

HP Business Availability Center

for the Windows and Solaris operating systems

Software Version: 8.01

Using TransactionVision

Document Release Date: March 2009

Software Release Date: March 2009



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Third-Party Web Sites

HP provides links to external third-party Web sites to help you find supplemental information. Site content and availability may change without notice. HP makes no representations or warranties whatsoever as to site content or availability.

Copyright Notices

© Copyright 2000 - 2009 Hewlett-Packard Development Company, L.P.

Trademark Notices

TransactionVision® is a registered trademark of the Hewlett-Packard Company.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Oracle® is a registered trademark of Oracle Corporation and/or its affiliates.

UNIX® is a registered trademark of The OpenGroup.

Documentation Updates

This guide's title page contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign-in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

You can visit the HP Software Support web site at:

<http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software Support Online provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the HP Software Support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Table of Contents

Welcome to This Guide11
How This Guide Is Organized11
Who Should Read This Guide12
Getting More Information12

PART I: TRANSACTIONVISION

Chapter 1: Introducing TransactionVision.....15
TransactionVision Overview16
TransactionVision Basic Concepts16
TransactionVision Data.....20
TransactionVision User Permissions22
Tips for Using Browsers23
Using the TransactionVision Trade Demo Sample Events24
TransactionVision Workflow26
Creating the Trade Project28
Setting Up Trade Demo Sample Events.....31
Importing the Trade Demo Sample Events.....33
Business Transactions in Dashboard User Interface34

PART II: ADMINISTRATION

Chapter 2: Introducing TransactionVision Administration39
TransactionVision Administration Overview39
Current Projects Page40
Administration Page.....40
Chapter 3: Projects.....43
Projects Overview43
Projects User Interface.....44
Troubleshooting and Limitations64

Chapter 4: Database Schemas	65
Database Schemas Overview	65
Database Schemas User Interface	66
Chapter 5: Analyzers	71
Analyzers Overview	71
Creating or Editing an Analyzer Host	72
Configuring the Analyzer.....	72
Stopping and Restarting the Analyzer Process.....	74
Failure Mode.....	74
Analyzers User Interface	83
Troubleshooting and Limitations	93
Chapter 6: Communication Links.....	105
Communication Links Overview	106
Understanding Communication Links	106
TransactionVision Queue Requirements	110
Communication Link Templates	117
Creating New Communication Links or Templates	120
Create a New Communication Link or Template for WebSphere MQ	121
Create a New Communication Link or Template for TIBCO EMS	130
Create a New Communication Link or Template for Progress SonicMQ.....	135
Create a New Communication Link or Template for BEA WebLogic JMS.....	140
Create a New Communication Link or Template for Other JMS Provider.....	142
Test a Communication Link	144
Communication Links User Interface	145
Troubleshooting and Limitations	212
Chapter 7: Data Collection Filters.....	215
Data Collection Filters Overview	215
Planning Data Collection Filters	216
Filter Conditions	219
Create a New Data Collection Filter.....	258
Edit a Data Collection Filter.....	260
Data Collection Filters User Interface	261

Chapter 8: Jobs	267
Jobs Overview	267
Create a New Job – Workflow	272
Make a Job Accessible in the UI/Job Server	274
Jobs User Interface	275
Chapter 9: Transaction Definition Editor.....	283
Transaction Definition Editor Overview	284
Transaction Definition Editor User Interface.....	288
Troubleshooting and Limitations	309
Chapter 10: Alias Management	311
Alias Management Overview.....	311
Alias Management User Interface	312
Chapter 11: Queries	317
Queries Overview	317
Optimizing Query Performance	318
Add a New Query	319
Edit a Query	321
Query Conditions.....	322
Queries User Interface	357
Chapter 12: TransactionVision HP Business Settings	361
HP Business Availability Settings	362
TransactionVision Data in Business Availability Center	362
Drill downs from End User Management to TransactionVision	366
HP Business Process Insight Settings.....	369
Classify Events for Business Process Insight Integration	372
HP Business Settings User Interface	376
Troubleshooting and Limitations	381
Chapter 13: Additional Reference Information	385
Administration Utilities	385
APIs Monitored by Sensors and Agents.....	413
Technical References for Sensor and Agent Technologies	421
Troubleshooting	421

PART III: REPORTS AND TOPOLOGIES

Chapter 14: Introducing TransactionVision Reports and Topologies	441
TransactionVision Reports and Topologies Overview	441
User-created Reports	442

Chapter 15: Aggregated Topology	445
Aggregated Topology Overview	445
Response Time Calculation and Breakdown.....	447
Aggregated Topology User Interface	448
Chapter 16: Instance Topology	453
Instance Topology Overview.....	453
Calculating Response Time	455
Instance Topology User Interface	455
Troubleshooting and Limitations	463
Chapter 17: Transaction Topology Flow Map	465
Transaction Topology Flow Map Overview	465
Viewing the Transaction Topology Flow Map	466
Time Frame for the Transaction Topology Flow Map	473
Transaction Topology Flow Map User Interface	475
Chapter 18: Component Topology Analysis	491
Component Topology Analysis Overview	492
Viewing the Component Topology Analysis	492
Edges.....	494
Set Background Color	501
Modify Layouts and Properties	502
Printing the Component Topology Analysis Graphs	508
Adjusting the Zoom Level	512
Configure Edge and Node Labels	514
Modify Component Groupings	519
Viewing Statistics.....	524
Component Topology User Interface	525
Troubleshooting and Limitations	532
Chapter 19: Transaction Detail Reports	533
Transaction Detail Reports Overview.....	534
Event Analysis Overview	535
Run a Report.....	539
Enable Session Tracking	541
Set Event Analysis Display Options	542
Event Attributes.....	545
Transaction Detail Reports User Interface.....	550
Chapter 20: Transaction Infrastructure Reports	569
Transaction Infrastructure Reports Overview	570
Transaction Infrastructure Reports User Interface	571

Chapter 21: Transaction Summary Reports.....	587
Transaction Summary Reports Overview	588
Save Report Settings	590
Transaction Summary Reports User Interface	591
Index.....	615

Welcome to This Guide

This guide describes how to use HP Business Availability Center TransactionVision to track transactions and graphically show the interaction between all the components of your system.

This chapter includes:

- How This Guide Is Organized on page 11
- Who Should Read This Guide on page 12
- Getting More Information on page 12

How This Guide Is Organized

This guide contains the following parts:

Part I TransactionVision

Provides an overview of TransactionVision concepts, data collection, and a general workflow of using TransactionVision. It also describes how to use and set up the Trade Demo Sample Events, which provide simulated data for demonstration purposes to help you become familiar with using TransactionVision.

Part II Administration

Describes how to set up and configure TransactionVision to track transactions.

Part III Reports and Topologies

Describes how to view and customize reports and topologies of business transactions.

Who Should Read This Guide

This guide is intended for the following users of HP Business Availability Center:

- HP Business Availability Center administrators
- HP Business Availability Center application administrators
- HP Business Availability Center end users

Readers of this guide should be knowledgeable about navigating and using enterprise applications, and be familiar with HP Business Availability Center and enterprise monitoring and management concepts.

Getting More Information

For a complete list of all online documentation included with HP Business Availability Center, additional online resources, information on acquiring documentation updates, and typographical conventions used in this guide, see the *HP Business Availability Center Deployment Guide* PDF.

Part I

TransactionVision

1

Introducing TransactionVision

This chapter includes the main concepts for using TransactionVision.

Concepts

- TransactionVision Overview on page 16
- TransactionVision Basic Concepts on page 16
- TransactionVision Data on page 20
- TransactionVision User Permissions on page 22
- Tips for Using Browsers on page 23
- Using the TransactionVision Trade Demo Sample Events on page 24

Tasks

- TransactionVision Workflow on page 26
- Creating the Trade Project on page 28
- Setting Up Trade Demo Sample Events on page 31
- Importing the Trade Demo Sample Events on page 33

Reference

- Business Transactions in Dashboard User Interface on page 34

TransactionVision Overview

TransactionVision is the transaction tracking solution that graphically shows you the interaction between all the components of your system. TransactionVision non-intrusively records individual electronic events generated by a transaction flowing through a computer network. More importantly, TransactionVision's patented “Transaction Constructor” algorithm assembles those events into a single coherent business transaction.

Graphical analysis of business transactions enable you to:

- Find lost transactions
- Monitor and meet service level agreements
- Improve efficiencies of your business processes

TransactionVision Basic Concepts

To understand the tasks required to administer TransactionVision, you need to understand the TransactionVision components, terms and concepts, and some data concepts, described in the following sections:

- “TransactionVision Components” on page 16
- “TransactionVision Terms and Concepts” on page 17

TransactionVision Components

TransactionVision consists of the following major components:

- UI/Job Server
- Analyzer
- Sensors

- RDBMS
- Data Manager

See *the HP TransactionVision Deployment Guide* PDF for a complete description of each component.



TransactionVision Terms and Concepts

To administer TransactionVision effectively, you must be familiar with the following terms and concepts:

- “Communication Links” on page 17
- “Event Collection” on page 17
- “Projects” on page 19
- “Data Collection Filters” on page 19
- “Database Schemas” on page 20

Communication Links

Communication links enable a TransactionVision Sensor to communicate from a host on which an application is being monitored to the Analyzer. Two communication paths must be defined for each communication link: one path for configuration messages from the Analyzer to the Sensors and another path for captured events from the Sensors to the Analyzer. When you define a communication link, you specify the name of the queue manager and queues the Sensor monitors for configuration messages and sends event messages to, as well as the queue managers and queues the Analyzer sends configuration messages to and retrieves event messages from. See “Communication Links” on page 105 for more information.

Event Collection

The WebSphere MQ Sensor library implements all API entry points and has the same name as the standard library for the monitored technology (for example, mqm.dll for WebSphere MQ on Windows). It is installed in a different directory location, and your library search path is altered so that programs load the Sensor library at runtime instead of the standard technology library.

When a program running on the system where a Sensor is installed calls a WebSphere MQ API for the monitored technology, it actually calls the corresponding function in the Sensor library.

The Sensor function first generates a TransactionVision event, recording the API call and other details based on the data collection filters. It then invokes the actual API from the standard technology library. When the actual API returns, the Sensor function generates another TransactionVision event, this time representing the exit state of the function, and then forwards the return information to the calling program.

Note: For WebSphere MQ Sensors to intercept API calls from applications, the applications must link dynamically to the technology library as a shared library. Otherwise, Sensors cannot record events for the applications.

On the z/OS CICS platform, WebSphere MQ Sensors use the API-crossing exit mechanism provided by the CICS adapter of WebSphere MQ for z/OS. On z/OS batch and IMS platforms, applications are statically linked to, or dynamically invoke, TransactionVision stubs, which are replacements for the standard WebSphere MQ stubs.

A WebSphere API Exit Sensor is also available for other platforms.

The WebSphere Application Server servlet, JMS, and EJB Sensors use Java bytecode instrumentation to intercept servlet, JMS, and EJB API calls. At installation of the WebSphere servlet Sensor, the server JVM settings are modified to add a Sensor classloader plugin. This plugin intercepts servlet API calls and allows the TransactionVision Sensor to report them. The JMS API called are trapped using static bytecode instrumentation and the instrumented jar file must be added to the application CLASSPATH before any other JMS jar file. See “Additional Reference Information” on page 385 for more information.

Projects

Event collection projects enable you to easily group and manipulate communication links, data collection filters, database schemas, and Analyzers as one entity. An event collection project is used by an Analyzer to define the communication links, the data collection filters, and the database schema that data is written into.

Each TransactionVision project is assigned a single host running the Analyzer. Projects enable you to easily group and manipulate communication links, data collection filters, database schemas, and Analyzers as one entity. When you start a project, the Analyzer on the host assigned to the project is started automatically. You may also start the Analyzer on a host from the Analyzers page in the web user interface, in which case the Analyzer starts processing events on all active projects it is assigned to. See “Projects” on page 43 for more information.

Data Collection Filters

Data collection filters assigned to a project determine the amount and type of information collected by each Sensor. Data collection filters specify criteria such as the following:

- Which technologies, hosts, programs, or APIs to collect information about
- Which CICS regions, transactions, and job names to collect information about
- Which queues or queue managers to collect information about
- Which servlets, WebSphere applications, WebSphere servers and URIs to collect information about
- What time range to collect information for
- The level of detail to be collected, such as API name only, API name and call arguments, or API name, call arguments, and data buffer segment. For the Servlet Sensor, the default data buffer size is 1K. See “Data Collection Filters” on page 215 for more information.

Database Schemas

The project schema defines the tables into which the events collected by Sensors are stored. When the Analyzer retrieves and processes events collected by Sensors, it places them into event related tables. By using schemas to partition event data by project, you can control access to event data collected by each project. See “Database Schemas” on page 65 for more information.

TransactionVision Data

This section includes:

- “Business Transaction State Categories” on page 21
- “Business Transaction Properties” on page 21

TransactionVision tracks user-defined business transactions throughout their entire lifecycle. TransactionVision Sensors trace WebSphere MQ API, CICS API, J2EE servlet, EJB, and JMS calls across transaction-based systems, and track and analyze the interaction with each component.

The data for each transaction is sent to HP Business Availability Center in aggregated samples (see “BAC Aggregate TV Sample Data” on page 364 for details).

This data is mapped to the Business Transaction CIs (see “Map End-User Transactions to TransactionVision Business Transactions” in *Using End User Management*).

The samples from TransactionVision contain metrics for the completed and in-process (backlogged) business transaction instances, as well as volume, value and latency. See “TransactionVision Data in Business Availability Center” on page 362 for details on completed and in-process business transactions. For information on the volume, value and latency metrics, see “Transaction Topology Flow Map Metrics” on page 471.

These metrics are further broken down by various state categories for the data, according to definitions in the TransactionVision application.

Business Transaction State Categories

The state categories are:

- **Failed.** Business transaction instances that did not meet a success threshold defined by the TransactionVision Failures attribute rule. For example, a stock buy transaction would fail if the account balance was less than \$10,000, which is required to do the purchase, or a margin buy would fail if the account ratio is below a minimum balance threshold.
- **Successful.** The total count of transaction instances, minus the failed count.
- **Late.** (Also referred to as **delayed**.) Business transaction instances with a response time that exceeded the threshold set in TransactionVision. The threshold used can be seen in the properties for the TV Monitor CI.
- **Exceptions.** Business transaction instances that did not follow the expected flow path, as defined by the TransactionVision Exception attribute rule.

Note: Transactions can be classified in more than one state category, for example, as both late and as successful.

Business Transaction Properties

The properties for the TV Monitor CI include attributes that are populated from the transaction class in TransactionVision. The properties received from TransactionVision are used by the KPIs and reports for the TV Monitor CI. You can view the properties by selecting **Properties** in the context menu for a TV Monitor CI in **Applications > Dashboard > Console**. The business transaction dialog box displays CI properties in the General Properties panel and TV Monitor CI properties in the Other Properties panel.

For descriptions of the CI properties in the General Properties panel, see “Configuration Item Properties Dialog Box” in *Model Management*.

For descriptions of the TV Monitor CI properties, see “Other Properties Dialog Box” on page 34.

TransactionVision User Permissions

User permissions determine what options are available and what operations users can perform in the **Admin > TransactionVision > Administration** page or **Admin > TransactionVision > Current Projects** page. User or group permissions can be set by assigning a pre-defined role or by granting individual permissions for specific options (see “Current Projects Page” on page 40 for details).

By default, TransactionVision permission changes go into affect the next time the user logs in. For more information about permissions and setting them, see “User Management” in *Platform Administration*.

The following table shows the pre-defined roles, what operations can be performed by each role, and which Administration page can be accessed:

Role	Operations	Admin Pages Allowed
TransactionVision Administrator	FULLCONTROL on Projects, Analyzer and User-created Reports	Current Projects and Administration
TransactionVision Operator	<ul style="list-style-type: none"> ➤ VIEW on Projects and User-created Reports ➤ EXECUTE on Analyzer 	Current Projects
TransactionVision User	VIEW on Projects and User-created Reports	Current Projects
Business Availability Center Administrator	FULLCONTROL on Projects, Analyzer, User-created Reports and User Data	Current Projects and Administration
Business Availability Center System Modifier	<ul style="list-style-type: none"> ➤ VIEW and CHANGE on Projects ➤ VIEW on User-created Reports ➤ EXECUTE on Analyzer 	Current Projects
Business Availability Center System Viewer	VIEW on Projects and User-created Reports	Current Projects

CMDB View Permissions Requirements

- You must grant view permission to the CMDB Business Transactions View resource (located in the CMDB permissions context) when you grant permissions only by assigning any of the three TransactionVision roles (described in the table), and/or grant permissions on individual TransactionVision resources. Otherwise, the user has no access to the Transaction Over Time, Performance Summary, Aggregated Topology or Instance Topology reports.
- The following reports require view permissions to the CMDB Business Transactions View resource.
 - Transaction Over Time
 - Performance Summary
 - Aggregated Topology
 - Instance Topology
- A user that has CMDB view permissions, but no project permissions, can access the reports listed above.
- All other reports can only be accessed if users have project view permissions.

Tips for Using Browsers

The following information provides useful tips when using browsers to view your reports. See the *HP TransactionVision Deployment Guide* PDF for supported browsers.

- Users who have browser pop-up blockers must disable them for the TransactionVision user interface. Otherwise, certain reports may open windows with Session Expired error messages, and help pages are inaccessible. See the documentation for your pop-up blocker to determine how to add TransactionVision to the list of web applications for which pop-ups are enabled.

- Many TransactionVision web client pages contain session information and are temporary. Therefore, creating browser bookmarks to return to these pages at a later time may not work.
- If you set up the TransactionVision host and port setting in Business Availability Center to be **https://**, then you must login to Business Availability Center with **https://**.

Using the TransactionVision Trade Demo Sample Events

To help you become familiar with TransactionVision's features and capabilities, and the types of transaction information you can track, TransactionVision provides a sampling of trade events called *Trade Demo Sample Events* that can be processed by the Analyzer and placed into the TransactionVision database. The Trade Demo Sample contains simulated transaction data spanning over one month. This simulated data is for demonstration purposes only.

Note: Corrections and additional instructions for importing the Trade Demo Sample Events can be found in the <TVISION_HOME>\samples\readme.txt file, which is included with the TransactionVision Analyzer installation package. See this file before importing the sample events.

To get started, see “Creating the Trade Project” on page 28.

Files Installed with the Trade Demo Sample Events

The following Trade Demo Sample Events files are in the `<TVISION_HOME>/samples/trade` directory.

File	Description
tvision_trade_events_800.zip	Zip file jar file containing the Trade Demo Sample Events.
TradeTransaction.xdm	Trade demo XML database mapper transaction definition file defining trade-specific business transaction attributes.
TransactionDefinition.xml	The classification rule file.
TradeTransactionClass.xdm	Trade demo XML database mapper transaction class definition file defining trade-specific class attributes.
TransactionClass.xdm	Trade demo XML database mapper transaction class defining the trade transaction classes.

TransactionVision Workflow

This task describes a suggested working order for using the TransactionVision application to configure your collected events and view them graphically.

For a scenario of this task, see “Creating the Trade Project” on page 28.

Note:

- Some of the steps are optional.
 - The optional steps can generally be performed at any point after you have created the project after step 5, “Create and Configure Communication Links” on page 27.
-

This task includes the following steps:

- “Prerequisites” on page 27
- “Create and Configure a New Project or Select an Existing Project” on page 27
- “Create a Database Schema” on page 27
- “Configure Analyzers” on page 27
- “Create and Configure Communication Links” on page 27
- “Configure Data Collection Filters” on page 27
- “Define and Configure Scheduled Jobs” on page 27
- “Create aAliases for System Objects” on page 27
- “View and Configure Transaction Topologies” on page 27
- “View and configure transaction detail reports” on page 28
- “View and configure transaction infrastructure reports” on page 28
- “View and configure summary reports” on page 28

1 Prerequisites

Before you begin, you need to have completely installed and configured the TransactionVision UI. See the *HP TransactionVision Deployment Guide* PDF for details.

2 Create and Configure a New Project or Select an Existing Project

The Projects Wizard steps you through the workflow as shown in this task when you create a new project. See “Projects” on page 43 for details.

3 Create a Database Schema

See “Database Schemas” on page 65 for details.

4 Configure Analyzers

See “Analyzers” on page 71 for details.

5 Create and Configure Communication Links

See “Communication Links” on page 105 for details.

6 Configure Data Collection Filters

See “Data Collection Filters” on page 215 for details.

7 Define and Configure Scheduled Jobs

See “Jobs” on page 267 for details.

8 Create aAliases for System Objects

See “Alias Management” on page 311 for details.

9 View and Configure Transaction Topologies

You can view the transaction topologies using the default settings, or you can specify settings to view what meets your needs.

- View the Aggregated Topology, which displays the flow path for the aggregation of instances of a business transaction. See “Aggregated Topology” on page 445 for details.

- View the Instance Topology, which displays the flow path for an instance of a business transaction. See “Instance Topology” on page 453 for details.
- View the Component Topology, which displays the interaction between all system components for which sensors collect event information.

10 View and configure transaction detail reports

You can view the transaction detail reports using the default settings, or you can specify settings to view what meets your needs.

These reports display transaction activity for web sessions, transaction tracking, event analysis, and service level analysis. See “Transaction Detail Reports” on page 533 for details.

11 View and configure transaction infrastructure reports

You can view the transaction infrastructure reports using the default settings, or you can specify settings to view what meets your needs.

These reports display information for servers, message queue latency, infrastructure performance, SQL statement statistics and WebSphere MQ and JMS components statistics. See “Transaction Infrastructure Reports” on page 569 for details.

12 View and configure summary reports

You can view the transaction summary reports using the default settings, or you can specify settings to view what meets your needs.

These reports display transaction performance. See “Transaction Summary Reports” on page 587 for details.

Creating the Trade Project

This task describes how to use the Project Wizard to create a project called **Trade** with the schema **TRADE** (see “Project Wizard” on page 53 for details):

This task includes the following steps:

- “Open the Project Page” on page 29

- “Specify Project Name” on page 29
- “Create a New Database Schema” on page 29
- “Select an Analyzer” on page 29
- “Select a Global Communication Link Template” on page 29
- “Assign Data Collection Filters to the Communication Link” on page 30
- “Create Jobs from Template” on page 30
- “Save the New Project” on page 31

1 Open the Project Page

Choose **Admin > TransactionVision > Administration > Projects** and click **New Project** to start the Project Wizard.

2 Specify Project Name

In Step1 of the Project Wizard, enter **Trade** and click **Next**.

3 Create a New Database Schema

In Step2 of the Project Wizard, enter **TRADE** as the new database schema, uncheck **Apply timeskew to collected events**, and click **Next**.

4 Select an Analyzer

In Step3 of the Project Wizard, if an Analyzer has not already been created, click **New Analyzer...**, enter the Analyzer host name and then click **Finish**.

Select the created (or existing) Analyzer then click **Next**.

5 Select a Global Communication Link Template

In Step4 of the Project Wizard, if a Global Communication Link Template does not already exist, create a new one by clicking **Create New Template**.

Note: Because the Trade Demo Sample Events are processed directly by the Analyzer and bypass the Queue Manager (such as WebSphere MQ and Sonic MQ), the Communication Link Template is not actually used. However, Transaction Vision Projects require a communication link and therefore you need to create one if one does not already exist.

The following instructions describe how to create a WebSphere MQ communication link using the Communication Link Creating Wizard (which opens when you select **Create New Template**). If you use a different messaging system, see Chapter 6, “Communication Links,” on how to create a communication link.

- a** In step 1 of the Communication Link Creation Wizard, enter a New Communication Link Name such as **Trade Comm**, select **IBM WebSphere MQ**, then click **Next**.
- b** In step 2 of the Communication Link Creation Wizard, enter the queue manager name **TRADING** and click **Next**.
- c** In steps 3, 4 & 5 of the Communication Link Creation Wizard, click **Next**.
- d** In the final step of the Communication Link Creation Wizard, click **Finish**.
- e** Proceed with step 4 of the Project Wizard by selecting the newly created (or existing) Communication Link Template, click the **Create from Template** button, then click **Next**.

6 Assign Data Collection Filters to the Communication Link

In step 5 of the Project Wizard, **Collect All** data collection filter definition is already created and assigned. Click **Next**.

7 Create Jobs from Template

In step 6 of the Project Wizard, accept the defaults and click **Next**.

8 Save the New Project

In the final step of the Project Wizard, view the project summary to be sure these are the settings you want, and click **Finish**.

You can now set up the Trade Demo Sample Events. See “Setting Up Trade Demo Sample Events” on page 31 for instructions.

Setting Up Trade Demo Sample Events

This task describes how to set up the Trade Demo Sample Events.

This task includes the following steps:

- “Stop the Analyzer” on page 31
- “Copy and Unzip the tvision_trade_events_800.zip File to Your Hard Drive” on page 31
- “Set the read_from_jar Property in the Analyzer.properties File” on page 32
- “Copy TransactionDefinition.xml, TradeTransactionClass.xdm and TransactionClass.xdm files to a Temporary Directory” on page 32
- “Import the Trade Business Transactions and Classification Rules” on page 32
- “Restart the Analyzer” on page 32
- “Restart the UI” on page 33

1 Stop the Analyzer

Enter the following command to stop the Analyzer:

```
<TVISION_HOME>\bin\nanny.bat stopService Analyzer
```

2 Copy and Unzip the tvision_trade_events_800.zip File to Your Hard Drive

- a Copy **tvision_trade_events_800.zip** from the **<TVISION_HOME>/samples/trade** directory to a location on your system's hard drive.

- b** Unzip the **tvision_trade_events_800.zip** archive. On Windows, use either the Windows Compressed Folder utility or WinZip. On UNIX platforms, use the command:

```
unzip tvision_trade_events_800.zip
```

- c** Copy the resulting **trade_raw_events.jar** file into <TVISION_HOME>/logs directory.

3 Set the read_from_jar Property in the Analyzer.properties File

Modify the <TVISION_HOME>/config/services/Analyzer.properties file and uncomment the statement **read_from_jar=true**.

