new document types that are independent of the TransactionVision event and analysis process and entirely controlled by custom code. This allow user written code to store and retrieve XML data in a convenient way without having to code SQL on the JDBC level. The custom code can insert and modify data for this document type by using the various interface methods of the XMLDatabasMapper instance directly, or by using the corresponding interface methods in QueryServices which allow query based updates. Data query and retrieval can be accomplished by creating a QueryDoc containing conditions based on the document type, and using the QueryServices class to retrieve the data. The necessary steps are:

- 1. Create a new XDM file with the documentType set to the root tag of the XML data to handle. The key definition should be set to the primary key of the new table.
- 2. Define a XPath-column mapping for each XML value that has to be stored into the table.
- 3. If the new table should be automatically created for new projects, register the new document type with the TableManager (see chapter 9.3.1).
- 4. Get a XMLDatabaseMapper instance for the document type with XMLDatabaseMappper.instance(newDocType) and use its method to insert/modify the XML data, or use the corresponding QueryServices interface.
- 5. Use the QueryServices interface to query and retrieve the XML data (see chapter 4).