4 Copy TransactionDefinition.xml, TradeTransactionClass.xdm and TransactionClass.xdm files to a Temporary Directory

Copy **TransactionDefinition.xml**, **TradeTransactionClass.xdm** and **TransactionClass.xdm** files from the <TVISION_HOME>/samples/trade directory to a temporary directory of your choosing.

5 Import the Trade Business Transactions and Classification Rules

From within the temporary directory, run:

```
<TVISION_HOME>\bin\DataUtil.bat -import_txndef TransactionDefinition.xml.
```

Be sure to include a space " " and a dot "." after **TransactionDefinition.xml** to import the Trade transaction class attributes.

6 Restart the Analyzer

Enter the following command to restart the Analyzer:

```
<TVISION_HOME>\bin\nanny.bat startService Analyzer
```

7 Restart the UI

Enter the following commands to restart the UI:

```
<TVISION_HOME>\bin\nanny.bat stopService Tomcat  
<TVISION_HOME>\bin\nanny.bat startService Tomcat
```

You can now import the Trade Demo sample events. See “Importing the Trade Demo Sample Events” on page 33 for instructions.

Importing the Trade Demo Sample Events

This task describes how to import the Trade Demo Sample Events.

Note: Before importing the sample events, see the `<TVISION_HOME>\samples\readme.txt` file for corrections and additional instructions for importing the Trade Demo Sample Events. This file is included with the TransactionVision Analyzer installation package.

This task includes the following steps:

- “Start the Project” on page 33
- “Stop the Project and the Analyzer” on page 34
- “Reset the read_from_jar property in the Analyzer.properties File” on page 34

1 Start the Project

From within the Projects page, click the **Start** button to start the Project. The Analyzer now processes the raw events and stores them into the database. Note that the eventcount on the commlink does not change since the events are bypassing the event queue. But you can follow the Analyzer progress by observing the event count in the database (click the **Update** button next to **Events in database schema**).

2 Stop the Project and the Analyzer

After all events have been processed (around 103,000), stop the project and the Analyzer.

3 Reset the read_from_jar property in the Analyzer.properties File

Modify the <TVISION_HOME>/config/services/Analyzer.properties file and comment the statement `read_from_jar=true`.

Business Transactions in Dashboard User Interface

This section describes:

- Other Properties Dialog Box on page 34

Other Properties Dialog Box

Description	Displays the TV Monitor CI attributes, which you can modify. The properties received from TransactionVision are used by the KPIs and reports for the TV Monitor CI. To access: On the Applications > Dashboard > Top View page or Console page, select the Business Transactions view, then right-click on a TV Monitor CI to open the Context menu and select Properties .
Useful Links	<ul style="list-style-type: none">➤ “TransactionVision Data” on page 20➤ “Business Transaction Properties” on page 21➤ Chapter 12, “TransactionVision HP Business Settings”

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Close	Closes the dialog box.
Help	Provides help for each field in this dialog box.
Aggregation delta	The time period allocated for the delay while the sample is created. The default is 360 seconds, defined in TransactionVision on the Admin > TransactionVision > Business Availability Settings page using the Time slice shift in job (in seconds) option. See “HP Business Availability Settings Page” on page 379 for details.
Collection interval	The time period, in seconds, covered by each sample received for this business transaction. The time is set to 300 seconds (5 minutes).
Edit Selector Expression	Enables you to edit the Selector Expression.
Selector Expression	A Boolean expression, which displays which rows of the data in the profile database (TransactionVision samples) are associated with the selected TV Monitor CI using the txn_id .
TransactionVision currency unit	The applicable currency for the monetary values received for this business transaction. This currency unit is displayed in the Dashboard tabs and reports for all monetary values associated with this business transaction.
TransactionVision Exception attribute rule	The check box for this property is selected if the Exception attribute rule is defined for the transaction class in TransactionVision. If the check box is selected, Business Availability Center assigns the Exceptions KPI to the business transaction.
TransactionVision Failures attribute rule	The check box for this property is selected if the Failures attribute rule is defined for the transaction class in TransactionVision. If the check box is selected, Business Availability Center assigns the Failures KPI to the business transaction.

GUI Element	Description
TransactionVision ID	The internal ID for the business transaction in TransactionVision.
TransactionVision threshold	The response time threshold defined for the business transaction in TransactionVision. Instances of the transaction that take longer than this time are considered late .
TransactionVision Value attribute rule	The check box for this property is selected if the Value attribute rule is defined for the transaction class in TransactionVision. If the check box is selected, Business Availability Center assigns the Value KPI to the business transaction.

Part II

Administration

2

Introducing TransactionVision Administration

This chapter introduces TransactionVision and the TransactionVision Administration pages.

This chapter includes:

Concepts

- TransactionVision Administration Overview on page 39
- Current Projects Page on page 40
- Administration Page on page 40

TransactionVision Administration Overview

The TransactionVision Administration pages enable users to set up and configure transactions for data collection and report viewing.

To access the TransactionVision Administration pages, select **Admin > TransactionVision**.

TransactionVision Administration includes the following pages:

- “Current Projects Page” on page 40
- “Administration Page” on page 40

Current Projects Page

The Current Projects page is available for all users. The options on this page limit users to work with the currently selected project. The user's permissions determine what options are available:

- Users with view only permissions on the project can view Select Project, Project, and Queries.
- Users with modify permissions on the project can view and modify all the options on the Current Projects page.

For more information about permissions and setting them, see “Configure Users and Permissions - Workflow” in *Platform Administration*.

The Current Projects page includes:

- Select Project and Projects, see “Projects” on page 43 for details.
- Jobs, see “Jobs” on page 267 for more details.
- Communication Links, see “Communication Links” on page 105 for details.
- Data Collection Filters, see “Data Collection Filters” on page 215 for details.
- Queries, see “Queries” on page 317 for details.

Administration Page

The Administration page allows users with administrative permissions only to work globally across projects. The Administration page is hidden from all other users. For more information about permissions and setting them, see “Configure Users and Permissions - Workflow” in *Platform Administration*.

The Administration page includes:

- Projects, see “Projects” on page 43 for details.
- Jobs, see “Jobs” on page 267 for more details.
- Analyzers, see “Analyzers” on page 71 for details.
- Communication Link Template, see “Communication Links” on page 105 for details.

- Database Schema, see “Database Schemas” on page 65 for details.
- Alias Management, see “Alias Management” on page 311 for details.
- Transaction Definition Editor, see “Transaction Definition Editor” on page 283 for details.
- HP Business Process Insight Settings, see “TransactionVision HP Business Settings” on page 361 for details.
- HP Business Availability Settings, see “TransactionVision HP Business Settings” on page 361 for details.

3

Projects

This chapter includes:

Concepts

- Projects Overview on page 43

Reference

- Projects User Interface on page 44

Troubleshooting and Limitations on page 64

Projects Overview

Event data collected by TransactionVision is organized into projects. When you are logged into TransactionVision, one project is the current project—any project-related changes are applied to the current project, and any Topologies or analysis views use data from the current project. Use commands under **Admin > TransactionVision > Current Projects** to view status or make changes to the current project.

If you have administrator permissions, you can use commands under **Admin > TransactionVision > Administration > Projects**. From here you can create new projects or modify any existing projects.



Projects User Interface

This section describes:

- Edit Projects Page on page 44
- Projects Page on page 48
- Project Status Page on page 51
- Project Wizard on page 53
- Welcome to TransactionVision Page on page 64



Edit Projects Page

Description	Displays details of the selected project, for editing. To access: Select Admin > TransactionVision > Administration > Projects , select a project and click Edit Project .
Important Information	Only users with Administrator permissions may start or stop data collection

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Finish	Click to save edits made to this project.
Cancel	Click to cancel edits made to this project.
Analyzer Assigned to Project	Displays the Analyzer host assigned to this project, required for correlating and performing transaction analysis on events in the schema.

GUI Element	Description
Apply timeskew to collected events	Check this to apply the timeskew detected by TransactionVision on the communication link to events. This box is checked by default. Note that disabling the timeskew calculation affects all projects that use the database schema assigned to your project. You must restart the Analyzer associated with your project for this change to take effect.
Database Schema	Displays the name of the database schema associated with the project.
Description	Description of the project, this field can be edited.
Only save events of failed transactions (requires transaction classification to be enabled in the Analyzer)	<p>Check to enable failure mode, in which the Analyzer only stores event data of failed business transactions (or transactions violating SLA). For successful business transactions, only the corresponding business transaction rows are saved. The main goal of the Analyzer Failure Mode is to reduce the data storage requirement for a project.</p> <p>Failure mode requires the following:</p> <ul style="list-style-type: none"> ➤ A transaction classification with suitable rules for the Result column must be in place. See Chapter 9, “Transaction Definition Editor,” for details on transaction classification. ➤ Failure mode processing must be enabled in the Analyzer. For instructions, see “Failure Mode” on page 74. <p>The failure mode can be enabled on a per-project basis, allowing to run certain projects in failure mode while running others in the conventional Standard Mode.</p>
Project Name	Displays the name of the selected project.

Communication Link Assignment

GUI Element	Description
Create	<p>The Analyzer used by a project must be assigned one or more communication links. This allows a Sensor to communicate from a host on which an application is being monitored to the Analyzer.</p> <p>To create a communication link for the project that is not based on a global template, click Create to start the Communication Link wizard. See “Project Communication Links Page” on page 211 for information about creating communication links.</p>
Create from Template	<p>To assign a communication link to the Analyzer, select a communication link template in the list of global communication link templates and click Create from Template. The communication link created for your project is a copy of the selected template; changes to the template do not affect the communication link used by your project.</p>
Create New Template	<p>Click to add a new global communication link template, and start the Communication Link wizard. See “Global Communication Link Templates Page” on page 209 for more information.</p>
Delete	<p>Select a communication link, then click Delete to delete it.</p> <p>To disable a communication link without deleting it, clear the check box to the left of the communication link. To enable a disabled communication link, check the box.</p>
Edit	<p>Select a communication link, then click Edit to make changes to it.</p>
Global Communication Link Templates	<p>Displays the available Communication Link Templates for selection, see “Communication Link Templates” on page 117 for more information.</p>
Project Communication Links	<p>Displays the communication links for the project. They can each be selected and edited or deleted.</p>

Communication Link - Data Collection Filter Assignment

GUI Element	Description
Collect All	By default this filter is selected, which collects all events.
Communication Link to Assign Filters to	Select the communication link to view/assign filters from the drop-down list, then check filters to assign them, or clear checks to unassign filters.
Create new Filter	Click to create a new data collection filter definition, and display the Add Filter page. See “Create a New Data Collection Filter” on page 258 for information about creating new data collection filters.
Data Collection Filters	You must assign one or more data collection filters to each communication link for the Analyzer assigned to your project.

Alias and Project Summaries

GUI Element	Description
Alias List Assignment	An alias maps a user-defined name to system model objects displayed in the Aggregate, Instance and Component Topologies, and Event Analysis view. Alias definitions are grouped into alias lists, which may be associated with one or more projects.
Alias List to assign to this project:	To associate an alias list with the project, select it from the drop-down list. See Chapter 10, “Alias Management,” for information about creating alias lists and defining alias mappings.
Project Summary	Any changes made to the project on this page display here, giving an opportunity to verify before clicking Finish or Cancel .



Projects Page

Description	<p>This page gives an overview of projects available. When a project is selected, project details display and you can start, stop, edit, delete and create new projects from this page.</p> <p>To access: Select Admin > TransactionVision > Administration > Projects</p>
Useful Links	<ul style="list-style-type: none">➤ “Edit Projects Page” on page 44➤ “Project Status Page” on page 51➤ “Project Wizard” on page 53

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Delete Project	Click to delete the selected project.
Edit Project	Click to edit the selected project, see “Edit Projects Page” on page 44 for more details.
New Project	Click to create a new project, see “Project Wizard” on page 53 for details.
Refresh	Click to refresh the page.
Start	To activate the current project, click Start .
Stop	To stop data collection, click Stop .

GUI Element	Description
Analyzer	<p>The name, IP address, and status of the Analyzer host assigned to the schema associated with the project. Status values are:</p> <ul style="list-style-type: none"> ➤ COLLECTING EVENTS ➤ NOT COLLECTING EVENTS ➤ QUIESCING ➤ FAIL ➤ RECOVERY ➤ SENSOR COLLECTING EVENTS ONLY <p>If the project receives a stop command and its Analyzer is still cleaning out events in the event queue, the Analyzer status is set to QUIESCING and the Start and Stop buttons for the project are disabled.</p> <p>A FAIL status indicates that TransactionVision cannot connect to the Analyzer host.</p> <p>The RECOVERY status indicates that the Analyzer is performing statistics recovery. This can happen when the Analyzer has been stopped abnormally (for example, the process has been killed), and not all event statistics have been flushed to disk. In this case the Analyzer reprocesses the last events to get the statistics up-to-date before it starts the normal processing. If the status is RECOVERY, the number of events processed and the total number of events to process (for example, “230 of 2140 events processed”) is displayed. Click Refresh to update the number of events processed.</p> <p>The SENSOR COLLECTING EVENTS ONLY status indicates that the Analyzer has been stopped or has exited without sending stop configuration messages to the Sensors. Sensors continue to produce events, which collect on the event queue until the Analyzer and project are started again.</p>
Database Schema	Displays the name of the database schema associated with the project.
Events in Database Schema	The number of events processed by the schema associated with the project.

GUI Element	Description
Project communication links	<p>Displays a list of the communication links assigned to the Analyzer host. For each communication link, it shows:</p> <ul style="list-style-type: none"> ➤ The number of configuration messages sent ➤ The number of events processed ➤ The number of unprocessed events on the event queue ➤ The number of event packages sent ➤ The time skew ➤ The data collection filters assigned to the communication link. <p>To view or edit a communication link, click the communication link name to display the Edit Project Communication Link page. See “Edit Communication Link Template” on page 184 for more information.</p> <p>To view or edit an assigned data collection filter, click the filter name to open the Edit Data Collection Filter page. See “Edit Data Collection Filter Page” on page 263 for more information.</p>
Project Name	Displays the name of the selected project.
Status	Indicates whether the project is currently collecting events. If the status is ACTIVE , the project is collecting events according to the settings in the data collection filter. If the status is INACTIVE , the project is not currently collecting events.

Project Status Page

Description	<p>The Project Status page displays information about the current project. Use this page to start and stop data collection or edit project settings.</p> <p>To access: Select Admin > TransactionVision > Administration > Projects</p>
Important Information	Only users with Administrator permissions may start or stop data collection
Useful Links	<ul style="list-style-type: none"> ➤ “Edit Projects Page” on page 44 ➤ “Project Wizard” on page 53

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Clear Logs	Click to clear Analyzer log messages from the display.
Edit Project	<p>Click to display the Edit Project page and change any of the following:</p> <ul style="list-style-type: none"> ➤ Description ➤ Whether to apply timeskew to collected events ➤ Communication link assignments ➤ Data collection filter assignments
Refresh	<p>Click to update the view to reflect any changes since it was created.</p> <p>When you refresh the page, the event count is not updated.</p>
Start	To activate the current project, click Start .
Stop	To stop data collection, click Stop .
Update	Click Update (to the right of the event count) to update the number of events in the project.

GUI Element	Description
Analyzer	<p>A FAIL status indicates TransactionVision cannot connect to the Analyzer host.</p> <p>The RECOVERY status indicates that the Analyzer is performing statistics recovery. This can happen when the Analyzer has stopped abnormally (for example, the process has been killed), and not all event statistics have been flushed to disk. In this case the Analyzer reprocesses the last events to get the statistics up-to-date before it starts the normal processing. If the status is RECOVERY, the number of events processed and the total number of events to process (for example, “230 of 2140 events processed”) is displayed. Click Refresh to update the number of events processed.</p> <p>The SENSOR COLLECTING EVENTS ONLY status indicates the Analyzer has stopped or exited without sending stop configuration messages to the Sensors. Sensors continue to produce events, which collect on the event queue until the Analyzer and project are started again.</p>
Communication Links	<p>This page displays a list of the communication links assigned to the Analyzer host. For each communication link, it shows:</p> <ul style="list-style-type: none"> ➤ The number of configuration messages sent ➤ The number of events processed ➤ The number of unprocessed events on the event queue ➤ The number of event packages sent ➤ The time skew ➤ The data collection filters assigned to the communication link.
Database Schema	<p>The name of the database schema associated with the project.</p>

GUI Element	Description
Events in Database Schema	The number of events in the project database schema. TransactionVision calculates the event count the first time the Project Status page is displayed in a session. When the page is refreshed, the event count is not updated for performance reasons. To update the event count, click Update . Updating the event count does not refresh the other information on the Project Status page. Move your cursor over the Update button to view a tooltip that shows the date and time the event count was calculated.
Project	The name of the current project. To change the current project, click Home and select a new project on the TransactionVision Home page.
Status	Indicates whether the project is currently collecting events. If the status is ACTIVE , the project is collecting events according to the settings in the data collection filter. If the status is INACTIVE , the project is not currently collecting events.

Project Wizard

Description	Enables you to create a new project. To access: Select Admin > TransactionVision > Administration > Projects
Wizard Map	The Project wizard contains: Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page

Important Information	New projects may only be created by users with administrator permissions.
Useful Links	<ul style="list-style-type: none">➤ “Edit Projects Page” on page 44➤ “Project Status Page” on page 51➤ “Projects Page” on page 48

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Project Description	Enter a brief description for your project.
Project Name	Enter a name for your project.

 **Database Schema Setup Page**

Description	Each project must be assigned a database schema where the Analyzer stores events collected by Sensors. To access: Select Admin > TransactionVision > Administration > Projects , click New Project .
Important Information	General information about the Project wizard is available in “Project Wizard” on page 53.
Wizard Map	The Projects wizard contains: Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Alias List to assign to this schema	To assign an alias list to the schema, select it from the list of available alias lists. If you assign an alias list, alias names are used for objects in the list in TransactionVision views and reports. For more information on using aliases, see Chapter 10, “Alias Management.”
Apply timeskew to collected events	Check this to apply the timeskew detected by TransactionVision on the communication link to events. This box is checked by default. Note that disabling the timeskew calculation affects all projects that use the database schema assigned to your project.
Create a new Database Schema	To create a new schema, enter the schema name. Note that schema names must begin with a letter and may contain only letters, numbers, and underscore characters.
Database Tablespace	Select the database tablespace to use for the project from the list of available tablespaces. The default tablespace is the tablespace picked by the database to create the schema. Note: You do not see this option if you use Microsoft SQL Server as your database for TransactionVision.

GUI Element	Description
Only save events of failed transactions (requires transaction classification to be enabled in the Analyzer)	<p>Check this to enable failure mode, in which the Analyzer only stores event data of failed business transactions (or transactions violating SLA). For successful business transactions, only the corresponding business transaction rows are saved. Failure mode requires the following:</p> <ul style="list-style-type: none"> ➤ A transaction classification with suitable rules for the Result column must be in place. See Chapter 9, “Transaction Definition Editor,” for transaction classification information. ➤ Failure mode processing must be enabled in the Analyzer. For instructions, see “Configuring the Analyzer” on page 72. <p>The failure mode can be enabled on a per-project basis, allowing to run certain projects in failure mode while running others in the conventional Standard Mode.</p> <p>The main goal of the Analyzer Failure Mode is to reduce the data storage requirement for a project.</p>
Use Existing Database Schema	<p>You may create a new schema for the project or use an existing one. To use an existing schema, click the down arrow and select a schema from the list.</p> <p>Note that users of projects using the same schema see all events from all projects that the schema is assigned to. Creating separate schemas for different projects enables you to control access to event data.</p>

Assign Analyzer to Project Database Schema Page

Description	<p>This step involves assigning an Analyzer host to the project, for correlating and performing transaction analysis on events in the schema.</p> <p>To access: Select Admin > TransactionVision > Administration > Projects, click New Project.</p>
Important Information	<p>A single Analyzer host may serve multiple schemas, but each schema may be served by only one Analyzer host. Therefore, if you assigned an existing schema to your project, the Analyzer host used by the schema is assigned to your project automatically.</p> <p>General information about the Project wizard is available in “Project Wizard” on page 53.</p>
Wizard Map	<p>The Projects wizard contains:</p> <p>Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page</p>

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
New Analyzer	Click this button to define a new Analyzer host. See “Create Analyzer Page” on page 88 for details.
Analyzers	Select an existing Analyzer host from the list.

 **Assign Communication Links to Analyzer Page**

Description	<p>This step is only necessary if you selected an Analyzer in the previous step.</p> <p>A project's Analyzer must be assigned one or more communication links. These communication links are required for a Sensor to communicate from a host on which an application is being monitored, to the Analyzer.</p> <p>To access: Select Admin > TransactionVision > Administration > Projects, click New Project.</p>
Important Information	<p>General information about the Project wizard is available in "Project Wizard" on page 53.</p>
Wizard Map	<p>The Projects wizard contains:</p> <p>Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page</p>

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Delete	To delete a communication link assigned to the Analyzer, select it in the list of communication links and click Delete .
Create from Template ->	To assign a communication link to the Analyzer, select a communication link template in the list of global communication link templates and click Create from Template-> . The communication link created for your project is a copy of the selected template; changes to the template do not affect the communication link used by your project.

GUI Element	Description
Create New Template	To add a new global communication link template, click Create New Template to start the Communication Link wizard. See “Creating New Communication Links or Templates” on page 120 for information on creating communication link templates.
Global Communication Links Template	Displays the existing global communication link templates available to the project. See “Communication Link Templates” on page 117 for more information.

Assign Data Collection Filters to Communication Link Page

Description	<p>This step only occurs if you assigned a communication link in the previous step.</p> <p>You must assign one or more data collection filters to the communication links for the Analyzer assigned to your project. By default the Collect All filter, which collects all events, is selected.</p> <p>To access: Select Admin > TransactionVision > Administration > Projects, click New Project.</p>
Important Information	<p>You may assign any number of filters. Events that match any filter are collected by Sensors and forwarded to the Analyzer. If you create a new filter with specific filter conditions, make sure you remove the Collect All filter. Otherwise, all events continue to be collected because they match the conditions of a filter. To remove a data collection filter from the list, select it and click Remove.</p> <p>General information about the Project wizard is available in “Project Wizard” on page 53.</p>
Wizard Map	<p>The Projects wizard contains:</p> <p>Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page</p>

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Create New Filter	To create a new data collection filter definition, click Create New Filter to display the Add Filter page. See “Create a New Data Collection Filter” on page 258 for information on creating a filter.
Data Collection Filter Definitions	To assign a filter to the communication links for the Analyzer associated with your project, select the filter name in the Data Collection Filter Definitions list and click Assign Filter >. The filter name appears in the Data Collection Filters list. For more information on Filters, see “Data Collection Filters” on page 215.

Create Jobs from Template Page

Description	<p>A job is a task that runs at a specified frequency. You may assign one or more jobs from available job templates to your project.</p> <p>To access: Select Admin > TransactionVision > Administration > Projects, click New Project.</p>
Important Information	<ul style="list-style-type: none"> ➤ A job instance is defined for this particular project; that instance can only be run on this project. However, multiple instances of the job bean may be run concurrently, each configured for a different project. For more information about jobs, see Chapter 8, “Jobs.” ➤ If the project you are creating shares a schema with another project, note that it also shares jobs with that project. For example, if you create a delete events job in project A and one in project B, if both projects share the same schema, the effects of both delete event jobs is seen in both projects. When creating jobs for projects that share a schema, keep this in mind to prevent one project's jobs from having unintended consequences on data that is shared between the projects.
Wizard Map	<p>The Projects wizard contains:</p> <p>Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page</p>

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<List of Jobs>	Check the check box for any job you wish to create an instance of for your new project.
<Making changes to Jobs>	<p>To change the Startup Mode, Job Interval, or startup settings for jobs after your project is created, choose Administration > Current Projects > Jobs to display the Job Status page. For more information, see “Jobs Page” on page 276. For the deleteEvents job, in particular, you may wish to make the following changes to the startup options:</p> <ul style="list-style-type: none"> ➤ By default, the value for the -commit option for the deleteEvents job is 100, meaning that the commit is done after the deletion of every 100 events. If you are deleting a large number of events, explicitly specify the parameter -commit n to commit after n event deletions. Otherwise, the database transaction log may overflow or the table may be locked for a long period of time. ➤ By default, the deleteEvents job and script do not allow deleting events in the last 48 hours. To override the default of 48 hours, use the -minage number-of-hours option. For example, -minage 10 prevents the job from deleting events from the last 10 hours.

 **Project Summary Page**

Description	The project summary lists the settings created for this new project. To access: Select Admin > TransactionVision > Administration > Projects , click New Project .
Important Information	General information about the Project wizard is available in “Project Wizard” on page 53.
Wizard Map	The Projects wizard contains: Project Wizard > Database Schema Setup Page > Assign Analyzer to Project Database Schema Page > Assign Communication Links to Analyzer Page > Assign Data Collection Filters to Communication Link Page > Create Jobs from Template Page > Project Summary Page

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
< Back	Click < Back to return to previous pages in the Project wizard.
Cancel	Click to exit the Project wizard without creating the new project.
Finish	Click Finish to confirm that the settings are correct and create the new project. The Projects page displays.



Welcome to TransactionVision Page

Description	Displays available projects and queries, for the user to select the active project and type of query. To access: Select Admin > TransactionVision > Administration > Select Project
-------------	---

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Projects	Lists the Projects available for selection, and their Active or Inactive status.
Queries	Lists available queries for selection, including Last 24 hours and All events .
User Information	Displays the User and the time of Last Logon.



Troubleshooting and Limitations

Projects missing new tables

When a new XDM file containing XML mappings to new tables is added, older projects not containing these new tables fail to get loaded by TransactionVision. To work around this, add the new tables to the older projects with the CreateSqlScript as follows:

```
CreateSqlScript[.sh|.bat] -c -e -t table-name schema-name
```

4

Database Schemas

This chapter includes the main concepts and reference information for TransactionVision database schemas.

Concepts

- Database Schemas Overview on page 65

Reference

- Database Schemas User Interface on page 66



Database Schemas Overview

When the Analyzer processes event data collected by sensors, it stores the results in database schemas assigned to TransactionVision projects. Each project must be assigned one database schema. For more information about Analyzers and Projects see “Analyzers” on page 71 and “Projects” on page 43.

Only users with administrative permissions to access the TransactionVision Administration page can manage database schemas. For information about permissions and setting them, see the *HP TransactionVision Deployment Guide* PDF.

Database Schemas User Interface

This section describes:

- Database Schema Manager on page 66
- Create a New Database Schema Page on page 68

Database Schema Manager

Description	Displays information for each schema available in your environment. To access: Select Admin > TransactionVision > Administration > Database Schemas
Useful Links	"Projects" on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Delete Database Schema	Click to delete a selected schema. You are prompted to confirm the deletion. Click Yes to proceed. NOTE: Deleting a database schema permanently erases all event data in projects associated with the database schema. This data is not recoverable, so use this command with caution. You may not delete a schema if an Analyzer is assigned to it.
New Database Schema	Click to create a new schema. See "Create a New Database Schema Page" on page 68 for field descriptions. You may also create a new schema when you create a project with the Project wizard. See "Projects" on page 43 for details.
Register Database Schema	Click to register an unregistered schema. Click the radio button next to the schema name to select it. The Registered in TransactionVision column then displays Yes for the selected schema.

GUI Element	Description
Analyzer	The host name and IP address of the Analyzer host that is assigned to the schema. Analyzer hosts are assigned to the schema by the Project wizard or the Update Analyzer page.
Create Tables	Click Create Tables to create tables for a selected schema that does not have project tables. This only displays if a schema is still registered with TransactionVision, but the schema tables in the database do not exist.
Database Schema	Schema name.
Has Project Tables	<p>Whether TransactionVision project tables have been created.</p> <p>In general, when a project is created using the TransactionVision UI/Job Server, the Registered in TransactionVision and Has Project Tables columns have the value Yes. However, if a database administrator creates or manages schemas/tables using database tools (instead of the TransactionVision UI/Job Server), a scenario may occur where the project tables exist, but not the corresponding entries in the administration tables of the TransactionVision web user interface.</p> <p>Similarly, a schema and its tables might be deleted using external database tools, but a project might still use the schema. In that case, this page provides the status of these schemas in relation to TransactionVision projects and allows you to fix any errors.</p>
Registered in TransactionVision	Whether the schema has been registered in TransactionVision. A schema is registered in TransactionVision when any project uses the schema. However, if all projects using the schema are deleted, it is still registered in TransactionVision.

GUI Element	Description
Tablespace	The database tablespace associated with the schema. This column displays only if you are using Oracle or DB2 as your database for TransactionVision. It does not appear if you are using Microsoft SQL Server. See “Create a New Database Schema Page” on page 68 for more information about tablespaces.
Used by Project(s)	The names of any projects to which the schema is assigned. Schemas are assigned to projects by the Project wizard. See “Projects” on page 43 for details.

Create a New Database Schema Page

Description	Enables you to create a new database schema. To access: On the Admin > TransactionVision > Administration > Database Schemas page, click New Database Schema .
--------------------	---

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to return to the Database Schema Manager without creating a new schema.
Database Schema Name	Enter the name of the new database schema. The schema name may contain only alphanumeric characters and the underscore character. It must begin with an alphabetical character.

GUI Element	Description
Database tablespace	<p>Select the database tablespace to use for the schema from the list of available tablespaces. The list depends on which tablespaces are defined in the database.</p> <p>The <Default> option is always in the list.</p> <ul style="list-style-type: none"> ➤ For Oracle, the list contains the default tablespace name (such as USERS) as well as the <Default> option. Therefore, you can select the default name or <Default> to choose the default tablespace. ➤ For DB2, the list does not contain the default tablespace name. Therefore, if no other tablespaces are defined, only <Default> appears in the list. In this case, the database determines which tablespace is the default. <p>NOTE: SQL Server does not have the concept of table spaces, but the file group is a similar concept. Though project tables can be distributed to different file groups by using custom DDL scripts, the creation of project tables in specific file groups through the TransactionVision web user interface or CreateSqlScript is not supported.</p>
Finish	<p>Click to register the new database schema in TransactionVision and to create all TransactionVision-related tables.</p>

5

Analyzers

This chapter includes:

Concepts

- Analyzers Overview on page 71
- Creating or Editing an Analyzer Host on page 72
- Configuring the Analyzer on page 72
- Stopping and Restarting the Analyzer Process on page 74
- Failure Mode on page 74

Reference

- Analyzers User Interface on page 83
- Troubleshooting and Limitations on page 93**

Analyzers Overview

TransactionVision uses the Analyzer service to perform the following tasks:

- Send configuration messages to Sensors and collect event messages from Sensors
- Process events from the Sensor, correlate them into transactions, and store them in the project schema.

The Analyzer uses an embedded RMI registry so that Analyzers may be controlled by the TransactionVision web user interface running on remote hosts.

Only users with required permissions to access the TransactionVision Administration page can manage services. See “TransactionVision User Permissions” on page 22 for information about permissions.

Creating or Editing an Analyzer Host

Create an Analyzer host to identify a host running the Analyzer service. You must create an Analyzer host to which you assign communication links for a TransactionVision project. You create an Analyzer host on the Create Analyzer page, which is accessed by clicking **New** from the “Analyzers Page” on page 84 or by using the “Project Wizard” on page 53.

Once an Analyzer host is created, you can modify the settings on the Update Analyzer page by clicking **Edit** from the “Analyzers Page” on page 84.

Configuring the Analyzer

This section describes:

- “Configuration Message Expiry” on page 72
- “Configure the Analyzer” on page 74

Configuration Message Expiry

The Analyzer depends on a database connection in order to maintain its configuration messages, which tell Sensors whether to collect events or not. When the database used by TransactionVision is shutdown, particularly for extended periods of time, configuration message expiry settings can affect how your Analyzer and Sensors perform. There are two ways the configuration message expiry can effect your collection of events:

- The configuration message expiry controls how long Sensors remain active after the database shuts down. If the database is down, the Analyzer is not able to generate new configuration messages; after the time specified in the configuration message expiry settings, Sensors are automatically disabled.
- The configuration message expiry setting also controls how often the Analyzer attempts to re-establish communication and send out a new configuration message. It is at the most configuration message expiry setting minutes after the database is restarted.

Reasons for a High Configuration Message Expiry

If it is critical that your Sensors do not miss collecting any events, even those that occur while the database is down, you would want to set your configuration expiry to a high value. This means that Sensors would continue to collect for this long after the Analyzer Goes Inactive, And Place Their Messages In The Message Queue. When The Analyzer finally reconnects and processes all the backed up messages on the queues.

Note: If the database were to be down for a long time, events would keep being backed up on the queues since they would never be processed. Make sure you have the necessary system resources set to be able to handle the backed up events if you have a high event volume, expect the database to be down for extended periods, and want to ensure that all events are saved during this period.

Reasons for a Low Configuration Message Expiry

On the other hand, if keeping all generated events during database downtime is not important, you could set your configuration expiry to a shorter value. In this way, Sensors stop collecting after the specified time period, and be disabled until the Analyzer is able to reconnect and send out a new message. This would prevent unnecessary events from backing up while the analyzer is unable to process them.

To force the Analyzer to immediately re-send a configuration message after the database has been shutdown, stop and restart the project.

 **Configure the Analyzer**

Analyzer configuration information is stored in the <TVISION_HOME>/**config/ services/Analyzer.properties** file. For more information about this file, see the *HP TransactionVision Deployment Guide* PDF.

 **Stopping and Restarting the Analyzer Process**

To completely stop the Analyzer process, use the following nanny command:

```
nanny stopService tv_analyzer
```

To restart the Analyzer, use the following nanny command:

```
nanny startService tv_analyzer
```

 **Failure Mode**

The Analyzer can process events in Failure Mode. In this mode, the Analyzer only stores event data of failed business transactions (or transactions violating SLA). For successful business transactions, only the corresponding business transaction rows are saved. Since this mode relies heavily on a meaningful value of the **result** attribute of the business transaction, it requires a transaction classification with suitable rules for the **result** column to be in place. The Failure Mode can be enabled on a per-project basis, allowing to run certain projects in Failure Mode while running others in the conventional Standard Mode.

The main goal of the Analyzer Failure Mode is to reduce the data storage requirement for a project. In a system that only carries a few percent of failed transactions, the Failure Mode can significantly reduce the amount of data that gets stored into the database. In addition, on less powerful database storage systems the reduced workload on the database level can also increase the performance of the Analyzer by up to 50%. The advantages of the Failure Mode comes at a price though: since all event data is only kept in memory, recovery from a system crash or processing errors is not possible.

A typical use case for the Failure Mode is for high-volume systems with short-lived transactions where statistics collection and failure alerting is the primary requirement.

This section describes:

- “Enabling Failure Mode” on page 75
- “How Failure Mode Works” on page 76
- “SLA Violations” on page 77
- “Tuning Considerations” on page 78
- “Advanced Configuration” on page 80

Enabling Failure Mode

To enable Failure Mode processing in the Analyzer, set the property **enable_failure_mode=on** in the `<TVISION_HOME>/config/services/Analyzer.properties` configuration file. By default, only the Standard Mode (**enable_dbcaching=on**) is enabled. In case all of your projects only use Failure Mode, the Standard Mode can be disabled to save some system resources.

Once the Failure Mode has been enabled in the Analyzer, one or more projects can be configured to use the new processing mode. This can either be done at project creation time by selecting **Only save events of failed transactions** in the project creation wizard or by editing an existing project and checking the same option in the project edit page. For more information on creating or editing a project, see “Projects” on page 43.

It is possible to switch the Analyzer processing mode while the Analyzer is running – once the mode is set and the Analyzer starts processing new events, those events are processed using the new processing mode.

How Failure Mode Works

When the Analyzer is processing a transaction in Failure Mode, all event, correlation and local transaction data is kept in memory until the transaction result is known. While the Standard Mode flushes all data continuously to the database, the Failure Mode only writes the business transaction data (**BUSINESS_TRANSACTION** and custom **XDM** transaction tables) for all incoming events to the database. The remaining data (event XML, user data, correlation and local transaction data) is stored in memory caches until either:

- The transaction result is set to **SUCCESS** by matching classification rules, in which case all data for this transaction kept in memory are discarded (except the business transaction data itself, which has already been flushed in the same way as in Standard Mode). Further events for this transaction are discarded as well.
- The transaction result is set to **FAILED** by the classification rule engine, in which case all data of this transaction is flushed to the database. Further events for this transaction are stored as well.
- The transaction violates a predefined SLA set for its class. The response time of a transaction is continuously monitored and transaction data is flushed to the database if this response times exceeds an SLA time value set in the business transaction table.
- The transaction result has still not been set even after a certain configurable time has passed, in which case the transaction is also considered to be “aged out”, and all data is flushed to the database. The SLA status column of the transaction (explained below) is set to **Aged Out**. In case an event is processed later which marks the transaction as **successful**, the SLA status is reset and the transaction is treated in the same way as a successful transaction - all further events are discarded.

Since the memory caches are limited in size, it is vital for the efficiency of the Failure Mode that the caches are cleared as quickly as new events are entering the Analyzer and being processed. A high system load, too small cache sizes, or keeping event data too long in the caches can all lead to cache overflows (indicated by a **FailureModeCacheOverflow** warning in the analyzer.log file). When this happens, the entire cache contents is written to the database slowing down the Analyzer event processing rate.

The timeout after which a transaction is considered to “age out” is dynamically computed from the sum of the transaction class SLA plus a configurable “process delay” value which is defined in **Analyzer.properties** file. This value specifies the maximum time the processing of any event can be ‘delayed’ within the Analyzer in respect to other events of the same transaction. There is a certain amount of lag between the time events in a transaction are generated at the sensor and the time events for a transaction are completely processed by the Analyzer. This lag is due to the effects of queuing in the WebSphere MQ infrastructure, event batching, multithreading, and context switches by the OS. So if for example, the transaction class SLA is defined as 5 seconds, and the “process_delay” is set to 10 seconds, any transaction whose result has not been determined 15 seconds after the first event has been processed are aged out.

Successful transactions still appear together with their attribute values on the Transaction Tracking Report in the GUI application, but it is not possible any more to navigate to any other views or drill down to the event level, since all this data has not been written to the database. But for failed and aged out transactions all data is available, and the report offers the same functionality as if those transactions had been processed in Standard Mode.

SLA Violations

In addition to the two scenarios mentioned above (failed and aged out), the Failure Mode can also be configured to write all event data in case a transaction violated the defined SLA for its class. In order to allow SLA violation tracking, the following section has to be enabled in Beans.xml:

```
<Module name="LogSLAViolationCtx" type="Context">
  <Module
class="com.bristol.tvision.services.analysis.eventanalysis.LogSLAViolationBean"
type="Bean">
    <Attribute name="WriteTxn" value="true"/>
    <Attribute name="LogLevel" value="Warning"/>
  </Module>
</Module>
```

When the context is enabled, the response time of each transaction is continuously monitored against the SLA value defined for its class, and the status is stored in the “sla_status” column of the business transaction. If the response time exceeds the SLA, a value of 1 (=Violated) is stored, otherwise a value of 0 (=None). For transactions which age out during the analysis process (as mentioned above), a value of 2 (=AgedOut) is stored into the column. When the attribute “WriteTxn” is set to true, transactions violating the SLA are treated like failed transactions, and their event data is stored in the database.

The attribute 'LogLevel' specifies the severity level of the log message and can be adjusted to the level of the corresponding log appender. Valid values are 'Info', 'Warning', and 'Error'. The default is 'Warning'.

Note: If you are using the SMTP log appender to send SLA alerts via email, the 'LogLevel' attribute has to be set to 'Error'.

Tuning Considerations

As mentioned previously, the Failure Mode only works efficiently if the Analyzer is able to clear out the memory caches as quickly as new data is coming in. If a cache overflow occurs, the system slows down considerably since all the data has to be written to the database. A system with frequent cache overflows does not perform as expected.

The ability of the Analyzer to keep the caches from overflowing is determined by multiple factors:

- The rate of which new events are coming into the system,
- The database performance,
- The configured size of the caches,
- The time an open transaction is kept in the system (determined by SLA + process delay).

The Analyzer Failure Mode is intended for systems with short living transactions, with SLA values in the range of seconds. As the SLA values get bigger, the transactions need to be kept in the caches much longer, and more data can accumulate in the caches. Such a system can still work properly - if the cache sizes are large enough, or if the load on the system is small enough.

There are certain issues with the Failure Mode one has to be aware of. First of all, since all event data is initially only kept in memory, this mode does not offer any recoverability. If an Analyzer crash, database error, or general system failure occurs, the data in the caches is lost. In standard processing mode, the XML event is always written to the database, so that in case of such a failure the missing analysis data can be recovered. In Failure Mode, the XML event data might have not been written to the database, so there is nothing to recover from.

Another issue arises when events for a successful transaction are being processed (or even entering the system) after the data for the transaction has been discarded from the caches. If these events contain an event which increases the response time of the transaction beyond the SLA value, and if writing of SLA violations is enabled, then this transaction only has partial data stored in the database, because previous event data has already been discarded. Note that this issue does not occur if the transaction classification contains time rules, since successful transaction are not be removed from the cache until the time rules have fired.

The second case in which late events can cause a problem occurs when the event is being processed so late that the correlation data for the corresponding transaction has already been pushed out of the lookup caches (which are, like the write caches, also limited in size). In this case, an event might not get correlated into the existing business transaction, and a new transaction is created instead. Increasing the size of the caches (event_based) could remedy the problem, but it might also be an indicator that the Failure Mode is not well suited for this kind of transactions.

Advanced Configuration

If Failure Mode is used, it is usually necessary to increase the cache sizes defined in **CacheSizes.properties**. We recommend to increase the sizes at least to twice of the size used in the standard processing mode.

The following other configuration parameters can be changed in **Analyzer.properties**:

- **failure_mode_thread_count**. Defines the number of flushing threads to use for Failure Mode. This number should be set high enough so that the cache data can be flushed fast enough, but low enough to not waste too many system resources. The default value is 2.
- **failure_mode_process_delay**. Defines the amount of time, in milliseconds, which gets added to the SLA value for determining the age out timeout and is a crucial value for the Failure Mode. If this value is set too low, a lot of successful transaction data unnecessarily ages out and is written to the database, which results in poorer performance, and wasted database disk space. If this value is set too high, too much data is held in the memory caches. If the cache sizes are not large enough for the system load, cache overflows occurs frequently and leads to very poor performance. The default value is 5 seconds.
- **failure_mode_discard_overflow**. In case of a cache overflow all the data in the caches is normally be written to the database. If the overflow was caused, for whichever reason, by a longer lasting system slowdown, the constant flushing of the cache could further worsen the situation and slow down the Analyzer even more. If the main concern is to avoid any chance of filling up the event queue in such situations, you can set the above parameter to true and force the discarding of all cache data in case of an overflow. The default value is false.

The diagnostic status output of the Analyzer can assist in finding appropriate configuration values for the Failure Mode. The following steps show how to get a meaningful status.

To get a meaningful status:

- 1** Set **detail=on** in the file <TVISION_HOME>/config/services/**Performance.properties**.
- 2** Enable **Failure Mode** in **Analyzer.properties**, and change the processing mode of the project in the GUI.
- 3** Run the project in Failure Mode for a sufficient amount of time on typical data.
- 4** Run **ServicesManager -diagstatus**.
- 5** Look for the section beginning with **FailureMode Cache Manager statistics for schema**.

Sample output

FailureMode Cache Manager statistics for schema TEST

[...]

Number of forced flushes: 0

FailureMode FlushThread#0: last flush: 4 rows 0 ms
 current flush time = 0 ms
 total flush cycles: 213
 cycles skipped: 3
 max scan depth: 0

FailureMode FlushThread#1: last flush: 0 rows 0 ms
 current flush time = 0 ms
 total flush cycles: 309
 cycles skipped: 309
 max scan depth: 0

Max txn processing time in schema TEST: 16188 ms (business txn 858)
 Avg txn processing time in schema TEST: 4842 ms
 Number of txns (success/failed/agedOut) : 0/129/6
 Number of overflow cache writes in schema TEST: 0

The values of interest are:

- **Number of forced flushes.** This is the number of times the Analyzer processing had to wait for the caches to be flushed to disk. If there are no forced flushes, the event processing and data flush happen mostly asynchronously with maximum performance. If the number of forced flushes is high, the processing threads have to suspend the analysis process until the flushing threads have emptied the caches often, which results in poor Analyzer performance. To minimize forced flushes, increase the number of flushing threads, or the size of the caches.
- **Total flush cycles / cycles skipped.** The ratio of both numbers gives an indication of how well the flushing threads are used. If the number of cycles skipped is very close to the total number of flush cycles, the threads are idle most of the time and the number of flushing threads can be reduced to save some system resources. If all flushing threads show a very low number of skipped cycles, it indicates that all threads are fully utilized, and additional flush threads might be necessary.
- **Maximum transaction processing time.** This is the longest time a transaction has been processing inside the Analyzer (from first to last event). The Failure Mode process_delay value should be chosen at least so big that the maximum processing time is still less than the sum of transaction class SLA and process_delay value to avoid unnecessary aging out of transaction data. If the value is set too high or too low, it can lead to performance issues as described earlier.
- **Number of overflow cache writes.** The number of times the complete cache contents had to be flushed to the database because a cache overflow occurred. Since overflow cache writes have a severe impact on the performance, it should always be tried to keep this value at zero by configuring the parameters explained above as required. If this still won't succeed it might indicate that the system load or transaction structure may not be well suited to be processed in Failure Mode.

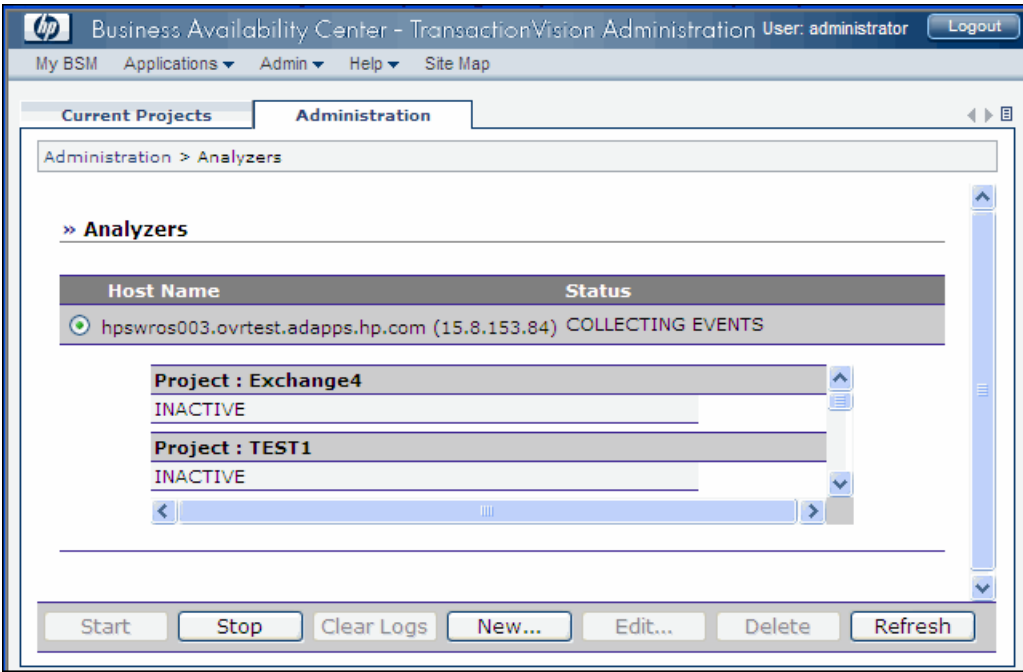
Analyzers User Interface

This section describes:

- Analyzers Page on page 84
- Create Analyzer Page on page 88
- Update Analyzer Page on page 92

 **Analyzers Page**

The following is an example of the Analyzers page.



Description	Displays information about each Analyzer available in your environment. To access: Select Admin > TransactionVision > Administration > Analyzers .
Useful Links	<ul style="list-style-type: none">➤ “Project Wizard” on page 53➤ “Create Analyzer Page” on page 88➤ “Update Analyzer Page” on page 92

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Clear Logs	Click to clear Analyzer log messages from the display.
Delete	Click to remove a selected Analyzer host from your environment. You may only delete an Analyzer host if no projects are assigned to it. You are prompted to confirm the deletion; click Yes to proceed.
Edit	Click to change the Analyzer configuration settings for a host. The Update Analyzer page appears. See “Create or Edit an Analyzer Host” on page 37 for information on modifying settings for an Analyzer host. If the Analyzer is collecting events, the Edit button is disabled; click Stop to stop the Analyzer to edit it.
New	Click to add a new Analyzer host to your environment. The Create Analyzer page appears. See “Create Analyzer Page” on page 88 for details.
Refresh	Click to update the view to reflect any changes since it was created.
Start	Click to begin collecting events on the selected Analyzer host. The Analyzer connects to assigned communication links and starts collecting events from the event queues. The status changes to COLLECTING EVENTS and details about event collection are displayed for each communication link for each active project.

GUI Element	Description
Stop	<p>Click to stop collecting events from communication link event queues on an Analyzer host.</p> <p>The Analyzer shuts down all collection threads without waiting for them to clean out events in the event queue. While the Analyzer is shutting down all the threads, its status is set to QUIESCING. Once all collection threads are shut down, the status changes to NOT COLLECTING EVENTS.</p> <p>NOTE: If an Analyzer ends abnormally, there may be XML events in the database that have not been processed for transaction analysis yet. When the Analyzer is started, it scans the event table for unprocessed events and recovers the statistics cache. To recover events for analysis, put failed events from the exception message queue back to the event queue with the ManageQueue utility. For more information about this utility, see “ManageQueue” on page 401.</p>
Host Name	Name and IP address of the host on which the TransactionVision Analyzer component is installed.
Status	Analyzer status. See “Analyzer Status” on page 87 for details.
<name and status>	Project name and status to which the Analyzer is assigned.

Analyzer Status

The following table describes possible status values:

Collecting Events	Description
COLLECTING EVENTS	<p>The Analyzer is started and collecting events. The following details display for each communication link for each active project:</p> <ul style="list-style-type: none"> ➤ Communication link name ➤ Number of configuration messages sent since the Analyzer was started or refreshed ➤ Number of events collected on the communication link ➤ Number of event packages sent on the communication link ➤ Amount of time skew across the communication link
FAIL	TransactionVision cannot connect to the Analyzer; an error message is provided.
RECOVERY	<p>The Analyzer is performing event analysis and statistics data recovery. This can happen when the Analyzer has been stopped abnormally (for example, the process has been killed), and not all event analysis and statistics data has been flushed to disk. In this case, the Analyzer reprocesses the last events to get the analysis and statistics data up-to-date before it starts the normal processing.</p> <p>If the status is RECOVERY, the number of events processed and the total number of events to process (for example, “230 of 2140 events processed”) is displayed. Click Refresh to update the number of events processed.</p>
QUIESCING	<p>The Analyzer is shutting down all collection threads in response to a stop command. Once all collection threads are shut down, the status changes to NOT COLLECTING EVENTS.</p>

Collecting Events	Description
NOT COLLECTING EVENTS	The Analyzer has been stopped and is not collecting events.
SENSOR COLLECTING EVENTS ONLY	The Analyzer has been stopped or has exited without sending stop configuration messages to the Sensors. Sensors continue to produce events, which collect on the event queue until the Analyzer and project are started again.



Create Analyzer Page

Description	<p>Enables you to create an Analyzer host to identify a host running the Analyzer service. You must create an Analyzer host to assign to the hose, communication links for a TransactionVision project.</p> <p>To access:</p> <ol style="list-style-type: none">1 On the Admin > TransactionVision > Administration > Analyzers, Analyzers page, click New2 Use the “Project Wizard” on page 53.
Useful Links	<ul style="list-style-type: none">➤ “Project Wizard” on page 53➤ “Update Analyzer Page” on page 92

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Analyzer Host	Enter the name of the host to add (for example, host1.bristol.com). The Analyzer service component must be installed on this host. Only one instance of the Analyzer may run on a given host; however, it can handle multiple projects.
Analyzer Port	Enter the port number where the Analyzer listens for RMI calls. The default value is 21100. If this system is changed on the server so that the Analyzer listens for RMI calls on another port, you must specify the correct port number.
Client Time Skew Servlet URL	Specify the URL of the servlet used by the Analyzer to determine the time skew between the Analyzer host and the host running the TransactionVision application server. The default URL points to the time server servlet in the Business Availability Center installation (http://bac_host:80/topaz/services/technical/time?alt=text/plain)
Configuration Message Expiry	Specify the expiration time in minutes of configuration messages sent from the Analyzer to the Sensors. When this value is set, configuration messages are sent two minutes before the old configuration message expires. Setting an expiration time limits how long Sensors continue to report events if the Analyzer is stopped in a non-proper fashion. When the configuration messages expire, Sensors stop reporting events. Set the Configuration Message Expiry to a value small enough so that the amount of events generated by an “orphaned” configuration message can fit into the event queues without causing major production issues. Specify value of -1 if you do not want configuration messages to expire. The Analyzer only sends configuration messages if filter criteria change, it has been told to start, or a communication link is changed. For guidance on setting this value, see “Configuration Message Expiry” on page 72.

GUI Element	Description
Connection Retry Delay	Specify the number of seconds to wait between attempts to connect to the queue manager for the communication link event queue. The default is 10 seconds.
Connection Retry Delay	Specify the number of seconds to wait between attempts to connect to the queue manager for the communication link event queue. The default is 10 seconds.
Database Exception Threshold	Specify the number of exceptions from the database to accept in the event collection threads before attempting to reconnect to the database.
Database Reconnect Interval	Specify the time in seconds to wait between tries to connect to the database. This value is used if the first attempt to connect to the database fails. The default is 10 seconds.
Database Schema	Check the box for each database schema you wish to assign to the Analyzer. The Analyzer is automatically assigned to all projects that use selected database schemas. You may only select project schemas that are not currently assigned to an Analyzer host. This field is available when editing an Analyzer, but not when creating a new Analyzer.
Event Get Wait Interval	Specify the number of seconds to wait for an event message. Increasing this number slows down the Analyzer response to commands. The default is 5 seconds.
Time Skew Check Interval	Specify the number of seconds between sending messages to check determine the time skew across the communication link. The default is 3600 seconds (one hour). At each specified interval, the Analyzer takes time skew samplings across the communication links and, if it is running on a different server, the TransactionVision application server.

GUI Element	Description
Time Skew History Size	Specify the number of previous time skew samplings to consider for the best time skew to use. The default is 24. Having a larger history size improves the chance of finding a time skew with smaller latency, but it also increases the chance of picking a time skew that may have occurred long ago to represent a time skew that may be more off due to clock drifting. A history size of 0 causes the time skew that was taken with the smallest latency to be used. This time skew is not discarded until a new time skew with equal or less latency is taken. Using this value disregards clock drift, which occurs frequently.
Time Skew Latency Threshold	Specify the amount of time in milliseconds to accept when sampling time skews. The default value is 100 milliseconds. If more than the specified time passes when waiting for a time skew reply, another time skew sampling is taken in hopes it has a lower latency. Any time skew samples that surpass the threshold are still kept in the time skew history because it is possible no samplings are under the suggested latency threshold. This threshold affects time skew sampling across communication links and to the application server.
Time Skew Message Priority	Specify the WebSphere MQ Message Priority on the Confirmation of Arrival (COA) time skew messages. The default value is -1, which causes the default priority to be used. Since configuration and event messages use the default priority, it is a good idea to set this value at a value higher than the default. Setting a higher value ensures that time skew messages travel through channels faster than other messages that may be backed up on transmission queues. Therefore, latency times when sampling time skews is reduced and results are more accurate. This value must not exceed the MaxPriority attributes of the queue managers.

GUI Element	Description
Time Skew Reply Wait Interval	Specify the number of seconds to wait for the message sent to determine the time skew across a communication link. The default is 30 seconds. If no response is received in this time, the time skew is disregarded.
Time Skew Retry Threshold	Specify the maximum number of time skew samples to take at a single time skew check interval. At each time skew check interval, TransactionVision continues to retry up to this number of times if the sampling surpasses the time skew latency threshold. The default number of retries is 8.



Update Analyzer Page

Description	Enables you to edit Analyzer host settings. To access: On the Admin > TransactionVision > Administration > Analyzers , Analyzers page, click Edit .
Useful Links	“Create Analyzer Page” on page 88

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Cancel	Click to return to the Analyzers page without making any changes.
Finish	Click to save the settings.
<all fields>	See “Create Analyzer Page” on page 88 for descriptions of each fields.

Troubleshooting and Limitations

This section provides the following troubleshooting and limitation information:

- “Failing to Process Events” on page 93
- “Exception Message Queue” on page 94
- “Storing and Using Raw Events” on page 95
- “Using the ManageQueue Utility” on page 95
- “Configuring the Analyzer” on page 96
- “Configuration for Storing Raw Events into a JAR File” on page 96
- “Configuration for Processing Raw Events from a JAR File” on page 96
- “Runtime Configuration” on page 97
- “Analyzer Recovery Mode” on page 97
- “Memory Resource Problems” on page 99
- “Analyzer States” on page 100
- “Event and Transaction Correlation Problems” on page 101
- “Batching Multiple Listener Events in a Transaction” on page 102
- “JDBC Batch Error Due to Incorrect Database Version” on page 104

Failing to Process Events

There are several reasons why the Analyzer might fail to process events. The most common are:

- Database errors, such as the database is down for a period longer than the reconnect setting or storage space is full
- Memory resource problems, such as OutOfMemory exceptions and cache overflows
- Errors caused by custom analysis beans

TransactionVision provides several methods for minimizing the impact if the Analyzer fails to process events for any reason:

- Placing events on the exception message queue
- Storing and using raw events
- Starting recovery mode for the Analyzer

Exception Message Queue

If the Analyzer fails for any reason, TransactionVision does not lose any analysis data. Instead, failed events are placed on the exception message queue defined in the communication link from which the events have been retrieved. See “Communication Links” on page 105, for instructions on setting the exception message queue when you create a communication link.

The number of events to be placed on the exception message queue depends on the value of the `batch_commit_count` property in the **Analyzer.properties** file. See the *HP TransactionVision Deployment Guide* PDF for more information about this file.

Once the cause of the failure is resolved, use the `ManageQueue` utility to replay the events on the exception message queue to the event queue, as in the following example:

```
ManageQueue.bat -iq TVISION.EXCEPTION.QUEUE -oq TVISION.EVENT.QUEUE  
-qm YOUR_QUEUE_MANAGER
```

For more information about this utility, see “`ManageQueue`” on page 401.

If the exception message queue contains events that have been failed after the XML data of the event has been written to the database (in the event analysis stage), those events cause a “Duplicate XML event” warning when they are replayed to the event queue. You can safely ignore this warning.

Storing and Using Raw Events

The TransactionVision sensors generate event messages and put them on the event queue. The Analyzer collects these event messages, processes them, and stores them in the database. The unprocessed event messages are called raw events.

TransactionVision provides two ways to store and re-process raw events in a JAR file: either using the ManageQueue utility or configuring the Analyzer during processing.

Using the ManageQueue Utility

- ManageQueue can get raw events from a queue (such as the event queue or exception queue) to a JAR file. This can be used to obtain events that the Analyzer failed to process (if those events are still available on the exception queue) for initial investigation.
- ManageQueue can read a JAR file and put the events back to a queue (such as the event queue).

For more information about using the ManageQueue utility, see “ManageQueue” on page 401. Specifically, see the following examples:

- Archive all messages on the event queue into a jar file while leaving them on the queue.
- Replay all events stored in a jar file for the TVISION schema to the **TVision – tv.qmgr** communication link.

Configuring the Analyzer

The Analyzer can be configured to:

- Automatically store raw events that are collected during Analyzer processing into a JAR file. This is done by setting the property **write_to_jar** in **Analyzer.properties**. This should be used after a problem in the Analyzer has occurred to obtain all events that are collected for a specific project.

Note: Since writing the JAR file has a performance penalty it should never be enabled during production.

- Read raw events from the JAR file and process them as if they originated from the event queue by setting the property **read_from_jar** in **Analyzer.properties**.
- Enable JAR reading and writing during the Analyzer runtime.

The following sections describe how to configure the Analyzer.

Configuration for Storing Raw Events into a JAR File

To configure a JAR file for storing raw events:

- 1 Edit **Analyzer.properties** and set **write_to_jar=true**.
- 2 Restart the Analyzer.

Once you start a project and collect events, a JAR file with the name **SCHEMA_raw_events.jar** is created in **\$TVISION_HOME/logs**.

Any events for this project are stored in the JAR file.

- 3 After you stop the project, the JAR file closes and can be used.

Configuration for Processing Raw Events from a JAR File

To configure a JAR file for processing raw events:

- 1 Edit **Analyzer.properties** and set **read_from_jar=true**.
- 2 Restart the Analyzer.

3 Start the project.

4 Copy the JAR file with the name **<SCHEMA>_raw_events.jar** into **\$TVISION_HOME/logs**. Replace **<SCHEMA>** with the actual schema name for the project.

The Analyzer gets the JAR file, processes all events in the file, and then renames the JAR file to **<SCHEMA>_raw_events.jar.processed.TIMESTAMP**. This avoids the Analyzer from processing the JAR file again the next time the project is started. Also, this allows multiple JARs to be processed (after each one has finished processing) for one project.

Runtime Configuration

Both JAR reading and writing can be enabled at runtime while the Analyzer is running.

- To enable JAR writing, execute:
ServicesManager -debugcmd write_jar on
- To disable JAR writing, execute:
ServicesManager -debugcmd write_jar off
- To enable JAR reading, execute:
ServicesManager -debugcmd read_jar on
- To disable JAR reading, execute:
ServicesManager -debugcmd read_jar off

Analyzer Recovery Mode

The Analyzer uses in-memory caches to accumulate and store all event analysis data as well as the statistics data for the Static Topology Analysis view. This data is flushed and written to disk frequently by separate flush threads. The analysis data is flushed continuously during Analyzer processing, and the statistics data is flushed at certain intervals which are specified by the **flush_interval** property in the **StatisticsCache.properties** file. See Appendix B in the *HP TransactionVision Deployment Guide* PDF for more information about this file.

While the Analyzer is running, it keeps track of which events the analysis and statistics data has been flushed for. If the writing of the cache data fails for any reason, the transaction and statistics data is incomplete, and the project schema is marked for RECOVERY mode. Failures that can lead to RECOVERY mode include database errors, event processing errors, and abnormal termination of the Analyzer process.

When a project whose schema has been marked for RECOVERY is started, the Analyzer goes into RECOVERY mode to regenerate the analysis and statistics data from the stored XML event. The time needed for the recovery process is proportional to the number of events that need to re-analyzed, and is usually slower than the normal event processing rate.

When the Analyzer is in RECOVERY mode, the Analyzer status displays the number of events processed and the total number of events to process (for example, “230 of 2140 events processed”). Click **Refresh** to update the number of events processed.

A log entry in the analyzer.log file also shows how many events need to be recovered. If TRACE Logging is configured for the Analyzer, the progress of the recover process is also logged. To enable trace logging for the TransactionVision Analyzer, set the value of the trace property in the Analyzer.properties file to on. For more information about this file, see Appendix B in the the *HP TransactionVision Deployment Guide* PDF.

Once the data has been flushed to disk successfully, the recovery mode of the schema is reset, and the project starts up normally the next time.

Note: If the data is not flushed because of a processing error which cannot be resolved, the Analyzer is not able to re-analyze the corresponding events and continue to start up in RECOVERY mode. In such a case, the recovery state can be cleaned out with the FlushStatusUtil utility. After the state has been cleaned, the Analyzer does not attempt to recover those events again. For more information about this utility, see “FlushStatusUtil” on page 399.

Memory Resource Problems

The amount of memory the Analyzer can use for event processing is bound by the JVM parameter “-Xmx” which defines the maximum JVM heap size. The default value is 256 Megabytes, but this value may be changed via the **service_jvm_flags** parameter in the **Analyzer.properties** configuration file. If the in-memory caches in the Analyzer are configured too large, the Analyzer tries to keep more data in memory than the JVM is capable of and at some point throws an OutOfMemory exception. In this case the Analyzer stops processing, and the error is logged to the **analyzer.log**. To resolve this problem either reduce the cache sizes in **CacheSize.properties** or increase the JVM heap size in **Analyzer.properties**.

Another possible memory related error is the overflow of the analysis caches. When this error occurs, the following error is logged to the analyzer.log file:

```
TransactionVision Error(EventAnalyzeFailed): Failed to analyze events from ('CSQ1'
'TVISION.EVENT.QUEUE') in database ('TVISION'); TransactionVision
Error(CacheOverflow): Overflow of DB cache RelationLookup in schema MQIMS
```

A cache overflow means that the amount of data for one or more transactions that the Analyzer is trying to store in memory exceeds the configured capacity of the caches. The caches most likely to overflow are “LocalTransaction” (the size of which may be configured by the ‘transaction_based’ parameter), “RelationLookup”, and “EventRelation” (both of which may be configured by the **event_based** size parameter).

The follow table summarizes possible causes of a cache overflow and includes instructions regarding how the problem may be corrected:

Cause of cache overflow	Corrective action required
The monitored transaction data is very complex and requires larger cache sizes.	Increase the appropriate cache size parameter in CacheSize.properties.
An erroneous custom correlation or transaction bean generates a huge amount of incorrect relational or transactional data for one event.	Check the logic of any custom beans, especially the generation of custom relation lookup keys (see “Event and Transaction Correlation Problems” on page 101 for details).

Analyzer States

You can view Analyzer status in either of the following ways:

- Choose **Admin > TransactionVision > Administration > Analyzers** to display the Analyzers page.
- Choose **Admin > TransactionVision > Current Projects > Project Status** to display the “Project Status Page” on page 51.

The following table describes the status values:

Status	Description
COLLECTING EVENTS	The Analyzer is started and collecting events.
FAIL	TransactionVision cannot connect to the Analyzer; an error message is provided.
RECOVERY	<p>The Analyzer is performing event analysis and statistics data recovery. This can happen when the Analyzer has been stopped abnormally (for example, the process has been killed), and not all event analysis and statistics data has been flushed to disk. In this case the Analyzer reprocesses the last events to get the analysis and statistics data up-to-date before it starts the normal processing.</p> <p>If the status is RECOVERY, the number of events processed and the total number of events to process (for example, “230 of 2140 events processed”) is displayed. Click Refresh to update the number of events processed.</p>
QUIESCING	The Analyzer is shutting down all collection threads in response to a stop command. Once all collection threads are shut down, the status changes to NOT COLLECTING EVENTS .
NOT COLLECTING EVENTS	The Analyzer has been stopped and is not collecting events.

Status	Description
SWITCHING PROCESSING MODE	The Analyzer is switching the current processing mode (normal / failure mode). All event processing is currently halted, and continues as soon as the switch is completed.
STATUS UNAVAILABLE – PREVIOUS COMMAND PENDING	The analyzer is still executing a previous command, and the current status can not be determined until this command completes.

Event and Transaction Correlation Problems

The Analyzer issues the following warnings when it detects potential problems correlating events into local transactions or correlating local transactions into business transactions.

Event Correlation

```
TransactionVision Warning(): Number of correlated events exceeds threshold: event ID='(1245,23)', count=30.
```

The event with the specified ID has been correlated into a message path together with an unusually large number of related events. If you are using custom correlation beans, this warning could indicate that the logic in the custom beans does not work correctly.

If no custom correlation logic is in place, this warning could indicate that the message data in the events used for correlation are not well-suited for the standard correlation, and custom correlation may be necessary. For example, applications may re-use WebSphere MQ message IDs.

Transaction Correlation

TransactionVision Warning(): Number of local transactions in one business transaction exceeds threshold: business txn ID='32', count=1000.

The business transaction with the specified ID consists of an unusually large number of local transactions. If you are using custom analysis beans, this warning could indicate that the logic in the custom beans does not work correctly.

If no custom analysis logic is in place, this warning could indicate that the nature of the message data might not be well-suited for the standard analysis algorithms, and custom correlation/transaction logic may be necessary.

Batching Multiple Listener Events in a Transaction

In environments where the WebSphere MQ channels have a non-zero BATCHINT (batch interval) parameter, events are batched up by the WebSphere MQ listener process and committed after n events. This causes TransactionVision to group together business transactions that would normally be separate.

To work around this issue, a local transaction definition file can be defined that splits up listener messages into separate local transactions or unit of works (UOW), regardless of their actual commit point.

The following is a sample definition file. Refer to the *HP TransactionVision Advanced Customization Guide* PDF for more information about this file. The following rules may be further customized to include specific queues, queue managers or event brokers.

```
<LocalTransactionDefinition>

  <LocalTransactionType dbschema="" hasMultiTracking= "false" >
    <Match xpath="/Event/StdHeader/TechName" operator= "EQUAL"
value="MQSERIES"/>
    <Match xpath="/Event/StdHeader/ProgramName" operator= "EQUAL"
value="AMQRMPPA"/>

  <LocalTransactionAttributes>

    <Attribute name="LookupKey">
      <Path>/LocalTransaction/LookupKey</Path>

      <ValueRule name="SetLookupKey">

        <Value type="XPath">/Event/EventID/ @programInstID</Value>

        <Value type="Constant">-</Value>
        <Value type="XPath">/Event/EventID/ @sequenceNum</Value>

      </ValueRule>

    </Attribute>

  </LocalTransactionAttributes>

</LocalTransactionType>

</LocalTransactionDefinition>
```

JDBC Batch Error Due to Incorrect Database Version

The following exception in the Analyzer logs indicates that DB2 8.1 FixPack 11 (also known as DB2 8.2 FixPack 4) is not installed:

```
2006-07-05 15:15:28,562 [MQJava0-July05-11] ERROR AppLog - TransactionVision
Error(DBFailedEventStore): Failed to store events from ('chenhpc.tv1.manager'
'TVISION.EVENT.QUEUE') in database ('tv'): Non-atomic batch failure. The batch was
submitted, but at least one exception occurred on an individual member of the batch.
Use getNextException() to retrieve the exceptions for specific batched elements.
```

```
2006-07-05 15:15:28,562 [MQJava0-July05-11] ERROR AppLog - TransactionVision
Error(ProcessEventError): Error while processing event.
```

```
2006-07-05 15:15:28,562 [MQJava0-July05-11] ERROR AppLog - TransactionVision
Error(JDBCBatchError): Error while executing JDBC batch operation 'INSERT INTO
JULY05.USER_DATA (proginst_id, sequence_no, data_num, type, ccsid, event_time,
user_data) VALUES (?, ?, ?, ?, ?, ?)' - SQL Error : com.ibm.db2.jcc.b.vd: Non-atomic
batch failure. The batch was submitted, but at least one exception occurred on an
individual member of the batch. Use getNextException() to retrieve the exceptions for
specific batched elements.
```

6

Communication Links

This chapter includes:

Concepts

- Communication Links Overview on page 106
- Understanding Communication Links on page 106
- TransactionVision Queue Requirements on page 110
- Communication Link Templates on page 117
- Creating New Communication Links or Templates on page 120

Tasks

- Create a New Communication Link or Template for WebSphere MQ on page 121
- Create a New Communication Link or Template for TIBCO EMS on page 130
- Create a New Communication Link or Template for Progress SonicMQ on page 135
- Create a New Communication Link or Template for BEA WebLogic JMS on page 140
- Create a New Communication Link or Template for Other JMS Provider on page 142
- Test a Communication Link on page 144

Reference

- Communication Links User Interface on page 145

Troubleshooting and Limitations on page 212

Communication Links Overview

TransactionVision uses WebSphere MQ, TIBCO EMS or SonicMQ for communication between Sensors and the Analyzer. Before a Sensor can collect events and forward them to the Analyzer, you must create a communication link between the Sensor and the Analyzer. A communication link defines the message queues used to pass configuration and event messages between the Sensor and Analyzer.

Only users with required permissions may create or edit communication links. See “TransactionVision User Permissions” on page 22 for information about permissions.

Understanding Communication Links

This section describes:

- “Communication Link Queues” on page 106
- “Communication Link Configurations” on page 109

Communication Link Queues

TransactionVision communication links require two message queues for the exchange of messages between a Sensor and the Analyzer:

- A configuration queue that Sensors monitor for configuration messages from the Analyzer
- An event queue that the Analyzer monitors for event messages from sensors

If an event fails to be processed by the Analyzer, the event is removed from the event queue and put on the exception message queue.

This section describes the following queues:

- “Configuration Queues” on page 107
- “Event Queues” on page 107
- “Exception Message Queues” on page 108

Configuration Queues

When TransactionVision begins recording events for a project, the Analyzer sends a configuration message to the configuration queues defined in the active communication links for the project. This configuration message defines the data collection filter conditions for the project, the name of the queue that event messages should be sent to, and the expiration time of the configuration message itself.

When a sensed event occurs, the Sensor checks to see whether a configuration queue exists. If no configuration queue exists, the Sensor skips the event generation code and returns control to the sensed application. If a configuration queue does exist, the Sensor reads all configuration messages on the queue, deleting any that have expired.

Event Queues

The Sensor evaluates the monitor event against the data collection filter conditions specified in the configuration message. If all filter conditions are met, the Sensor creates an event entry message about the API call. The Sensor then delegates control to the user application to finish the real work. When the work finishes, the Sensor creates an event exit message, and then passes the return on to the calling application. Entry and exit events are grouped into a single message to be put on the specified event queue. There may be more than one configuration message, so event messages may be put on more than one event queue.

For WebSphere MQ monitoring, Sensors use the same configuration queue for all applications that make Sensed API calls on the same queue manager. However, data collection filter conditions such as host name or program name may limit the number of event messages sent to the Analyzer.

After the first successfully monitored event, the Sensor continues to check the configuration queue periodically for new configuration messages.

While the Analyzer is recording events, it reads messages from the event queues into TransactionVision project schemas. It also sends new configuration messages to replace expiring configuration messages or to change the data collection filter conditions. When the Analyzer stops recording events, it sends a new configuration message to the configuration queue indicating that no more events should be recorded. The Analyzer may stop recording events for any of the following reasons:

- A user stops event recording from the TransactionVision web user interface.
- The end time specified in the data collection filter is reached.

Note: The channel that TransactionVision uses for event queue communications should have the channel property **CONVERT** set to **NO**. If this property is set to **YES**, event messages are discarded to the exception message queue. If the channel is being shared by queues that have messages that require the property to be set to **YES**, please create a separate channel specifically for the event queue.

Exception Message Queues

If an event fails to be processed by the Analyzer, either because of an invalid event or an unexpected error or exception, the event is removed from the event queue and put on the exception message queue defined in the communication link.

Any fatal Java errors (such as **OutOfMemoryError**) causes the event to be put on the exception message queue and events are no longer pulled off the event queue until the project is restarted. If the exception message queue name is blank, the failed event are discarded.

You can use the **ManageQueue** utility to transfer messages from the specified exception message queue to enable the Analyzer to attempt to reprocess the failed events. See “ManageQueue” on page 401, for more information.



Communication Link Configurations

This section describes configuration requirements for the following communication link types:

- “WebSphere MQ Configurations” on page 109
- “TIBCO EMS Configurations” on page 109
- “SonicMQ Configurations” on page 110
- “WebLogic JMS Configurations” on page 110

WebSphere MQ Configurations

For Sensors that use WebSphere MQ for communication with the Analyzer, configuration and event queues may reside on the same queue manager or on different queue managers. The following are possible configurations for communication links that use WebSphere MQ:

- Both the Sensor and the Analyzer communicate directly with the same queue manager.
- The Sensor uses one queue manager and the Analyzer uses a different one.
- The Sensor uses one queue manager, and the Analyzer sends configuration messages to one queue manager and monitors an event queue on a different queue manager.

The configuration queue used by the WebSphere MQ Sensor must be on the same queue manager used by the application being monitored.

TIBCO EMS Configurations

For Sensors that use TIBCO EMS for communication, the configuration queue and event queue must use the same EMS server to make a connection. However, you can use the TIBCO routing feature to access route home queues on different EMS servers. The following are possible configurations for Sensors that use TIBCO EMS:

- Both the Sensor and the Analyzer communicate directly with the TIBCO EMS server.

- The Sensor uses one TIBCO EMS server, and the Analyzer uses a different one. The configuration queue and/or event queue on one TIBCO EMS server is defined as a route to the TIBCO route home queue of the same name on the other TIBCO EMS server.
- The Sensor uses one TIBCO EMS server, and the Analyzer sends configuration messages to one TIBCO EMS server but use another TIBCO EMS server to retrieve events.

SonicMQ Configurations

For Sensors that use SonicMQ for communication, the configuration queue and event queue must use the same SonicMQ server to make a connection. However, you can use SonicMQ's Dynamic Routing Architecture to access queues on different SonicMQ brokers. The following are possible configurations for Sensors that use SonicMQ:

- Both the Sensor and the Analyzer communicate directly with the SonicMQ broker.
- The Sensor uses one SonicMQ broker, and the Analyzer uses a different one. The configuration queue and/or event queue on one SonicMQ broker is listed as a routing definition to the queue on the other SonicMQ broker.
- The Sensor uses one SonicMQ broker, and the Analyzer sends configuration messages to one SonicMQ broker but use another SonicMQ broker to retrieve events.

WebLogic JMS Configurations

For Sensors that use WebLogic JMS for communication, the configuration queue and event queue must use the same WebLogic JMS server to make a connection.

TransactionVision Queue Requirements

When creating queues for use with TransactionVision, you must specify attributes such as maximum message size and maximum queue depth. The settings you choose depend on factors such as the number of Analyzers using a queue manager, the number and size of messages used by monitored applications and the number of data collection filter conditions you specify.

In addition, TransactionVision event and configuration queues must have specific queue parameters and security permissions.

This section describes TransactionVision queue requirements for:

- “WebSphere MQ” on page 111
- “SonicMQ” on page 114
- “TIBCO EMS” on page 116



WebSphere MQ

Following are the TransactionVision queue requirements for WebSphere MQ.

Message Length

The message length requirements for configuration queues are small. A default configuration message uses approximately 500 bytes. However, specifying data collection filter conditions increases the size of configuration messages. A minimum message length of 10,000 bytes is required.

The length of event messages can vary greatly, depending on the parameters passed to the API, the user data passed through the calls, and data collection filter conditions. For event queues, it is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow an additional 4000 bytes for event information.

Queue Depth

For configuration queues, the default queue depth should be sufficient. Analyzers send configuration messages to start data collection, to change data collection filter conditions, to end data collections, and to determine time skews across hosts. The number of configuration messages sent to a configuration queue depends on the number of Analyzers sending messages to the queue.

Event queue depth should be set to a value that matches the queue manager storage space based on the average event message size. It should be adequate to handle the peak volume. It is recommended to use the largest queue depth possible for the event queue and, in the case of remote communication link configurations, all queues in between. In most cases, Sensors put event messages on the event queues faster than Analyzers retrieve them. Insufficient queue depth may result in event queues filling up and Sensors slowing down waiting for the Analyzer to catch up. Increasing the event queue depth helps prevent this situation.

Event Queue Manager

It is best to have the actual event queue hosted on a queue manager other than the production queue manager, so that issues related to event queues (such as running out of disk storage due to event backlog) does not affect normal operations.

Note: The channel on which the TransactionVision event queue is located should have the channel property **CONVERT** set to **NO**. If this property is set to **YES**, event messages are discarded to the exception message queue. If the channel is being shared by queues that have messages that require the property to be set to **YES**, please create a separate channel specifically for the event queue.

Event Queue Storage

Event queue storage should match the event queue depth based on the event message size. It should not exceed the storage capability of the queue manager host system. This is especially important when the event queue and TransactionVision exception message queue are hosted by the production queue manager because running out of storage space may stop the queue manager completely.

Event Queue Message Persistency

The event queue message persistency property should match the event collecting policy.

Queue Parameters

- TransactionVision configuration and event queues must have the Default Share Option (**DEFSOPT**) set to **SHARED**.
- TransactionVision configuration and event queues must allow Shared Access (**SHARED**).
- TransactionVision configuration and event queues must allow Get (**GET(ENABLED)**) and Put (**PUT(ENABLED)**) operations.
- TransactionVision configuration queues must allow messages of at least 10,000 bytes in length (**MAXMSGL**).
- TransactionVision event queues must allow messages of at least 10,000 bytes in length (**MAXMSGL**).
- All channels (including client channels) must have **MAXMSGL** of at least 10,000 bytes.

Sensor Security Permissions

Userids of programs using the Sensor require the following:

- GET and BROWSE authority to the TransactionVision configuration queue.
- PUT authority to the TransactionVision event queue.

Analyzer Security Permissions

Users of the Analyzer require the following:

- PUT, GET, BROWSE, and INQ authority to all TransactionVision configuration and event queues.
- Passid permission to test communication links.

Note that Analyzer users can be given authority to only subsets of the configuration/event queues if you want to limit access to certain Sensors.

If a user does not have the required access permission on the configuration queue (**TVISION.CONFIGURATION.QUEUE** by default), the following message is found in TransactionVision logs:

```
TransactionVision Sensor: cannot open event queue  
TVISION.CONFIGURATION.QUEUE on queue manager merce.es1.manager: Not  
authorized for access.
```

To set access permission of **TVISION.CONFIGURATION.QUEUE** for a user (tester in this example), run the following command:

```
setmqaut -m queuemanager -n TVISION.CONFIGURATION.QUEUE -t queue -p tester.  
+get +browse +put +inq
```

It returns a message that the command completed successfully.

To display the access permission of **TVISION.CONFIGURATION.QUEUE** for a user (tester in this example), run the following command:

```
dspmqaut -m queuemanager -n TVISION.CONFIGURATION.QUEUE -t queue -p tester
```

It returns output similar to the following example:

```
Entity tester has the following authorizations for object  
TVISION.CONFIGURATION.QUEUE:  
get  
browse  
put  
inq
```



SonicMQ

Following are the TransactionVision queue requirements for SonicMQ.

Event Queue SonicMQ Broker

It is best to have the actual event queue hosted on a SonicMQ broker other than the production broker, so that issues related to event queues (such as running out of disk storage due to event backlog) does not affect normal operations.

Event Queue Storage

Event queue storage should not exceed the storage capability of the SonicMQ broker host system. This is especially important when the event queue and TransactionVision exception message queue are hosted by the production SonicMQ broker. The behavior of the production SonicMQ broker may be unpredictable if you run out of storage space.

Queue Parameters

- TransactionVision configuration and event queues must allow receive, send, and browse operations that are defined at TransactionVision setup, to be granted to the user.
- TransactionVision configuration queues must allow messages of at least 10,000 bytes in length (**maxbytes**).
- TransactionVision event queues must allow messages of at least 10,000 bytes in length (**maxbytes**).

Sensor Security Permissions

The user used by the TransactionVision Sensor to connect to configuration and event queues must have the following permissions:

The user name that is used by the TransactionVision Sensor must have the following permissions to use the configuration and event queues:

- receive and browse permission to the TransactionVision configuration queue
- send permission to the TransactionVision event queue

Furthermore, to allow the TIBCO EMS sensor to resolve route and bridge information correctly, the user name provided in TIBCO EMS server setup must have administrative permission.

TIBCO EMS

Following are the TransactionVision queue requirements for TIBCO EMS.

Event Queue TIBCO EMS Server

It is best to have the actual event queue hosted on a TIBCO EMS server other than the production EMS server, so that issues related to event queues (such as running out of disk storage due to event backlog) do not affect normal operations.

Event Queue Storage

Event queue storage should not exceed the storage capability of the TIBCO EMS server host system. This is especially important when the event queue and TransactionVision exception message queue are hosted by the production TIBCO EMS server because running out of storage space may cause unpredictable behavior of the production EMS server.

Event Queue Message Persistency

The event queue message persistency property should match the event collecting policy.

Queue Parameters

- TransactionVision configuration and event queues must allow receive, send, browse operations are granted to the user that are defined at TransactionVision setup.
- TransactionVision configuration queues must allow messages of at least 10,000 bytes in length (**maxbytes**).
- TransactionVision event queues must allow messages of at least 10,000 bytes in length (**maxbytes**).

Sensor Security Permissions

The user used by the TransactionVision Sensor to connect to configuration and event queues must have the following permissions:

- receive and browse permission to the TransactionVision configuration queue

- send permission to the TransactionVision event queue

Furthermore, to allow the TIBCO EMS sensor to resolve route and bridge information correctly, the user name provided in TIBCO EMS server setup must have administrative permission

Analyzer Security Permissions

The user used by the TransactionVision Analyzer to connect to configuration and event queues must have the following permissions:

- send, receive, browse permission to all configuration and event queues
- view-destination permission on event queues to check the pending unprocessed event count

Communication Link Templates

TransactionVision provides two types of communication link templates that are displayed on the “Global Communication Link Templates Page” on page 209 (select **Admin > TransactionVision > Administration > Communication Link Templates**):

- User-created Global Communication Link Templates, see “User-Created Global Communication Link Templates” on page 117 for details.
- TransactionVision-Created Global Communication Link Templates, see “TransactionVision-Created Templates” on page 118 for details.

See “Global Communication Link Templates Page” on page 209 for details on the contents of this page.

User-Created Global Communication Link Templates

TransactionVision enables you to define global communication link templates to be used as the basis for project communication links. A global communication link template identifies the configuration and event queues and queue managers for the Sensor and the Analyzer.

You can create one or more communication link templates to be used as communication links for specific projects in the following ways:

- On the Global Communication Links page, click **New** to create a new template using the Communication Link Wizard. See “Creating New Communication Links or Templates” on page 120 for instructions.
- On the Global Communication Links page, select an existing user-created global communication link and click **Copy** to create a copy of a communication link named Copy of original_link_name. Select the copy and click **Edit** to change settings on the Edit Communication Link Template page (see “Edit Communication Link Template” on page 184 for details).
- In the Project Wizard to create a new project (see “Project Wizard” on page 53), select a global communication link template, and project communication links based on the template are created.

The project communication link is a copy of the global communication link template. Changes made to the template do not affect project communication links and changes to project communication links do not affect the template.



TransactionVision-Created Templates

TransactionVision provides the following three global communication link templates:

- “TransactionVision SonicMQ” on page 118
- “TransactionVision HTTP” on page 119
- “TransactionVision RUM” on page 120

TransactionVision SonicMQ

The TransactionVision SonicMQ template is a global communication link template that is provided to help you get started up quickly. It defines a communication link to the SonicMQ, which is installed with TransactionVision.

You can copy and edit this template in the same way you would any other user-created template. See “User-Created Global Communication Link Templates” on page 117 for information on copying and editing the global communication link template.

The TransactionVision SonicMQ communication link is used by the BEA Tuxedo, the NonStop TMF, and the .NET Sensors to send events to the Analyzer. It may also be used by any Java Agent as a turn-key solution for communication between the Analyzer and the Java Agent(s).

TransactionVision HTTP

The TransactionVision HTTP communication link is a special communication link implementing HTTP over TransactionVision's provided SonicMQ.

You can copy and edit this template in the same way you would any other user-created template. See “User-Created Global Communication Link Templates” on page 117 for information on copying and editing the global communication link template.

The TransactionVision HTTP communication link is used by the BEA Tuxedo and NonStop TMF Sensors to send events to the Analyzer. This template can also be used by custom User Event Sensors to do advanced filtering in a custom, user-designed Sensor.

The event queue is called **HTTP.EVENT.QUEUE**, and the exception queue is called **HTTP.EXCEPTION.QUEUE**.

Sensors post HTTP messages to the HTTP acceptor URL **http://<hostname>:21113/tv_http** and SonicMQ broker receives the messages and forwards them to **HTTP.EVENT.QUEUE** as JMS messages, which are then read by the Analyzer.

To configure the HTTP communication link to use SSL, you would change the URL to be **https://<hostname>:21114/tv_http**.

For more information on the HTTP acceptor, see the Sonic MQ product's Deployment Guide (*mq_deploy.pdf*).

TransactionVision RUM

The TransactionVision RUM template is for projects that collect Real User Monitor (RUM) events. Currently the integration of TransactionVision with Real User Monitor data is only supported for a single TransactionVision project. Therefore, you can only have one TransactionVision project that is set up to be integrated for RUM event processing, to which one RUM communication link is assigned.

You can copy and edit this template in the same way you would any other user-created template. See “User-Created Global Communication Link Templates” on page 117 for information on copying and editing the global communication link template.

Creating New Communication Links or Templates

The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template for the following types (see the corresponding tasks for instructions):

- WebSphere MQ — see “Create a New Communication Link or Template for WebSphere MQ” on page 121
- TIBCO EMS — see “Create a New Communication Link or Template for TIBCO EMS” on page 130
- Progress SonicMQ — see “Create a New Communication Link or Template for Progress SonicMQ” on page 135
- BEA WebLogic JMS — see “Create a New Communication Link or Template for BEA WebLogic JMS” on page 140
- Other JMS Provider — see “Create a New Communication Link or Template for Other JMS Provider” on page 142

Note: The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues.

Create a New Communication Link or Template for WebSphere MQ

This task describes how to create a new communication link or template for the WebSphere MQ type:

This task includes the following steps:

- “Prerequisites” on page 122
- “Assign Name” on page 122
- “Sensor Connection for WebSphere MQ” on page 122
- “Analyzer Connection for WebSphere MQ” on page 123
- “Remote Queue Connection Choice for WebSphere MQ” on page 124
- “Remote Queue Connection for WebSphere MQ” on page 124
- “Maximum Message Length for WebSphere MQ” on page 125
- “Event Retry for WebSphere MQ” on page 126
- “Miscellaneous Information for WebSphere MQ” on page 127

1 Prerequisites

A project has been selected. See “Projects” on page 43 for details.

2 Assign Name

Enter the name of the new communication link and select the communication link type, then click **Next** to define the Sensor connection.

For the communication link type, select the type of transport to use for communication between the Sensor and the Analyzer. See “Communication Link Wizard” on page 146 for field descriptions of this page.

3 Sensor Connection for WebSphere MQ

If you selected IBM WebSphere MQ as the communication link type, the Sensor Connection page displays. As you provide information on this page, the diagram is updated to provide a graphical representation of the Sensor connection for the new communication link.

Following is an example of the Sensor Connection page.

HomeViewsCurrent ProjectAdministrationReportsLogoutHelp

» Global Communication Link Template Creation Wizard -- Step 2: Sensor Connection

Configuration Queue

Queue manager sensor will connect to:localhost:1414

Queue Sensor will be receiving configuration messages from:TVISION.CONFIGURATION.QUEUE

☒ Use Default Configuration Queue

Event Queue

Queue Sensor will be sending event messages to:TVISION.EVENT.QUEUE

Is event queue in a cluster?☐ Cluster Queue

Transmission Queue/Queue Manager Alias that Sensor will use to open event queue:☒ Use Local Queue Manager

Configuration of WSMQ[1]

localhost:1414

Sensor

Configuration Queue
TVISION.CONFIGURATIO...

Event Queue
TVISION.EVENT.QUEUE

Analyzer

< Back

Next >

Cancel

- a Specify the name of the queue manager which Sensors connect to. This queue manager must be used by applications to be monitored.

- b** Specify the name of the queue that Sensors should check for configuration messages. This queue must be a local queue on the specified queue manager. Check **Use Default Configuration Queue** to use the default configuration queue name TVISION.CONFIGURATION.QUEUE. To use a different queue name, uncheck **Use Default Configuration Queue** and enter the queue name.
- c** Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.

If this queue is in a cluster, check **Cluster Queue**. If you check this setting, the Transmission Queue setting is not displayed.

If this queue is not in a cluster and is a local queue, check **Use Local Queue Manager**. If it is not a local queue, uncheck **Use Local Queue Manager** and specify either the name of the transmission queue that has the same name as the remote queue manager or a queue manager alias that resolves to the remote queue manager. This method of sending messages is described in the *WebSphere MQ Application Programming Guide* (SC33-0807-08).

- d** Click **Next** to specify whether the Analyzer should connect to the same queue manager to which the Sensors connect.

4 Analyzer Connection for WebSphere MQ

For the WebSphere MQ communication link type, the Remote Queue Connection Question page displays giving you a choice.

If Yes	If No
<p>Select Yes if the Analyzer connects to the same queue manager that WebSphere MQ applications connect to.</p> <p>Click Next to specify the maximum message length in “Maximum Message Length for WebSphere MQ” on page 125.</p>	<p>Select No if the Analyzer connects to a different queue manager than WebSphere MQ applications connect to.</p> <p>Click Next to specify whether to use one or two remote queue managers in Remote Queue Connection Choice for WebSphere MQ.</p>

5 Remote Queue Connection Choice for WebSphere MQ

If you selected **No** on the Analyzer Connection Choice page, the Remote Queue Connection page displays:

If Yes	If No
Select Yes if the Analyzer should use the same queue manager for both configuration and event messages. See Same WebSphere MQ Server in the “Remote Queue Connection for WebSphere MQ” on page 124 step.	Select No if the Analyzer should use different queue managers for configuration and event messages. See Separate WebSphere MQ Managers in the “Remote Queue Connection for WebSphere MQ” on page 124 step.

Click **Next** to specify remote queue connections on the Remote Queue Connection page.

6 Remote Queue Connection for WebSphere MQ

Specify the same WebSphere MQ server or separate WebSphere MQ managers, depending on your answer in “Remote Queue Connection Choice for WebSphere MQ” on page 124.

Same WebSphere MQ Server

If you answered **Yes** on the Remote Queue Connection Choice Question page, specify a single remote queue server.

- a** Specify the name of the queue manager the Analyzer is to use for both configuration and event messages.
- b** Specify the name of the queue that the Analyzer should send configuration messages to. This queue must be a remote queue for the configuration queue on the Sensor queue manager.
- c** To use a transmission queue or queue manager alias, check the box next to Use Transmission Queue or Queue Manager Alias. Enter the name of the transmission queue that has the same name as the remote queue manager or a queue manager alias that resolves to the remote queue manager.

- d** Specify the name of the queue that the Analyzer should monitor for event messages. This queue must be a local queue and should not be used by applications other than TransactionVision.
- e** Click **Next** to specify the maximum message length.

Separate WebSphere MQ Queue Managers

If you answered **No** on the Remote Queue Connection Choice Question page, specify separate remote queue managers.

- a** Specify the name of the queue manager the Analyzer is to use for sending configuration messages to Sensors.
- b** Specify the name of the queue that the Analyzer should send configuration messages to. This queue must be a remote queue for the configuration queue on the Sensor queue manager.
- c** Specify the name of the queue manager the Analyzer is to monitor for event messages from Sensors.
- d** Specify the name of the queue on the event queue manager that the Analyzer should monitor for event messages. This queue must be a local queue and should not be used by applications other than TransactionVision.
- e** Click **Next** to specify the maximum message length.

7 Maximum Message Length for WebSphere MQ

Specify the maximum event message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.

It is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow an additional 4000 bytes for event information.

Click **Next** to specify event delivery retry information for the communication link on the Event Retry Options page.

8 Event Retry for WebSphere MQ

The WebSphere MQ Sensor relies on the sensed application to handle connections to the event and configuration queue manager. However, other Sensors can be configured to try to reconnect if the event or configuration queue manager goes down and the Sensor loses the connection. Choose one of the following actions if the event or configuration connection is lost:

- If a Sensor is unable to connect to the event queue or is unable to put events on the event queue, events can be lost. To avoid this situation, you can configure retry options for event connection and delivery.
- Select one of the following actions for the Sensor to take if connection to the event queue fails:

Option	Description
Retry forever	Continues to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.
Retry once per TransactionVision event	Retries once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across transaction events	Continues to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision event	Continues to retry at the specified interval until the timeout is reached.

- The default setting is to retry forever every 10 seconds.

- To enable Sensors to retry if they fail to put an event on the event queue, check **Retry Event Package Delivery on Failure**. If you enable retry mode, set the following options to control retry attempts:

Option	Description
Retry Timeout	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.

- The default setting is not to retry event package delivery on failure. Set a value appropriate for your event collecting strategy. In a production environment, uncheck **Retry Event Package Delivery on Failure** for the safest setting.
- To avoid any delays in the application when a failure occurs, setting Event Delivery Retry on Failure to not retry overrides the Event Connection Retry on Failure option set on the same Communication Link.
- Click **Next** to specify time server information for WebSphere MQ communication link types or miscellaneous information for other communication links.

9 Miscellaneous Information for WebSphere MQ

- a** Specify the name of the exception message queue for the queue manager. The default is **TVISION.EXCEPTION.QUEUE**. If an event fails to be processed by the Analyzer, either because of an invalid event or an unexpected error or exception, the event is removed from the event queue and put on the exception message queue defined in the communication link.

Any fatal Java errors (such as **OutOfMemoryError**) cause the event to be put on the exception message queue and events are no longer pulled off the event queue until the project is restarted. If the exception message queue name is blank, the failed event are discarded.

You can use the **ManageQueue** utility to transfer messages from the specified exception message queue to enable the Analyzer to attempt to reprocess the failed events. See “ManageQueue” on page 401, for more information.

- b** Specify the number of event collection threads, from 1 to 100. The default value is 1. This value should match the results from the DBMS insert test. For more information about performing this test, see the *HP TransactionVision Advanced Customization Guide* PDF.

Note: Multiple threads performing concurrent database operations along with a DB2 behavior called Next Row Locking may lead to database deadlocks. To avoid deadlocks when using multiple event-processing threads, set the DB2 variable DB2_RR_TO_RS to ON with the following command:

```
db2set DB2_RR_TO_RS=ON
```

This setting is only effective after a DB2 restart. It affects all DB2 applications and cannot be used if other non-TransactionVision applications that require “Repeatable Read” semantics (Transaction Isolation Level RR) are using the same DB2 instance. If other applications using the same instance require “Repeatable Read” semantics, you must either create a separate DB2 instance for TransactionVision or set the number of event collection threads to 1.

- c** Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.
- d** Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see *HP TransactionVision Advanced Customization Guide* PDF. By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.

- e** Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.
- f** Specify whether the analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this communication link. Not storing the event XML reduces the space required for the project and may increase the analyzer performance. Note that the analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
- g** Specify whether the analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the analyzer performance.
- h** Specify whether to use a client or server connection type. The default connection type is server.

Note: For TIBCO EMS communication links and generic JMS communication links, the client connection concept does not apply.

If you select **client**, enter the following information:

Field	Description
Configuration Channel	Specify the channel to use when the Analyzer connects to the configuration queue manager.
Configuration Host	Specify the host running the configuration queue manager.
Configuration Port	Specify the listener port for the configuration queue manager.

Field	Description
Configuration CCSID	Specify the coded character set identifier (CCSID) to be used by the configuration queue manager. You are only required to set this value for queue managers using international double-byte character sets. To identify the correct CCSID for the queue manager, run the following command: echo display qmgr ccsid runmqsc <queue_manager> The default CCSID on a Japanese system is 932.
Event Channel	Specify the channel to use when the Analyzer connects to the event queue manager.
Event Host	Specify the host running the event queue manager.
Event Port	Specify the listener port for the event queue manager.
Event CCSID	Specify the coded character set identifier (CCSID) to be used by the event queue manager. This value matches the CCSID for the configuration queue manager.

- i** Click **Finish** to create the communication link or **<Back** to return to the Communication Link Maximum Message Length page.

Create a New Communication Link or Template for TIBCO EMS

This task describes how to create a new communication link or template for the TIBCO EMS type:

This task includes the following steps:

- “Prerequisites” on page 131
- “Assign Name for TIBCO EMS” on page 131
- “Sensor Connection for TIBCO EMS” on page 131
- “Analyzer Connection for TIBCO EMS” on page 132
- “Remote Queue Connection Choice for TIBCO EMS” on page 132
- “Remote Queue Connection for TIBCO EMS” on page 132

- “Maximum Message Length for TIBCO EMS” on page 133
- “Event Retry TIBCO EMS” on page 134
- “Time Server for TIBCO EMS” on page 134
- “Miscellaneous Information TIBCO EMS” on page 134

1 Prerequisites

A project has been selected. See “Projects” on page 43 for details.

2 Assign Name for TIBCO EMS

Enter the name of the new communication link and select the TIBCO EMS communication link type, then click **Next** to define the Sensor connection.

3 Sensor Connection for TIBCO EMS

If you selected TIBCO EMS as the communication link type, the Step 2: Sensor Connection page appears for TIBCO EMS. As you provide information on this page, the diagram is updated to provide a graphical representation of the Sensor connection for the new communication link.

- a** Specify the name of the TIBCO EMS server host.
- b** Specify the port number of the TIBCO EMS server.
- c** Specify the user name and password the Sensor is to use for JMS connections. This user should have administrative privileges.
- d** Specify the name of the queue that Sensors should check for configuration messages. Check **Use Default Configuration Queue** to use the default configuration queue name TVISION.CONFIGURATION.QUEUE. To use a different queue name, uncheck **Use Default Configuration Queue** and enter the queue name.
- e** Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
- f** Click **Next** to specify whether the Analyzer should connect to the same queue manager that the Sensors connect to on the Step 3: Analyzer Connection page.

4 Analyzer Connection for TIBCO EMS

For the TIBCO EMS communication link type, the Analyzer Connection Choice page appears. This page is similar to the WebSphere MQ Analyzer Connection Choice page in “Analyzer Connection for WebSphere MQ” on page 123.

If Yes	If No
Select Yes if the Analyzer should connect to the same TIBCO EMS Server that the sensor connects to. Click Next to specify the maximum message length in “Maximum Message Length for TIBCO EMS” on page 133.	Select No if the Analyzer should connect to a different TIBCO EMS Server that the sensor connects to. Click Next to specify whether to use one or two TIBCO EMS servers in “Remote Queue Connection Choice for TIBCO EMS” on page 132.

5 Remote Queue Connection Choice for TIBCO EMS

Selecting **No** on the Analyzer Connection Choice page, displays the following page:

If Yes	If No
Select Yes if the Analyzer should use the same TIBCO EMS servers for both configuration and event messages. See Same TIBCO EMS Server in “Remote Queue Connection for TIBCO EMS” on page 132.	Select No if the Analyzer should use different TIBCO EMS servers for configuration and event messages. See Separate TIBCO EMS Servers in “Remote Queue Connection for TIBCO EMS” on page 132.

Click **Next** to specify remote queue connections on the Remote Queue Connection page.

6 Remote Queue Connection for TIBCO EMS

Specify the same or separate TIBCO EMS servers, depending on your answer in “Remote Queue Connection Choice for TIBCO EMS” on page 132.

Same TIBCO EMS Server

If you answered **Yes** in Remote Queue Connection Choice for TIBCO EMS, specify a single TIBCO EMS server:

- a** Specify the TIBCO EMS server host name and port number.
- b** Specify the user name and password to make the connection.
- c** Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable here.
- d** Click **Next** to specify the maximum message length.

Separate TIBCO EMS Servers

If you answered **No** in Remote Queue Connection Choice for TIBCO EMS, specify separate TIBCO EMS servers.

- a** Specify the host name and port number for configuration queue's TIBCO EMS server.
- b** Specify the user name and password to make the connection to the configuration queue.
- c** Specify the host name and port number for the event queue's TIBCO EMS server.
- d** Specify the user name and password to make connection to the event queue.
- e** Because route queues must have the same name as their home queues, the configuration queue and event queue names are not editable here.
- f** Click **Next** to specify the maximum message length.

7 Maximum Message Length for TIBCO EMS

See "Maximum Message Length for WebSphere MQ" on page 125 for details on completing this page.

Click **Next** to specify reconnect retry options if the Sensor loses the connection on the Event Retry page.

8 Event Retry TIBCO EMS

See “Event Retry for WebSphere MQ” on page 126 for details on completing this page.

Click **Next** to calculate the time skew information.

9 Time Server for TIBCO EMS

For TIBCO EMS communication link types, the Analyzer and Sensors require a time server to calculate the time skew information. Check **Use the Time Server Running in the Analyzer** if you are running the time server within the Analyzer. Otherwise, specify the name of the host running the time server and its port number.

The time server reports the current time on the host where it is running. Any TransactionVision component can connect to it via TCP/IP and get the time skew from the current host to the host running the time server.

The Analyzer itself can act as a time server, or it can be run on a different host. To run the time server within the Analyzer, modify the `Analyzer.properties` file as follows:

- Set the **time_server** property to **on**
- Optionally, set **time_server_port** to the listening port you would like it to use. The default port is **21104**.

Use the **TimeServer.[sh|bat][PORT]** utility to run the time server stand-alone. The default port is **21104**.

Click **Next** to specify miscellaneous information for the communication link.

10 Miscellaneous Information TIBCO EMS

See “Miscellaneous Information for WebSphere MQ” on page 127 for details on completing this page.

Create a New Communication Link or Template for Progress SonicMQ

This task describes how to create a new communication link or template for the Progress SonicMQ type:

This task includes the following steps:

- “Prerequisites” on page 135
- “Assign Name for SonicMQ” on page 135
- “Sensor Connection for SonicMQ” on page 135
- “Analyzer Connection for SonicMQ” on page 136
- “Remote Queue Connection Choice for SonicMQ” on page 137
- “Remote Queue Connection for SonicMQ” on page 137
- “SonicMQ Domain Manager Information” on page 138
- “Maximum Message Length for SonicMQ” on page 139
- “Event Retry for SonicMQ” on page 139
- “Time Server for SonicMQ” on page 139
- “Miscellaneous Information for SonicMQ” on page 139

1 Prerequisites

A project has been selected. See “Projects” on page 43 for details.

2 Assign Name for SonicMQ

Enter the name of the new communication link and select the SonicMQ communication link type, then click **Next** to define the Sensor connection.

3 Sensor Connection for SonicMQ

If you selected SonicMQ as the communication link type, the Step 2: Sensor Connection page appears for SonicMQ. As you provide information on this page, the diagram is updated to provide a graphical representation of the Sensor connection for the new communication link.

Provide the following information:

- a Specify the host and port of the SonicMQ Broker listener that the Sensor connects to. Specify in the form of *hostname:port*.

Note: To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, **ssl://hostname:21112**.

- b Specify the name of the queue that Sensors should check for configuration messages. Check **Use Default Configuration Queue** to use the default configuration queue name **TVISION.CONFIGURATION.QUEUE**. To use a different queue name, uncheck **Use Default Configuration Queue** and enter the queue name.
- c Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
- d Click **Next** to specify whether the Analyzer should connect to the same queue manager to which the Sensors connect.

4 Analyzer Connection for SonicMQ

The Analyzer Connection Choice page is similar to the WebSphere MQ Analyzer Connection Choice page in “Analyzer Connection for WebSphere MQ” on page 123. Select Yes or No.

If Yes	If No
Select Yes if the Analyzer should connect to the same SonicMQ Broker that the Sensor connects to. Click Next to specify the maximum message length in “Maximum Message Length for SonicMQ” on page 139.	Select No if the Analyzer should connect to a different SonicMQ Broker than the Sensor connects to. Click Next to specify whether to use one or two SonicMQ Brokers for the Analyzer to connect to, in “Remote Queue Connection Choice for SonicMQ” on page 137.

5 Remote Queue Connection Choice for SonicMQ

Selecting **No** on the Analyzer Connection Choice page, displays the Remote Queue Connection Choice page:

If Yes	If No
Select Yes if the Analyzer should use the same SonicMQ Brokers for both configuration and event messages. See “Same SonicMQ Server” on page 137.	Select No if the Analyzer should use different SonicMQ Brokers for configuration and event messages. See “Separate SonicMQ Servers” on page 138.

Click **Next** to specify remote queue connections on the Remote Queue Connection page.

6 Remote Queue Connection for SonicMQ

Specify the same or separate SonicMQ servers, depending on your answer in “Remote Queue Connection Choice for SonicMQ” on page 137.

Same SonicMQ Server

If you answered **Yes** in Remote Queue Connection Choice for SonicMQ, specify a single SonicMQ server:

- a** Specify the host and port of the SonicMQ Broker listener that the Sensor connects to. Specify in the form of *hostname:port*.

Note: To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, **ssl://hostname:21112**.

- b** Specify the user name and password the Sensor is to use for the SonicMQ connection.
- c** Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable here.
- d** Click **Next** to specify the maximum message length.

Separate SonicMQ Servers

If you answered **No** in “Remote Queue Connection Choice for SonicMQ” on page 137, specify separate SonicMQ servers.

- a** Specify the host and port for the SonicMQ Broker with the configuration queue the Analyzer is to send configurations to.
- b** Specify the user name and password to make the connection to this SonicMQ Broker.
- c** Specify the SonicMQ configuration queue route the Analyzer is to use to send configuration messages to the Sensor.
- d** Specify the host and port for the SonicMQ Broker with the event queue the Analyzer is to receive events from.
- e** Specify the user name and password to make connection to this SonicMQ Broker.
- f** Specify the event queue the Analyzer is to receive events from the Sensor.
- g** Click **Next** to specify the maximum message length.

7 SonicMQ Domain Manager Information

Enter the following information.

- a** Specify the SonicMQ Domain Name to which the broker(s) belong. The SonicMQ default value is **Domain1**.
- b** Specify the host and port of the SonicMQ Domain Manager listener to which the Analyzer connects. This allows you to query for the event queue depth. Specify in the form of hostname:port.
- c** Specify the user name and password the Analyzer uses for the SonicMQ Domain Manager connection.
- d** Specify the name of the broker to which the Analyzer connects to collect events. This name is used to query the event queue depth through the Domain Manager.

8 Maximum Message Length for SonicMQ

See “Maximum Message Length for WebSphere MQ” on page 125 for details on completing this page.

Click **Next** to specify reconnect retry options if the Sensor loses the connection on the Event Retry page.

9 Event Retry for SonicMQ

See “Event Retry for WebSphere MQ” on page 126 for details on completing this page.

Click **Next** to calculate the time skew information.

10 Time Server for SonicMQ

For the SonicMQ communication link type, the Analyzer and Sensor require a time server to calculate the time skew information. Check **Use the Time Server Running in the Analyzer** if you are running the time server within the Analyzer. Otherwise, specify the name of the host running the time server and its port number. See “Time Server for TIBCO EMS” on page 134 for details on completing this page.

Click **Next** to specify miscellaneous information for the communication link.

11 Miscellaneous Information for SonicMQ

See “Miscellaneous Information for WebSphere MQ” on page 127 for details on completing this page.

Create a New Communication Link or Template for BEA WebLogic JMS

This task describes how to create a new communication link or template for the BEA WebLogic JMS type:

This task includes the following steps:

- “Prerequisites” on page 140
- “Assign Name for WebLogic JMS” on page 140
- “Sensor Connection for WebLogic JMS.” on page 140
- “Maximum Message Length for WebLogic JMS” on page 141
- “Event Retry for WebLogic JMS” on page 142
- “Miscellaneous Information for WebLogic JMS” on page 142

1 Prerequisites

A project has been selected. See “Projects” on page 43 for details.

2 Assign Name for WebLogic JMS

Enter the name of the new communication link and select the BEA WebLogic JMS communication link type, then click **Next** to define the Sensor connection.

3 Sensor Connection for WebLogic JMS.

If you selected WebLogic JMS as the communication link type, the Step 2: Sensor Connection page appears for WebLogic JMS. As you provide information on this page, the diagram is updated to provide a graphical representation of the Sensor connection for the new communication link.

Note: Since the Sensor and Analyzer need to use the same WebLogic JMS server, the connection information (host name, port number, queue connection factory, user name and password) must agree between the Sensor and the Analyzer.

- a** Specify the name of the WebLogic JMS server host.
- b** Specify the port number of the WebLogic JMS server.
- c** Specify the JNDI name of the queue connection factory.
- d** Specify the user name and password that the Sensor uses for JMS connections. This user should have administrative privileges.
- e** Specify the name of the queue that Sensors should check for configuration messages. Check **Use Default Configuration Queue** to use the default configuration queue name **TVISION.CONFIGURATION.QUEUE**. To use a different queue name, uncheck **Use Default Configuration Queue** and enter the queue name.
- f** Specify the name of the queue to which the Sensors should send event messages. This queue should only be used by TransactionVision.
- g** Click **Next** to specify maximum message length.

4 Maximum Message Length for WebLogic JMS

See “Maximum Message Length for WebSphere MQ” on page 125 for details on completing this page.

Click **Next** to specify reconnect retry options if the Sensor loses the connection on the Event Retry page.

5 Event Retry for WebLogic JMS

See “Event Retry for WebSphere MQ” on page 126 for details on completing this page.

Click **Next** to specify miscellaneous information for the communication link.

6 Miscellaneous Information for WebLogic JMS

See “Miscellaneous Information for WebSphere MQ” on page 127 for details on completing this page.

Create a New Communication Link or Template for Other JMS Provider

This task describes how to create a new communication link or template for the Other JMS Provider type:

This task includes the following steps:

- “Prerequisites” on page 142
- “Assign Name for Other JMS Provider” on page 142
- “Sensor Connection for Other JMS Provider” on page 143
- “Maximum Message Length for Other JMS Provider” on page 143
- “Event Retry for Other JMS Provider” on page 144
- “Miscellaneous Information for Other JMS Provider” on page 144

1 Prerequisites

A project has been selected. See “Projects” on page 43 for details.

2 Assign Name for Other JMS Provider

Enter the name of the new communication link and select the Other JMS Provider communication link type. Click **Next** to define the Sensor connection.

3 Sensor Connection for Other JMS Provider

For the Other JMS Provider communication link type, the Step 2: Sensor Connection page appears for Other JMS Provider. As you provide information on this page, the diagram is updated to provide a graphical representation of the Sensor connection for the new communication link. Specify the name of the WebLogic JMS server host.

- a** Specify the initial context factory of the JMS provider.
- b** Specify the URL that the Sensor uses to connect to the JNDI naming service.
- c** Specify the JNDI name of the Queue Connection Factory for connecting to the JMS server.
- d** Specify the user name and password that the Sensor uses for JMS connections. This user should have administrative privileges.
- e** Specify the name of the queue that Sensors should check for configuration messages. This queue must be a local queue on the specified queue manager. Check **Use Default Configuration Queue** to use the default configuration queue name **TVISION.CONFIGURATION.QUEUE**. To use a different queue name, uncheck **Use Default Configuration Queue** and enter the queue name.
- f** Specify the name of the queue to which the Sensors should send event messages. This queue should only be used by TransactionVision.
- g** Click **Next** to specify maximum message length.

4 Maximum Message Length for Other JMS Provider

See “Maximum Message Length for WebSphere MQ” on page 125 for details on completing this page.

Click **Next** to specify reconnect retry options if the Sensor loses the connection on the Event Retry page.

5 Event Retry for Other JMS Provider

See “Event Retry for WebSphere MQ” on page 126 for details on completing this page.

Click **Next** to specify miscellaneous information for the communication link.

6 Miscellaneous Information for Other JMS Provider

See “Miscellaneous Information for WebSphere MQ” on page 127 for details on completing this page.

Test a Communication Link

This task describes how to test a communication link:

- 1 Select the communication link to test.
- 2 Select the Analyzer to test the selected communication link with.
- 3 Click **Test**.

If the selected communication link template uses WebSphere MQ queues, the selected Analyzer sends a request message to the configuration queue requesting a confirmation on arrival (COA). The COA report message is sent to the event queue by the configuration queue that the communication link is configured to monitor. The end result is much like the behavior of typical communication between the Analyzer and Sensors. The communication link entry in the list of communication links is updated to show the Analyzer and test status.

If the selected communication link template uses JMS queues, the test only checks whether the selected Analyzer can successfully make a connection to the given configuration queue and event queue.

If the test fails, the error message is displayed in the test status column. The following are the most common reasons for a communication link test to fail:

- Timeout waiting for a reply message

This error is typically caused by channel problems with remote configurations. The request message may not make it to the final queue manager that is to report, or the report message may not make it to the queue manager to which the Communication Link Editor is connected. If the link test fails for this reason, see “Troubleshooting and Limitations” on page 212.

- WebSphere MQ errors

WebSphere MQ problems within the communication link are reported with the WebSphere MQ error reason, which usually provides enough information to isolate the problem.

Communication Links User Interface

This section describes:

- Communication Link Wizard on page 146
- Edit Communication Link Template on page 184
- Global Communication Link Templates Page on page 209
- Project Communication Links Page on page 211

 **Communication Link Wizard**

Description	<p>Enables you to begin the process of creating a project communication link or a communication link template by adding a name for the new communication link.</p> <p>To access: On the Admin > TransactionVision > Administration > Communication Link Template > Global Communication Link Template page, click New</p>
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating and configuring a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. The Assign Name step is applicable to all communication link types.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135 ➤ “Create a New Communication Link or Template for BEA WebLogic JMS” on page 140 ➤ “Create a New Communication Link or Template for Other JMS Provider” on page 142

Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Cancel	Returns you to the Global Communication Link Templates page without saving settings.
Next	Click to define the Sensor connection for the new communication link.
Enter Communication Link Name	Enter the name of the new communication link.
Select Communication Link Type	<p>Select the type of transport to use for communication between the Sensor and the Analyzer. Select one of the following:</p> <ul style="list-style-type: none"> ➤ IBM WebSphere MQ ➤ TIBCO EMS ➤ Progress SonicMQ ➤ BEA WebLogic JMS ➤ Other JMS Provider

 **Sensor Connection**

Description	Enables you to specify the Sensor connection for the selected communication link type.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. This Sensor Connection step is applicable to all communication link types.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135 ➤ “Create a New Communication Link or Template for BEA WebLogic JMS” on page 140 ➤ “Create a New Communication Link or Template for Other JMS Provider” on page 142

Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following buttons are included (unlabeled GUI elements are shown in angle brackets):

GUI Buttons	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next	<p>Click to proceed to the next page.</p> <ul style="list-style-type: none"> ➤ For IBM WebSphere MQ, TIBCO EMS, and SonicMQ, click to proceed to “Analyzer Connection” on page 155. ➤ For BEA WebLogic JMS and Other JMS Provider, click to proceed to “SonicMQ Domain Manager Information” on page 171.

IBM WebSphere MQ Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Queue manager sensor will connect to:	Specify the name of the queue manager Sensors connects to. This queue manager must be used by applications to be monitored.
Queue Sensor will be receiving configuration messages from:	Specify the name of the queue that Sensors should check for configuration messages. This queue must be a local queue on the specified queue manager.
Use Default Configuration Queue	Check to use the default configuration queue name TVISION.CONFIGURATION.QUEUE . To use a different queue name, uncheck Use Default Configuration Queue and enter the queue name.
Queue Sensor will be sending event messages to:	Specify the name of the queue to which Sensors should send event messages. This queue should not be used by applications other than TransactionVision.
Is event queue in a cluster?	Check if the queue is in a cluster. If you check this setting, the Transmission Queue setting is not displayed.
Use Local Queue Manager	Check if this queue is not in a cluster and is a local queue. If it is not a local queue, uncheck Use Local Queue Manager and specify either the name of the transmission queue that has the same name as the remote queue manager or a queue manager alias that resolves to the remote queue manager. This method of sending messages is described in the <i>WebSphere MQ Application Programming Guide</i> .

TIBCO EMS Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
TIBCO Server Host Name:	Specify the name of the TIBCO EMS server host.
TIBCO Server Port Number:	Specify the port number of the TIBCO EMS server.
JMS Connection User Name:	Specify the user name the Sensor uses for JMS connections. This user should have administrative privileges.
JMS Connection User Password:	Specify the password the Sensor uses for JMS connections.
Queue Sensor will be receiving configuration messages from:	Specify the name of the queue that Sensors should check for configuration messages.
Use Default Configuration Queue	Check to use the default configuration queue name TVISION.CONFIGURATION.QUEUE . To use a different queue name, uncheck Use Default Configuration Queue and enter the queue name.
Queue Sensor will be sending event messages to:	Specify the name of the queue to which Sensors should send event messages. This queue should not be used by applications other than TransactionVision.

Progress SonicMQ Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Broker URL Sensor will connect to (hostname:port):	Specify the host and port of the SonicMQ Broker listener to which the Sensor connects. Specify in the form of <i>hostname:port</i> . To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112 .
User Name:	Specify the user name the Sensor uses for the SonicMQ connections.
Password:	Specify the password the Sensor uses for the SonicMQ connections.
Queue Sensor will be receiving configuration messages from:	Specify the name of the queue that Sensors should check for configuration messages.
Use Default Configuration Queue	Check to use the default configuration queue name TVISION.CONFIGURATION.QUEUE . To use a different queue name, uncheck Use Default Configuration Queue and enter the queue name.
Queue Sensor will be sending event messages to:	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.

BEA WebLogic JMS Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

Note: Since the Sensor and Analyzer need to use the same WebLogic JMS server, the connection information (host name, port number, queue connection factory, user name and password) must agree between the Sensor and the Analyzer.

GUI Element	Description
WebLogic Server Host Name:	Specify the name of the WebLogic JMS server host.
WebLogic Server Port Number:	Specify the port number of the WebLogic JMS server.
Queue Connection Factory:	Specify the JNDI name of the queue connection factory.
JMS Connection User Name:	Specify the user name the Sensor uses for JMS connections. This user should have administrative privileges.
JMS Connection User Password:	Specify the user password the Sensor uses for JMS connections. This user should have administrative privileges.
Queue Sensor will be receiving configuration messages from:	Specify the name of the queue that Sensors should check for configuration messages.

GUI Element	Description
Use Default Configuration Queue	Check to use the default configuration queue name TVISION.CONFIGURATION.QUEUE . To use a different queue name, uncheck Use Default Configuration Queue and enter the queue name.
Queue Sensor will be sending event messages to:	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.

Other JMS Provider Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
JNDI Provider Context Factory:	Specify the initial context factory of the JMS provider.
JNDI Provider URL Sensor will connect to:	Specify the URL that the Sensor uses to connect to the JNDI naming service.
Queue Connection Factory:	Specify the JNDI name of the Queue Connection Factory for connecting to the JMS server.
JMS Connection User Name:	Specify the user name that the Sensor uses to make the connection.
JMS Connection User Password:	Specify the password that the Sensor uses to make the connection.
Queue Sensor will be receiving configuration messages from:	Specify the name of the queue that Sensors should check for configuration messages.

GUI Element	Description
Use Default Configuration Queue	Check to use the default configuration queue name TVISION.CONFIGURATION.QUEUE . To use a different queue name, uncheck Use Default Configuration Queue and enter the queue name.
Queue Sensor will be sending event messages to:	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.

Analyzer Connection

Description	Enables you to choose whether the Analyzer connects to the same or different queue manager than the communication type applications connect to.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. This Analyzer Connection step is applicable only to IBM WebSphere MQ, TIBCO EMS, and Progress SonicMQ.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135

Wizard Map	The Communication Link wizard contains: Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information >
Useful Links	<ul style="list-style-type: none">➤ “Understanding Communication Links” on page 106➤ “TransactionVision Queue Requirements” on page 110➤ “Communication Link Templates” on page 117

The following buttons are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next >	For IBM WebSphere MQ, TIBCO EMS, and SonicMQ only, click to proceed to “Remote Queue Connection Question” on page 158.

IBM WebSphere MQ Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
No	Select No if the Analyzer should connect to a different queue manager than WebSphere MQ applications connect to.
Yes	Select Yes if the Analyzer should connect to the same queue manager that WebSphere MQ applications connect to.

TIBCO EMS Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
No	Select No if the Analyzer should connect to a different TIBCO EMS Server than the sensor connects to.
Yes	Select Yes if the Analyzer should connect to the same TIBCO EMS Server that the sensor connects to.

Progress SonicMQ Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
No	Select No if the Analyzer should connect to a different SonicMQ Broker than the Sensor connects to.
Yes	Select Yes if the Analyzer should connect to the same SonicMQ Broker that the Sensor connects to.



Remote Queue Connection Question

Description	<p>Enables you to choose whether to use a single remote connection or separate remote connections for both configuration and event messages.</p> <p>This page displays if you selected No on “Analyzer Connection” on page 155.</p>
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. This Remote Queue Connection Question step is applicable only to IBM WebSphere MQ, TIBCO EMS, and Progress SonicMQ.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135

Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following buttons are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next >	<p>Click to proceed to the next page.</p> <ul style="list-style-type: none"> ➤ If you select Yes, click to proceed to the “Remote Queue Connection (for a Single Remote Connection)” on page 161. ➤ If you select No, click to proceed to the “Remote Queue Connection (for Separate Remote Connections)” on page 166

IBM WebSphere MQ Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
No	Select No if the Analyzer should use different queue managers for configuration and event messages. See “Separate IBM WebSphere MQ Queue Managers” on page 168.
Yes	Select Yes if the Analyzer should use the same queue manager for both configuration and event messages. See “Same IBM WebSphere MQ Server” on page 163.

TIBCO EMS Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
No	Select No if the Analyzer should use different TIBCO EMS servers for configuration and event messages. See “Separate TIBCO EMS Servers” on page 133.
Yes	Select Yes if the Analyzer should use the same TIBCO EMS servers for both configuration and event messages. See “Same TIBCO EMS Server” on page 133.

Progress SonicMQ Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
No	Select No if the Analyzer should use different SonicMQ Brokers for configuration and event messages. See “Separate Progress SonicMQ Brokers” on page 170.
Yes	Select Yes if the Analyzer should use the same SonicMQ Brokers for both configuration and event messages. See “Same Progress SonicMQ Broker” on page 165.



Remote Queue Connection (for a Single Remote Connection)

Description	<p>Enables you to specify a single remote connection for both configuration and event messages.</p> <p>This page opens if you selected Yes in “Remote Queue Connection Question” on page 158.</p>
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. This Remote Queue Connection (for a Single Remote Connection) step is applicable only to IBM WebSphere MQ, TIBCO EMS, and Progress SonicMQ.

Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135
Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following buttons are included (unlabeled GUI elements are shown in angle brackets):

GUI Buttons	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next	Click to proceed to “SonicMQ Domain Manager Information” on page 171.

Same IBM WebSphere MQ Server

If you answered **Yes** in “Remote Queue Connection Question” on page 158, specify a single remote queue manager.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Queue Manager Analyzer will connect to:	Specify the name of the queue manager the Analyzer uses for both configuration and event messages.
Use Transmission Queue or Queue Manager Alias	Click to use a transmission queue or queue manager alias.
Queue Analyzer will be sending configuration messages to:	Specify the name of the queue that the Analyzer should send configuration messages to. This queue must be a remote queue for the configuration queue on the Sensor queue manager.
Queue Analyzer will be receiving events from:	Enter the name of the transmission queue that has the same name as the remote queue manager or a queue manager alias that resolves to the remote queue manager.

Same TIBCO EMS Server

If you answered **Yes** in “Remote Queue Connection Question” on page 158, specify a single TIBCO EMS server.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
TIBCO EMS Server Host Name:	Specify the TIBCO EMS server host name.
TIBCO EMS Server Port Number:	Specify the TIBCO EMS server host port number. The default is 7222.

GUI Element	Description
User Name:	Specify the user name to make the connection.
Password:	Specify the password to make the connection.
Queue Analyzer will be sending configuration message to:	This configuration queue field is not editable since route queues must have the same name as their home queue.
Queue Analyzer will be receiving events from:	This event queue field is not editable since route queues must have the same name as their home queue.

Same Progress SonicMQ Broker

If you answered **Yes** in “Remote Queue Connection Question” on page 158, specify a single Progress SonicMQ server.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Broker URL Analyzer will connect to (hostname:port):	<p>Specify the host and port of the SonicMQ Broker listener that the Sensor connects to. Specify in the form of <i>hostname:port</i>.</p> <p>The host and port you specified in “Sensor Connection” on page 148 is automatically entered.</p> <p>To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112.</p>
User Name:	Specify the user name the Sensor uses for the SonicMQ connection.
Password:	Specify the password the Sensor uses for the SonicMQ connection.
Queue Analyzer will be sending configuration messages to:	Specify the name of the queue that the Analyzer should send configuration messages to. This queue must be a remote queue for the configuration queue on the Sensor queue manager.
Queue Analyzer will be receiving events from:	Enter the name of the transmission queue that has the same name as the remote queue manager or a queue manager alias that resolves to the remote broker.



Remote Queue Connection (for Separate Remote Connections)

Description	<p>Enables you to specify separate remote connections for both configuration and event messages.</p> <p>This page opens if you selected No in “Remote Queue Connection Question” on page 158.</p>
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. This Remote Queue Connection (for Separate Remote Connections) step is applicable only to IBM WebSphere MQ, TIBCO EMS, and Progress SonicMQ.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135

Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following buttons are included (unlabeled GUI elements are shown in angle brackets):

GUI Buttons	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next	Click to proceed to “SonicMQ Domain Manager Information” on page 171.

Separate IBM WebSphere MQ Queue Managers

If you answered **No** in “Remote Queue Connection Question” on page 158, specify separate remote queue managers.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Queue Manager Analyzer will connect to for sending configuration messages:	Specify the name of the queue manager the Analyzer uses for sending configuration messages to Sensors.
Queue Manager Analyzer will be sending configuration messages to:	Specify the name of the queue that the Analyzer should send configuration messages to. This queue must be a remote queue for the configuration queue on the Sensor queue manager
Queue Manager Analyzer will connect to for receiving event messages:	Specify the name of the queue manager the Analyzer monitors for event messages from Sensors.
Queue Analyzer will be receiving event messages on:	Specify the name of the queue on the event queue manager that the Analyzer should monitor for event messages. This queue must be a local queue and should not be used by applications other than TransactionVision.

Separate TIBCO EMS Servers

If you answered **No** in “Remote Queue Connection Question” on page 158, specify separate TIBCO EMS servers.

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Configuration Queue TIBCO Server Host Name:	Specify the host name for configuration queue's TIBCO EMS server.
Configuration Queue TIBCO Server Port Number:	Specify the port number for configuration queue's TIBCO EMS server.
Configuration Queue User Name:	Specify the user name to make the connection to the configuration queue
Configuration Queue User password.	Specify the password to make the connection to the configuration queue
Configuration Queue Queue Analyzer will be sending configuration messages to:	This configuration queue field is not editable since route queues must have the same name as their home queues.
Event Queue TIBCO Server Host Name:	Specify the host name and port number for the event queue's TIBCO EMS server.
Event Queue TIBCO Server Port Number:	Specify the host name and port number for the event queue's TIBCO EMS server.
Event Queue User Name:	Specify the user name and password to make connection to the event queue.

GUI Element	Description
Event Queue User Password:	Specify the user name and password to make connection to the event queue.
Event Queue Analyzer will be receiving event messages on:	This event queue field is not editable since route queues must have the same name as their home queues.

Separate Progress SonicMQ Brokers

If you answered **No** in “Remote Queue Connection Question” on page 158, specify separate Progress SonicMQ brokers.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Broker URL Analyzer will use for sending configuration messages:	Specify the host and port for the SonicMQ Broker with the configuration queue the Analyzer sends configurations to. To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112 .
User Name:	Specify the user name to make the connection to this SonicMQ Broker.
Password:	Specify the password to make the connection to this SonicMQ Broker.
Queue Analyzer will be sending configuration messages to:	Specify the SonicMQ configuration queue route the Analyzer uses to send configuration messages to the Sensor.

GUI Element	Description
Broker URL Analyzer will use for receiving event messages:	Specify the host and port for the SonicMQ Broker with the event queue the Analyzer receives events from. To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112 .
User Name:	Specify the user name to make connection to this SonicMQ Broker.
Password:	Specify the password to make connection to this SonicMQ Broker.
Queue Analyzer will be receiving event messages on:	Specify the event queue the Analyzer receives events from the Sensor.



SonicMQ Domain Manager Information

Description	Optional information that enables you to define the SonicMQ domain and domain manager connection. When provided, this information is used by the Analyzer to obtain the event queue depth to be displayed on the Projects page.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ This step applies only to Progress SonicMQ.
Included in Task	“Create a New Communication Link or Template for Progress SonicMQ” on page 135

Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Buttons	Description
Domain Name	Specify the SonicMQ Domain Name to which the broker(s) belong. SonicMQ default value is Domain .
Domain Manager URL Analyzer will connect to (hostname:port)	Specify the host and port of the SonicMQ Domain Manager listener to which the Analyzer connects. This allows you to query for the event queue depth. Specify in the form of <i>hostname:port</i> .
User Name	Specify the user name the Analyzer uses for the SonicMQ Domain Manager connection.
Password	Specify the password the Analyzer uses for the SonicMQ Domain Manager connection.
Name of Broker with event queue that the Analyzer will be connecting to	Specify the name of the broker to which the Analyzer connects to collect events. This name is used to query the event queue depth through the Domain Manager.

 **Specify Maximum Message Length**

Description	Enables you to specify the maximum message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ This step applies to all communication link types.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135 ➤ “Create a New Communication Link or Template for BEA WebLogic JMS” on page 140 ➤ “Create a New Communication Link or Template for Other JMS Provider” on page 142

Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Buttons	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next	Click to proceed to “Event Retry Options” on page 175.
Event Queue Maximum Message Length (10000 - 2147483647)	<p>Enter the maximum message length.</p> <p>It is recommended to use the default maximum message length of 4194304 bytes. A minimum message length of 10,000 bytes is required.</p> <p>To restrict the message length based on your applications, consider the length of messages sent by your application and allow an additional 4000 bytes for event information.</p>

 **Event Retry Options**

Description	Enables you to specify the number of retries.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ This step applies to all communication link types.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135 ➤ “Create a New Communication Link or Template for BEA WebLogic JMS” on page 140 ➤ “Create a New Communication Link or Template for Other JMS Provider” on page 142
Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Elements	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next	Click to proceed to the next page. ► For TIBCO EMS and Progress SonicMQ, click to proceed to “Time Server” on page 178. ► For BEA WebSphere MQ, IBM WebLogic JMS, and Other JMS Provider, click to proceed to “Miscellaneous Information” on page 180.
Retry on Failure	Check to enable Sensors to retry if they fail to put an event on the event queue. The Event Delivery Retry on Failure elements are enabled. Uncheck to disable the Event Delivery Retry on Failure elements. NOTE: The default setting is not to retry event package delivery on failure. Set a value appropriate for your event collecting strategy. In a production environment, uncheck Retry on Failure for the safest setting.

Event Connection Retry on Failure (does not apply to WebSphere MQ Sensors)

Note: The default setting is to retry forever every 10 seconds.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Elements	Description
Retry forever	Click to continue to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.
Retry once per TransactionVision event	Click to retry once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across TransactionVision events	Click to continue to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision even	Click to continue to retry at the specified interval until the timeout is reached.

Event Delivery Retry on Failure

If you selected **Retry on Failure**, set the following elements to control retry attempts (unlabeled GUI elements are shown in angle brackets):

GUI Elements	Description
Retry Timeout:	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.

 **Time Server**

Description	Enables you to specify the time server to calculate the time skew information for the Analyzer and Sensors.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ The Communication Link Wizard steps vary depending on the communication link type and your choices within some steps. This Time Server step is applicable only to TIBCO EMS and Progress SonicMQ. ➤ The time server reports the current time on the host where it is running. Any TransactionVision component can connect to it through TCP/IP and get the time skew from the current host to the host running the time server. ➤ The Analyzer itself can act as a time server, or it can be run on a different host. To run the time server within the Analyzer, modify the Analyzer.properties file as follows: <ul style="list-style-type: none"> ➤ Set the time_server property to on. ➤ Optionally, set time_server_port to the listening port you would like it to use. The default port is 21104. ➤ Use the TimeServer.[sh bat][PORT] utility to run the time server stand-alone. The default port is 21104.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135

Wizard Map	The Communication Link wizard contains: Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information >
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Elements	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Next	Click to proceed to “Miscellaneous Information” on page 180.
Time Server Host:	Enter the host name if Use the Time Server Running the Analyzer is unchecked.
Time Server Port:	Enter the server port if Use the Time Server Running the Analyzer is unchecked.
Use the Time Server Running in the Analyzer	Check if you are running the time server within the Analyzer. Otherwise, specify the name of the host running the time server and its port number.

 **Miscellaneous Information**

Description	Enables you to specify miscellaneous information.
Important Information	<ul style="list-style-type: none"> ➤ The Communication Link Wizard guides you through the process of creating a project communication link or a communication link template. ➤ The Communication Link Wizard does not check whether a queue exists when creating communication links. However, you must create all queues before TransactionVision can use the communication link. These queues include event queues, configuration queues, and exception message queues. See “TransactionVision Queue Requirements” on page 110 for details. ➤ This step applies to all communication link types.
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Communication Link or Template for WebSphere MQ” on page 121 ➤ “Create a New Communication Link or Template for TIBCO EMS” on page 130 ➤ “Create a New Communication Link or Template for Progress SonicMQ” on page 135 ➤ “Create a New Communication Link or Template for BEA WebLogic JMS” on page 140 ➤ “Create a New Communication Link or Template for Other JMS Provider” on page 142
Wizard Map	<p>The Communication Link wizard contains:</p> <p>Communication Link Wizard > Sensor Connection > Analyzer Connection > Remote Queue Connection Question > Remote Queue Connection (for a Single Remote Connection) > Remote Queue Connection (for Separate Remote Connections) > SonicMQ Domain Manager Information > Specify Maximum Message Length > Event Retry Options > Time Server > Miscellaneous Information ></p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Elements	Description
Back	Click to return to the previous page.
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Finish	Click to create the communication link.
Exception message queue name for <communication link queue manager>	<p>Specify the name of the exception message queue. The default is TVISION.EXCEPTION.QUEUE.</p> <p>If an event fails to be processed by the Analyzer, either because of an invalid event or an unexpected error or exception, the event is removed from the event queue and put on the exception message queue defined in the communication link.</p> <p>Any fatal Java errors (such as OutOfMemoryError) causes the event to be put on the exception message queue and events are no longer pulled off the event queue until the project is restarted. If the exception message queue name is blank, the failed event are discarded.</p> <p>You can use the ManageQueue utility to transfer messages from the specified exception message queue to enable the Analyzer to attempt to reprocess the failed events. See “ManageQueue” on page 401 for more information.</p>

GUI Elements	Description
Number of event collection threads (1-100):	<p>Specify the number of event collection threads, from 1 to 100. The default value is 1. This value should match the results from the DBMS insert test.</p> <p>NOTE: Multiple threads performing concurrent database operations along with a DB2 behavior called Next Row Locking may lead to database deadlocks. To avoid deadlocks when using multiple event-processing threads, set the DB2 variable DB2_RR_TO_RS to ON with the following command:</p> <pre>db2set DB2_RR_TO_RS=ON</pre> <p>This setting is only effective after a DB2 restart. It affects all DB2 applications and cannot be used if other non-TransactionVision applications that require “Repeatable Read” semantics (Transaction Isolation Level RR) are using the same DB2 instance. If other applications using the same instance require “Repeatable Read” semantics, you must either create a separate DB2 instance for TransactionVision or set the number of event collection threads to 1.</p>
Sensor trace logging:	Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.
Event persistence:	Specify PERSISTENT delivery mode to ensure that an event is not lost in transit in case of a JMS provider failure. An event sent with this delivery mode is logged to a stable storage. Event persistence is off by default.
Direct user event processing:	<p>Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.</p> <p>By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.</p>

GUI Elements	Description
Store User Data:	Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.
Store Event Document:	Specify whether the Analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this communication link. Not storing the event XML reduces the space required for the project and may increase the Analyzer performance. Note that the Analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
Store Raw Events:	Specify whether the Analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the Analyzer performance.

Edit Communication Link Template

Description	<p>Enables you to edit the communication link type's settings.</p> <p>To access:</p> <ul style="list-style-type: none"> ➤ Admin > TransactionVision > Administration > Communication Link Template > click Edit ➤ Admin > TransactionVision > Current Project> Communication Links > click Edit
Important Information	<p>Depending on your changes, the following message may be displayed:</p> <p>This significant change to the communication link could leave unprocessed events intended for the previous communication link definition on the event queue. It is recommended to only make this change while the event queue is empty. Proceed with this change?</p> <p>[Yes] [No]</p> <p>View the Project Status to ensure that the event queue is empty, then click Yes to confirm your changes.</p> <p>Note that if the project event packaging (specified in the data collection filter) is greater than 1, the Sensor may have events buffered but not yet sent to the event queue. If you are packaging multiple events, stop the project to cause the Sensor to send any buffered events to the event queue. Then restart the project to process the events on the event queue.</p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “TransactionVision Queue Requirements” on page 110 ➤ “Communication Link Templates” on page 117

The following buttons are included (unlabeled GUI elements are shown in angle brackets>):

GUI Buttons	Description
Cancel	Click to return to the Global Communication Link Templates page without saving settings.
Finish	Click to save your updated settings.

The following sections describe the communication link type edit settings:

- “IBM WebSphere MQ Edit Settings” on page 185
- “TIBCO EMS Edit Settings” on page 190
- “Progress SonicMQ Edit Settings” on page 195
- “BEA WebLogic JMS Edit Settings” on page 201
- “Other JMS Provider Edit Settings” on page 205

IBM WebSphere MQ Edit Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Elements	Description
Communication Link Name	Enter the name of the communication link.
Sensor Connection	
Configuration Queue Manager	Specify the name of the queue manager Sensors connect to. This queue manager must be used by applications to be monitored.
Configuration Queue	Specify the name of the queue that Sensors should check for configuration messages. This queue must be a local queue on the specified queue manager. The default configuration queue name is TVISION.CONFIGURATION.QUEUE.

GUI Elements	Description
Event Transmission Queue	This field only appears if the event queue is not in a cluster. Leave this field blank to indicate that the event queue is on the same queue manager as the configuration queue. If not, specify either the name of the transmission queue that has the same name as the remote queue manager or the queue manager alias that resolves to the remote queue manager. This method of sending messages is described in the <i>WebSphere MQ Application Programming Guide</i> (SC33-0807-08).
Event Queue	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
Cluster Queue	Check if the event queue is in a cluster. When this field is checked, the Event Transmission Queue field is not displayed.
Analyzer Connection	
Configuration Queue Manager	Specify the name of the queue manager the Analyzer uses for sending configuration messages to Sensors.
Configuration Transmission Queue	Leave this field blank if the configuration queue is on the same queue manager as the event queue. If not, specify either the name of the transmission queue that has the same name as the remote queue manager or the queue manager alias that resolves to the remote queue manager.
Configuration Queue	Specify the name of the queue that the Analyzer should send configuration messages to.
Event Queue Manager	Specify the name of the queue manager the Analyzer is to monitor for event messages from Sensors.
Event Queue	Specify the name of the queue on the event queue manager that the Analyzer should monitor for event messages. This queue must be a local queue and should not be used by applications other than TransactionVision.

GUI Elements	Description
Event Retry on Failure	
Retry on Failure	Check to retry if the event connection or delivery fails. Then specify event connection and delivery failure retry options. Otherwise, events are lost if the event connection or delivery fails.
Event Connection Retry on Failure (does not apply to WebSphere MQ Sensors)	
Retry forever	Continues to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.
Retry once per TransactionVision event	Retries once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across transaction events	Continues to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision event	Continues to retry at the specified interval until the timeout is reached.
Event Delivery Retry	
Retry Timeout	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.

GUI Elements	Description
Miscellaneous Information	
Event Queue Maximum Message Length	<p>Specify the maximum event message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.</p> <p>It is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow and additional 4000 bytes for event information.</p>
Exception Message Queue Name	<p>Specify the name of the exception message queue for the queue manager where the Analyzer place configuration messages. The default is TVISION.EXCEPTION.QUEUE. If left blank, the Analyzer discards events that fail to be placed in the database. For more information about the exception message queue, see “Miscellaneous Information” on page 180.</p>
Number of event collection threads (1-100)	<p>Specify the number of event collection threads, from 1 to 100. The default value is 1. For more information about event collection threads, see “Miscellaneous Information” on page 180.</p>
Sensor trace logging:	<p>Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.</p>
Direct user event processing	<p>Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF. By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.</p>

GUI Elements	Description
Store User Data	Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.
Store Event Document	Specify whether the analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this commlink. Not storing the event XML reduces the space required for the project and may increase the analyzer performance. Note that the analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
Store Raw Events	Specify whether the analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the analyzer performance.

GUI Elements	Description
Connection Type	<p>Specify whether to use a client or server connection type. If you choose a client connection, you must also specify the following information:</p> <ul style="list-style-type: none">➤ The channel to use when the Analyzer connects to the configuration queue manager➤ The host running the configuration queue manager➤ The listener port for the configuration queue manager➤ The coded character set identifier (CCSID) to be used by the configuration queue manager. You are only required to set this value for queue managers using international double-byte character sets (for example, Shift-JIS). The default CCSID on a Japanese system is 932. To identify the correct CCSID for the queue manager, run the following command: <code>echo display qmgr ccsid runmqsc <queue_manager></code>➤ The channel to use when the Analyzer connects to the event queue manager➤ The host running the event queue manager➤ The listener port for the event queue manager➤ The coded character set identifier (CCSID) to be used by the event queue manager. This value matches the CCSID for the configuration queue manager.

TIBCO EMS Edit Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Communication link name	Enter the name of the communication link.
Sensor Connection	
TIBCO Server Host Name	Specify the name of the TIBCO EMS server host.

GUI Element	Description
TIBCO Server Port number	Specify the port number of the TIBCO EMS server host.
User Name	Specify the user name the Sensor uses for JMS connections. This user should have administrative privileges.
Password	Specify the password for the specified user.
Config Queue Name	Specify the name of the queue that Sensors should check for configuration messages. The default configuration queue name is TVISION.CONFIGURATION.QUEUE .
Event Queue Name	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
Analyzer Connection	
Configuration Queue TIBCO Server Host Name	Specify the name of the TIBCO EMS server host the Analyzer uses for configuration messages.
Configuration Queue TIBCO Server Port number	Specify the port number of the TIBCO EMS server the Analyzer uses for configuration messages.
Configuration Queue User Name	Specify the user name the Analyzer uses for JMS connections. This user should have administrative privileges.
Configuration Queue Password	Specify the password for the specified user name.
Configuration Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.
Event Queue TIBCO Server Host Name	Specify the name of the TIBCO EMS server host the Analyzer uses for event messages.
Event Queue TIBCO Server Port Number	Specify the port number of the TIBCO EMS server the Analyzer uses for event messages.

GUI Element	Description
Event Queue User Name	Specify the user name the Analyzer uses for JMS connections. This user should have administrative privileges.
Event Queue Password	Specify the password for the specified user.
Event Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.
Event Retry on Failure	
Retry on Failure	Check to retry if the event connection or delivery fails. Then specify event connection and delivery failure retry options. Otherwise, events are lost if the event connection or delivery fails.
Event Connection Retry on Failure	
Retry forever	Continues to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.
Retry once per TransactionVision event	Retries once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across transaction events	Continues to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision event	Continues to retry at the specified interval until the timeout is reached.

GUI Element	Description
Event Delivery Retry on Failure	
Retry Timeout	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.
Time Server Definition	
Use the Time Server running in the Analyzer	Check to use the time skew server.
Host name	Specify the host name of the time skew server.
Host port	Specify the port number of the time skew server.
Miscellaneous Information	
Event Queue Maximum Message Length	<p>Specify the maximum event message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.</p> <p>It is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow and additional 4000 bytes for event information.</p>
Exception Message Queue Name	Specify the name of the exception message queue for the queue manager where the Analyzer places configuration messages. The default is TVISION.EXCEPTION.QUEUE . If left blank, the Analyzer discards events that fail to be placed in the database. For more information about the exception message queue, see “Miscellaneous Information” on page 180.

GUI Element	Description
Number of Event Collection Threads	Specify the number of event collection threads, from 1 to 100. The default value is 1. For more information about event collection threads, see “Miscellaneous Information” on page 180.
Sensor Trace Logging	Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.
Event persistence:	Specify PERSISTENT delivery mode to ensure that an event is not lost in transit in case of a JMS provider failure. An event sent with this delivery mode is logged to a stable storage. Event persistence is off by default.
Direct user event processing	Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF. By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.
Store User Data	Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.

GUI Element	Description
Store Event Document	Specify whether the analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this commlink. Not storing the event XML reduces the space required for the project and may increase the analyzer performance. Note that the analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
Store Raw Events	Specify whether the analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the analyzer performance.

Progress SonicMQ Edit Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Communication Link Name	Enter the name of the communication link.
Sensor Connection	
Broker URL	Specify the host and port of the SonicMQ Broker. To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112 .
User Name	Specify the user name the Sensor uses for JMS connections. This user should have administrative privileges.
Password	Specify the password for the specified user.
Configuration Queue Name	Specify the name of the queue that Sensors should check for configuration messages. The default configuration queue name is TVISION.CONFIGURATION.QUEUE .

GUI Element	Description
Event Queue Name	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
Analyzer Connection	
Configuration Queue Broker URL	Specify the name of the SonicMQ Broker host and port the Analyzer uses for configuration messages. To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112 .
Configuration Queue User Name	Specify the user name the Analyzer uses for JMS connections. This user should have administrative privileges.
Configuration Queue Password	Specify the password for the specified user.
Configuration Queue Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.
Event Queue SonicMQ Broker URL	Specify the name of the SonicMQ Broker host and port the Analyzer uses for event messages. To configure the TransactionVision SonicMQ communication link to use SSL, you need to specify the SSL protocol when setting the broker host. For example, ssl://hostname:21112 .
Event Queue User Name	Specify the user name the Analyzer uses for JMS connections. This user should have administrative privileges.
Event Queue Password	Specify the password for the specified user.
Event Queue Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.

GUI Element	Description
Domain Manager Information Domain Name	Specify the SonicMQ Domain Name to which the broker(s) belong. SonicMQ default value is Domain1 .
Domain Manager Info Broker URL	Specify the host and port of the SonicMQ Domain Manager listener to which the Analyzer connects. This allows you to query for the event queue depth. Specify in the form of <i>hostname:port</i> .
Domain Manager Information User Name	Specify the user name the Analyzer uses for the SonicMQ Domain Manager connection.
Domain Manager Information Password	Specify the password the Analyzer uses for the SonicMQ Domain Manager connection.
Domain Manager Name of Broker with event queue that the Analyzer will be connecting to	Specify the name of the broker to which the Analyzer connects to collect events. This name is used to query the event queue depth through the Domain Manager.
Event Retry on Failure	
Retry on Failure	Check to retry if the event connection or delivery fails. Then specify event connection and delivery failure retry options. Otherwise, events are lost if the event connection or delivery fails.
Event Connection Retry on Failure	
Retry forever	Continues to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.

GUI Element	Description
Retry once per TransactionVision event	Retries once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across transaction events	Continues to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision event	Continues to retry at the specified interval until the timeout is reached.
Event Delivery Retry	
Retry Timeout	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.
Time Server Definition	
Host name	Specify the host name of the time skew server.
Host port	Specify the port number of the time skew server.

GUI Element	Description
Miscellaneous Information	
Event Queue Maximum Message Length	<p>Specify the maximum event message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.</p> <p>It is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow an additional 4000 bytes for event information.</p>
Exception Message Queue Name	<p>Specify the name of the exception message queue for the queue manager where the Analyzer places configuration messages. The default is TVISION.EXCEPTION.QUEUE. If left blank, the Analyzer discards events that fail to be placed in the database. For more information about the exception message queue, see “Miscellaneous Information” on page 180.</p>
Number of Event Collection Threads	<p>Specify the number of event collection threads, from 1 to 100. The default value is 1. For more information about event collection threads, see “Miscellaneous Information” on page 180.</p>
Sensor Trace Logging	<p>Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.</p>
Event persistence:	<p>Specify PERSISTENT delivery mode to ensure that an event is not lost in transit in case of a JMS provider failure. An event sent with this delivery mode is logged to a stable storage. Event persistence is off by default.</p>

GUI Element	Description
Direct user event processing	Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF. By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.
Store User Data	Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.
Store Event Document	Specify whether the analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this commlink. Not storing the event XML reduces the space required for the project and may increase the analyzer performance, Note that the analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
Store Raw Events	Specify whether the analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the analyzer performance.

BEA WebLogic JMS Edit Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Communication Link Name	Enter the name of the communication link.
Sensor Connection	
WebLogic Server Host Name	Specify the host name of your WebLogic JMS server.
WebLogic Server Port Number	Specify the port number of your WebLogic JMS server.
Queue Connection Factory	Specify the Queue connection factory.
Config Queue Name	Specify the name of the queue that Sensors should check for configuration messages. The default configuration queue name is TVISION.CONFIGURATION.QUEUE .
Event Queue Name	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
Analyzer Connection	
Configuration Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.
Event Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.
Event Retry on Failure	
Retry on Failure	Check to retry if the event connection or delivery fails. Then specify event connection and delivery failure retry options. Otherwise, events are lost if the event connection or delivery fails.

GUI Element	Description
Event Connection Retry on Failure	
Retry forever	Continues to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.
Retry once per TransactionVision event	Retries once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across transaction events	Continues to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision event	Continues to retry at the specified interval until the timeout is reached.
Event Delivery Retry	
Retry Timeout	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.

GUI Element	Description
Miscellaneous Information	
Event Queue Maximum Message Length	<p>Specify the maximum event message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.</p> <p>It is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow and additional 4000 bytes for event information.</p>
Exception Message Queue Name	<p>Specify the name of the exception message queue for the queue manager where the Analyzer places configuration messages. The default is TVISION.EXCEPTION.QUEUE. If left blank, the Analyzer discards events that fail to be placed in the database. For more information about the exception message queue, see “Miscellaneous Information” on page 180.</p>
Number of Event Collection Threads	<p>Specify the number of event collection threads, from 1 to 100. The default value is 1. For more information about event collection threads, see “Miscellaneous Information” on page 180.</p>
Sensor Trace Logging	<p>Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.</p>
Event persistence:	<p>Specify PERSISTENT delivery mode to ensure that an event is not lost in transit in case of a JMS provider failure. An event sent with this delivery mode is logged to a stable storage. Event persistence is off by default.</p>

GUI Element	Description
Direct user event processing	Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF. By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.
Store User Data	Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.
Store Event Document	Specify whether the analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this commlink. Not storing the event XML reduces the space required for the project and may increase the analyzer performance. Note that the analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
Store Raw Events	Specify whether the analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the analyzer performance.

Other JMS Provider Edit Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Buttons	Description
Communication Link Name	Enter the name of the communication link.
Sensor Connection	
Connection Provider Context Factory	Specify the initial context factory of the JMS provider.
Connection URL	Specify the URL that the Sensor uses to connect to the JNDI naming service.
Queue Connection Factory	Specify the JNDI name of the Queue Connection Factory for connecting to the JMS server.
User Name	Specify the user name the Sensor uses to make the connection.
Password	Specify the password the Sensor uses to make the connection.
Config Queue Name	Specify the name of the queue that Sensors should check for configuration messages. The default configuration queue name is TVISION.CONFIGURATION.QUEUE .
Event Queue Name	Specify the name of the queue that Sensors should send event messages to. This queue should not be used by applications other than TransactionVision.
Analyzer Connection	
Configuration Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.
Event Queue Name	Because route queues must have the same name as their home queue, the configuration queue and event queue are not editable.

GUI Buttons	Description
Event Retry on Failure	
Retry on Failure	Check to retry if the event connection or delivery fails. Then specify event connection and delivery failure retry options. Otherwise, events are lost if the event connection or delivery fails.
Event Connection Retry on Failure	
Retry forever	Continues to retry at the specified interval as long as event collection is in progress. Note that this option may cause every application API call to slow down in the event of a connection failure, reducing application performance significantly. This option is recommended for an audit environment where all events must be logged.
Retry once per TransactionVision event	Retries once for each event. This option reduces the impact on application performance if a connection is lost, but TransactionVision events are lost if the retry attempt is not successful.
Retry across transaction events	Continues to retry at the specified interval as long as event collection is in progress. However, the reconnection time delay is not incurred for each API call, but at a given frequency.
Retry for a specified interval at every TransactionVision event	Continues to retry at the specified interval until the timeout is reached.
Event Delivery Retry	
Retry Timeout	To continue to retry as long as event collection is in progress, choose Retry Forever . Otherwise, specify the number of seconds, minutes, or hours to continue to retry. The default is to retry forever.
Retry Interval	Specify the number of milliseconds or seconds that the Sensor should wait between failed retry attempts.

GUI Buttons	Description
Time Server Definition	
Host name	Specify the host name of the time skew server.
Host port	Specify the port number of the time skew server.
Miscellaneous Information	
Event Queue Maximum Message Length	<p>Specify the maximum event message length. This value can be no larger than the maximum message length of any of the event queues, channels, transmission queues, and queue managers used in this communication link.</p> <p>It is recommended to use the default maximum message length of 4194304 bytes; a minimum message length of 10,000 bytes is required. To restrict the message length based on your applications, consider the length of messages sent by your application and allow an additional 4000 bytes for event information.</p>
Exception Message Queue Name	<p>Specify the name of the exception message queue for the queue manager where the Analyzer places configuration messages. The default is TVISION.EXCEPTION.QUEUE. If left blank, the Analyzer discards events that fail to be placed in the database. For more information about the exception message queue, see “Miscellaneous Information” on page 180.</p>
Number of Event Collection Threads	Specify the number of event collection threads, from 1 to 100. The default value is 1. For more information about event collection threads, see “Miscellaneous Information” on page 180.
Sensor Trace Logging	Specify whether to enable Sensor trace logging. Sensor trace logging is off by default.
Event persistence:	Specify PERSISTENT delivery mode to ensure that an event is not lost in transit in case of a JMS provider failure. An event sent with this delivery mode is logged to a stable storage. Event persistence is off by default.

GUI Buttons	Description
Direct user event processing	Specify whether to support user-defined events processing. User events are events created by user applications beyond those originating from the standard TransactionVision Sensors. For information about implementing user events in applications, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF. By default, the communication link is configured to process standard Sensor events. If your application implements user events, select On to enable the communication link to process user events as well.
Store User Data	Specify whether the analyzer should store the user message data in the database. If this is disabled, the user data is not available in the Event Detail view. Not storing user data reduces the space required for the project and may increase the analyzer performance.
Store Event Document	Specify whether the analyzer should store the event XML data in the database. If this is disabled, the Event Detail view is not available for events collected on this commlink. Not storing the event XML reduces the space required for the project and may increase the analyzer performance. Note that the analyzer might not be able to recover all event analysis data for currently processing events in the case of a crash if the event detail is not available.
Store Raw Events	Specify whether the analyzer should store a copy of the raw, unprocessed events in a special database table. This option should only be used if advised by HP Software Support for troubleshooting purposes, as it can degrade the analyzer performance.

Global Communication Link Templates Page

Description	<p>Enables you to define global communication link templates to be used as the basis for project communication links. A global communication link template identifies the configuration and event queues and queue managers for the Sensor and the Analyzer.</p> <p>To access: Select Admin > TransactionVision > Administration > Communication Link Templates</p>
Important Information	<p>When you create a new project and select a global communication link template, project communication links based on the template are created. The project communication link is a copy of the global communication link template. Changes made to the template do not affect project communication links and changes to project communication links do not affect the template.</p>
Useful Links	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “Communication Link Templates” on page 117 ➤ “Projects” on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
New	Starts the Communication Link Wizard. This wizard guides you through the entire process of creating a new communication link template. See “Communication Link Wizard” on page 146 for information about the first step in creating a new global communication link template.
Edit	Displays the Edit Communication Link Template page on which you can make any desired changes. For details, see “Edit Communication Link Template” on page 184.
Copy	Select an existing communication link and click Copy to create a copy of a communication link named Copy of original_link_name. Select the copy and click Edit to change settings on the Edit Communication Link Template page.
Delete	Click to delete the selected communication link template. You are prompted to confirm the deletion. Click Yes to proceed.
Analyzer drop-down	Select an Analyzer with which to test the selected communication link. See “Test a Communication Link” on page 144 for details.
Test	Click to test the selected communication link and Analyzer. The result displays in the Test Status column. See “Test a Communication Link” on page 144 for details.
<diagram>	Provides a graphical representation of the selected communication link template.



Project Communication Links Page

<p>Description</p>	<p>Enables you to define the communication links for the current project. Use this page to view, create, edit, test, and delete communication links for the current project. This page is similar to the Global Communication Link Templates page, except that it shows the communication link and diagram for the current project instead of all projects — the buttons and functionality are the same. See “Global Communication Link Templates Page” on page 209 for descriptions of the template content.</p> <p>To access: Select Admin > TransactionVision > Administration > Communication Link Templates</p>
<p>Important Information</p>	<p>When you create a new project and select a global communication link template, project communication links based on the template are created. The project communication link is a copy of the global communication link template. Changes made to the template do not affect project communication links and changes to project communication links do not affect the template.</p>
<p>Useful Links</p>	<ul style="list-style-type: none"> ➤ “Understanding Communication Links” on page 106 ➤ “Communication Link Templates” on page 117 ➤ “Global Communication Link Templates Page” on page 209 ➤ “Projects” on page 43

Troubleshooting and Limitations

If your project is not receiving events that you expect it to receive, there may be a problem with the communication link or links you are using. Perform the following steps to diagnose the problem:

- 1** Ensure that you are not using a data collection filter that prevents Sensors from collecting the expected events and that you are not using a project query that prevents you from seeing the events even though they are in the project database.
- 2** Choose **Admin > TransactionVision > Current Projects > Projects** to verify that you have activated at least one communication link for the project.
- 3** Choose **Admin > TransactionVision > Current Projects > Communication Links**, select each activated communication link, pick an Analyzer from the drop-down list, and click **Test**. If the communication link test fails, the error message should indicate which queue, queue manager or event broker failed — make sure the specified queue, queue manager or event broker has been created.
- 4** If the communication link test passes, ensure that a Sensor is intercepting WebSphere MQ API calls from the monitored application. On UNIX and Windows platforms, set the **TVISION_BANNER** environment variable, then start the application. A banner indicating that the application is loading the Sensor should appear. To disable this behavior, unset **TVISION_BANNER**. This environment variable can be set to any value. On Windows, it must be set to a value other than an empty string. On i5/OS, **TVISION_BANNER** does not display the library path as it does on UNIX.

If not, set the environment variable for your platform to cause the application to load the Sensor library instead of the standard WebSphere MQ library.

The following table shows the appropriate environment variables:

Platform	Environment Variable
Windows	PATH
Sun Solaris	LD_LIBRARY_PATH
HP-UX	SHLIB_PATH
Linux	LD_LIBRARY_PATH
IBM AIX	LIBPATH

On z/OS, the CICS region prints a message when it connects to the queue manager or event broker, indicating that the WebSphere MQ crossing exit was found and used. If the WebSphere MQ crossing exit was not found or loaded, a message indicating the error appears.

- 5 If the Sensor library is being loaded by the application, check the application host's system log for any error messages.
- 6 If there are no error messages in the system log, enable trace logging for the active communication links. Look at the trace messages in the application host's UNIX system log, Windows event log, z/OS operator console log, or i5/OS user's job log for any WebSphere MQ-related errors or messages that indicate that the data collection criteria such as collection time or collection entities were too restrictive.

On UNIX platforms, you can specify the log facility by setting the **TVISION_SYSLOG** environment variable to one of the following values: **user**, **local0**, **local1**, **local2**, **local3**, **local4**, **local5**, **local6**, or **local7**. If **TVISION_SYSLOG** is not set or is set to a value other than those listed, TransactionVision uses **local0**. The target log file must already exist.

To view the job log on i5/OS, use the **DSPJOBLOG** command.

If you are still unable to resolve a communication link problem, visit the HP Software Support web site at <http://www.hp.com/go/hpsoftwaresupport>.

7

Data Collection Filters

This chapter includes:

Concepts

- Data Collection Filters Overview on page 215
- Planning Data Collection Filters on page 216
- Filter Conditions on page 219

Tasks

- Create a New Data Collection Filter on page 258
- Edit a Data Collection Filter on page 260

Reference

- Data Collection Filters User Interface on page 261

Data Collection Filters Overview

Collecting all possible event information about all API calls can cause the project database to grow quite large. It also adds to your communication link overhead — depending on project settings, the Sensor may send a message for each event. However, you typically do not need to collect all information about every event that occurs in your system. For example, you may only need to collect events that occur on a specific host, events for a specific API, events with a specific MQI reason code, or events with a specific value in the user data buffer. You may require only a specific amount of user data for each event.

Data collection filters limit the number of events and the amount of data for each event that Sensors collect and forward to the Analyzer. The Analyzer communicates data collection filter criteria to Sensors via the configuration queue and queue manager defined in the communication links assigned to the Analyzer. The Sensor only collects events that match the filter criteria specified in the configuration message.

Only users with required permissions can create or edit data collection filters. See the *HP TransactionVision Deployment Guide* PDF for information about permissions.

Note: If you are collecting from 5.0 Java sensors, you need to be aware of the following limitation: If you wish your 5.0 TransactionVision Java Sensors to filter on an application server, you must manually specify in the data collection filter an entry that only contains the server name. That is, for a 5.0 sensor, if you wished to limit filtering to a WAS instance **WASNode02Cell/WASNode02/server1**, you would enter a custom entry of **server1** in the filter. While this would not allow you to filter on the cell or node name, it would still allow you to match the server name component of the application server's full name.

Planning Data Collection Filters

Data collection filters are unique for each project. You may create any number of filters for a project. You may assign a filter to multiple communication links, and you may assign multiple filters to a single communication link. When planning filters, keep the following in mind:

- An event only meets filter criteria if it matches all criteria specified for the filter. The AND operation is used within each filter.
- The Sensor collects events that meet the criteria for any assigned filter. The OR operation is used between filters.

- For each filter condition, you may specify whether to include or exclude events that meet the filter criteria.
- You may need to define more than one filter to meet your specific requirements.

Note: For the WebSphere MQ and CICS Sensors, a blank field is treated as a match when a data collection filter condition is specified. For example, if you create a filter that collects events with the user name J_SMITH, events with an empty user name are also collected by the WebSphere MQ and CICS Sensors.

Identifying Collection Needs

Before creating data collection filters, first identify exactly which events you need to collect. For example, you may only need to collect events for a particular Sensor type, host, or application. Or you may only need to collect events that use a specific WebSphere MQ object or application server. Or you may need to collect only events for a specific API, with a specific reason code, or with specific data in the user data buffer.

Once you determine which type of events you need to collect, it is helpful to convert your criteria to an algebraic formula and generate a truth table to verify that the equation satisfies the requirement.

For example, suppose you want to collect all events from application X except MQCMIT events, and you also don't want to collect any event with the MQRC_NO_MSQ_AVAILABLE reason code. To state these requirements as an algebraic formula, use the following variables:

A = application X

B = APICODE MQCMIT

C = reason code MQRC_NO_MSQ_AVAILABLE

The resulting formula is as follows:

$$(A \ \& \ !B \ \& \ !C) \ | \ (!A \ \& \ !C)$$

This formula shows two data collection filters. The first specifies the following conditions:

- Program = **include** events for application X
- API = **exclude** MQCMIT
- Reason code = **exclude** MQRC_NO_MSQ_AVAILABLE The second filter specifies the following conditions:
- Program = **exclude** events for application X
- Reason code = **exclude** MQRC_NO_MSQ_AVAILABLE

The following truth table shows that this combination of data collection filters satisfies the filter requirements:

A	B	C	(A & !B & !C)	(!A & !C)	(A & !B & !C) (!A & !C)
0	0	0	0	1	1
0	0	1	0	0	0
0	1	0	0	1	1
0	1	1	0	0	0
1	0	0	1	0	1
1	0	1	0	0	0
1	1	0	0	0	0
1	1	1	0	0	0

Data Collection Filters and Communication Links

Data collection filters assigned to communication links should have only the required MQ objects, applications, hosts, and APIs specified for collection. MQGET calls with NO_MSG_AVAIL reason code should generally be filtered out. For information about communication links, see “Communication Links” on page 105.

Filter Conditions

The filter settings are set by default for each technology event type. You can change the default filter settings in the Edit Data Collection Filter page for all or individual technologies. Sensors only collect and send events that meet their respective filter criteria.

This section describes the following technology criteria:

- “Common Options” on page 219
- “CICS Sensor Criteria” on page 225
- “EJB Sensor Criteria” on page 229
- “JDBC Sensor Criteria” on page 231
- “JMS Sensor Criteria” on page 235
- “Servlet Sensor Criteria” on page 237
- “User Event Sensor Criteria” on page 239
- “WebSphere MQ Sensor Criteria” on page 248

Common Options

The following criteria apply to all Sensors. However, if you specify a filter criterion for **All Sensors**, you have the option to specify a different filter criterion for an individual technology (Sensor) as an exception. The setting for the individual Sensor overrides the setting for all Sensors. For example, you select **All Hosts** for **All Sensors**, but on the Host page for the JMS Sensor, you may select only the host Host1. JMS events originating on Host1 and other Sensor events originating on any host all match the filter criteria.

Following are the common criteria to filter events for all Sensors:

Host

- Determine whether to Include or Exclude selected hosts in the filter criteria.
- Check the box next to each host you want to include specifically in the filter criteria, or check the box next to **All Host** to include all hosts in the filter criteria. The host list is generated from events collected in the project database.
- To manually add a new host to the list, enter the host name and click **Add**. To remove a host name from the list, check the box next to the host name and click **Delete**. Only hosts that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Program/Servlet/EJB

Note: For the NonStop Transaction Management Facility (TMF) Sensor, the Program/Servlet/EJB criterion is not evaluated.

- Determine whether to **Include** or **Exclude** selected programs, servlets, and/or EJBs in the filter criteria.
- Check the box next to each program, servlet, or EJB you want to include in the filter criteria, or check the box next to **All Program/Servlet/EJB** to include all programs, servlets, and EJBs in the filter criteria. The program list is generated from events collected in the project database.
- To manually add a new program to the list, enter the program name and click **Add**. To remove a program from the list, check the box next to the program name and click **Delete**. Only programs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Note: For the BEA Tuxedo Sensor the Program/Servlet/EJB criterion must be set to include the name of the executable process that made the Tuxedo call. This can be done by either including the name explicitly or by including All Program/Servlet/EJB.

Time

Note: For the .NET Sensor, BEA Tuxedo Sensor, and NonStop TMF Sensor the **Time** criterion is not evaluated.

Determine the desired time mode:

- If you select **Any Time**, the events are collected with no time restrictions.
 - If you select **From a set start to end time**, specify the start date and time, the end date and time, and select the time zone. Specify a time range (based on the clock of the host running the Sensor) to collect events. To select the date from a calendar, click the calendar button and then click the desired date. The events are collected from the specified start time until the specified end time.
 - If you select **From a set start time**, specify the start date and time and select the time zone. Specify a time range (based on the clock of the host running the Sensor) to collect events. The events are collected from a specific start time until you explicitly stop recording.
-

Note: Even if you specify a start time, you must initialize event recording before that start time. Initializing event recording causes the Analyzer to send a configuration message with filter information to all active communication links.

User

Note: For the .NET Sensor, BEA Tuxedo Sensor, and NonStop TMF Sensor the **User** criterion is not evaluated.

- Specify whether to **Include** or **Exclude** events with selected users.
- Check the box next to each user you want to include specifically in the filter criteria, or check the box next to **All User** to include all users in the filter criteria. The user list is generated from events collected in the project database.
- To manually add a new user to the list, enter the user and click **Add**. To remove a user from the list, check the box next to the user and click **Delete**. Only users that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Data Range

Note: For the BEA Tuxedo Sensor and NonStop TMF Sensor the **Data Range** criterion is not evaluated.

The options in the drop-down list vary depending on the technology (Sensor):

- For All Sensors, EJB, and WebSphere MQ technologies:
Specify whether to collect the **API Names Only**, the **API Names + Call Arguments**, or the **API Names + Call Arguments + User Buffer**.
If you collect API names only, some TransactionVision functionality is not available:
 - You are not able to view event details for events collected with this filter condition.

- The Transaction Analysis is not able to correlate MQPUT and MQGET calls.
- You are not able to query these events in the project based on API parameter values.
- For CICS, JDBC, JMS, Servlet, User Event Sensors technologies:

Specify whether to **Collect User Data** or **Do not Collect User Data**.

If you specify a data range that includes user buffer data, you may select a range to specify how many bytes of the user data buffer to collect. Check the box next to **All Data (no data ranges applied)** to collect the entire user data buffer.

To manually add a new range, enter the starting byte number (0 based index) in the Start Value field and the ending byte number in the End Value field and click **Add**. For example, use a start value of 0 and an end value of 99 to collect the first 100 bytes of the user data buffer. Specify a * in the End Value field to collect to the end of the user data buffer.

Sensors collect data in all selected data range definitions, so remove any data ranges you do not want to collect. To remove an existing range, select it and click **Delete**.

Event Packaging

Note: For the NonStop TMF Sensor the **Event Packaging** criterion is not evaluated.

- To enable Sensors to package multiple events together before sending them to the Analyzer, select **Use Event Packaging**.

When you use event packaging, the Sensor buffers events until the event packaging criteria is met (or other conditions within the Sensor force it to flush the package), then sends the packaged events to the event queue. If the Sensor has buffered events for event packaging and the project is stopped, it sends the buffered events to the event queue as soon as it is able to (depending on the type of Sensor).

Event Packaging directly improves the performance of the Sensor. However, in a development environment where application crashes could be more common, event packaging may result in some collected events not being sent to the Analyzer.

- Specify the following options in the Event Packaging area:

Option	Description
Maximum Package Size	<p>Specify the limit on the amount of memory an Sensor can set aside for storing events to be packaged. When the Sensor reaches or exceeds this limit while packaging events together, it sends the package to the event queue. Leave this field blank or specify a value of zero to use the number of events field exclusively.</p> <p>Note that if the Maximum Package Size is smaller than the event size, each event is sent immediately. In this case, Event Packaging is effectively turned off. The minimum event size is approximately 4000 bytes.</p>
Number of Events	<p>Specify the number of events that stored in the Sensor before sending a package to the Analyzer. As soon as the Sensor reaches this number of events while packaging events together, it sends the package to the event queue.</p> <p>The number of event should be set to a value suitable to the environment, usually 10 events per message.</p> <p>Leave this field blank or specify a value of zero to use the maximum package size field exclusively.</p>

If you specify a value for both the number of events and the maximum package size, Sensors combine messages into packages based on which limit is reached first. If both are zero or blank, the default behavior of no event packaging is used.

CICS Sensor Criteria

Following are the CICS criteria to filter CICS events:

API Name

This page allows you to filter CICS events by their API names as well as enable or disable CICS exits to filter a group of APIs.

- Determine whether to **Include** or **Exclude** events for selected APIs. To filter on selected exits, select Include and the API or API type check boxes for exits you wish to filter on. Selected APIs are enabled; all other exits are disabled. To disable specific exits, select **Exclude** and the API or API type check boxes for exits you wish to disable. All other exits are enabled.
- Check the box next to each API you want to enable or disable, or check the box next to **All API** to enable or disable all APIs. The APIs are grouped by API type. To select all APIs for a specific type, check the box next to the API type.

For example, if you select only the File Control exit to Include, all other exits are stopped and information messages are shown on the operator console. Similarly, if you exclude File Control exits, all other exits are enabled.

Transaction ID

- Determine whether to **Include** or **Exclude** events for selected CICS transaction IDs.
- Check the box next to each transaction ID you want to include specifically in the filter criteria, or check the box next to **All Transaction ID** to include all transaction IDs in the filter criteria. The list of transaction IDs is generated from events collected in the project database.
- To manually add a transaction ID to the list, enter its name and click **Add**. To remove a transaction ID from the list, check the box next to it and click **Delete**. Only transaction IDs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

User ID

- Determine whether to **Include** or **Exclude** events for selected user IDs.
- Check the box next to each user ID you want to include specifically in the filter criteria, or check the box next to **All User ID** to include all user IDs in the filter criteria. The list of user IDs is generated from events collected in the project database.
- To manually add a user ID to the list, enter its name and click **Add**. To remove a user ID from the list, check the box next to it and click **Delete**. Only user IDs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

File Name

- Determine whether to **Include** or **Exclude** events for selected file names.
- Check the box next to each file name you want to include specifically in the filter criteria, or check the box next to **All File Name** to include all file names in the filter criteria. The list of file names is generated from events collected in the project database.
- To manually add a file name to the list, enter its name and click **Add**. To remove a file name from the list, check the box next to it and click **Delete**. Only file names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Response Code

- Determine whether to **Include** or **Exclude** events for selected response codes.
- Check the box next to each response code you want to include specifically in the filter criteria, or check the box next to **All Response Code** to include all response codes in the filter criteria.

SYSID

- Determine whether to **Include** or **Exclude** events for selected CICS SYSIDs.
- Check the box next to each SYSID you want to include specifically in the filter criteria, or check the box next to **All SYSID** to include all SYSIDs in the filter criteria. The list of SYSIDs is generated from events collected in the project database.
- To manually add a SYSID to the list, enter its name and click **Add**. To remove a SYSID from the list, check the box next to it and click **Delete**. Only SYSIDs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

TD Queue Name

- Determine whether to **Include** or **Exclude** events for selected transient data (TD) queue names.
- Check the box next to each TD queue name you want to include specifically in the filter criteria, or check the box next to **All TD Queue Name** to include all TD queue names in the filter criteria. The list of TD queue names is generated from events collected in the project database.
- To manually add a TD queue name to the list, enter its name and click **Add**. To remove a TD queue name from the list, check the box next to it and click **Delete**. Only TD queue names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

TS Queue ID

- Determine whether to Include or Exclude events for selected temporary storage (TS) queue IDs.
- Check the box next to each TS queue ID you want to include specifically in the filter criteria, or check the box next to **All TS Queue ID** to include all TS queue IDs in the filter criteria. The list of TS queue IDs is generated from events collected in the project database.
- To manually add a TS queue ID to the list, enter its name and click **Add**. To remove a TS queue ID from the list, check the box next to it and click **Delete**. Only TS queue names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Terminal ID

- Determine whether to **Include** or **Exclude** events for selected terminal IDs.
- Check the box next to each terminal ID you want to include specifically in the filter criteria, or check the box next to **All Terminal ID** to include all terminal IDs in the filter criteria. The list of terminal IDs is generated from events collected in the project database.
- To manually add a terminal ID to the list, enter its name and click **Add**. To remove a terminal ID from the list, check the box next to it and click **Delete**. Only TD queue names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

 **EJB Sensor Criteria**

Following are the EJB criteria to filter EJB events:

Application Server

- Determine whether to **Include** or **Exclude** events for selected application servers.
- Check the box next to each application server you want to include specifically in the filter criteria, or check the box next to **All Application Server** to include all application servers in the filter criteria. The list of application servers is generated from events collected in the project database.
- To manually add an application server to the list, enter its name and click **Add**. To remove an application server from the list, check the box next to it and click **Delete**. Only application servers that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Note: If you are collecting from 5.0 Java sensors, you need to be aware of the following limitation: If you wish your 5.0 TransactionVision Java Sensors to filter on an application server, you must manually specify in the data collection filter an entry that only contains the server name. That is, for a 5.0 Sensor, if you wished to limit filtering to a WAS instance WASNode02Cell/WASNode02/server1, you would enter a custom entry of server1 in the filter. While this would not allow you to filter on the cell or node name, it would still allow you to match the server name component of the application server's full name.

Application Name

- Determine whether to **Include** or **Exclude** events for selected application names.
- Check the box next to each application name you want to include specifically in the filter criteria, or check the box next to **All Application Name** to include all application names in the filter criteria. The list of application names is generated from events collected in the project database.
- To manually add an application name to the list, enter the name and click **Add**. To remove an application name from the list, check the box next to it and click **Delete**. Only application names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

EJB Name

- Determine whether to **Include** or **Exclude** events for selected EJB names.
- Check the box next to each EJB name you want to include specifically in the filter criteria, or check the box next to **All EJB Name** to include all application names in the filter criteria. The list of EJB names is generated from events collected in the project database.
- To manually add a EJB name to the list, enter its name and click **Add**. To remove an EJB name from the list, check the box next to it and click **Delete**. Only EJB names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

EJB Method

- Determine whether to **Include** or **Exclude** events for selected EJB methods.
- Check the box next to each EJB method you want to include specifically in the filter criteria, or check the box next to **All EJB Method** to include all application names in the filter criteria. The list of EJB methods is generated from events collected in the project database.
- To manually add an EJB method to the list, enter its name and click **Add**. To remove an EJB method from the list, check the box next to it and click **Delete**. Only EJB methods that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Exception Setting

Determine whether to include both normal and exception events, normal events only, or exception events only.



JDBC Sensor Criteria

Following are the JDBC criteria to filter JDBC events:

Class

- Determine whether to Include or Exclude events for selected JDBC classes.
- Check the box next to each JDBC class you want to include specifically in the filter criteria, or check the box next to **All JDBC Class** to include all classes in the filter criteria.

Database

- Determine whether to **Include** or **Exclude** events for selected JDBC database.
- Check the box next to each database you want to include specifically in the filter criteria, or check the box next to **All Database** to include all databases in the filter criteria. The list of databases is generated from events collected in the project database.
- To manually add a custom database to the list, enter its name and click **Add**. To remove a custom database entry from the list, check the box next to it and click **Delete**. Only databases that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Database Object

- Determine whether to **Include** or **Exclude** events that make use of certain database objects
- Check the box next to each database object you want to include specifically in the filter criteria, or check the box next to **All Entries** to include all objects in the filter criteria. The list of database objects is generated from events collected in the project database.

Database objects are grouped by the schema that they belong to. For example, selecting a database table 'X' under schema 'Y' only selects JDBC events that access table 'X', if table 'X' is being used under schema 'Y'. Additionally by selecting a schema, you can filter on all tables that might occur under that schema.

- To manually add a custom database object to the list, enter the schema name you want this object to be under and the database object name and click **Add**. To remove a custom database object from the list, check the box next to it and click **Delete**. Only database objects that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Method

- Determine whether to **Include** or **Exclude** events for selected JDBC methods.
- Check the box next to each JDBC method you want to include specifically in the filter criteria, or check the box next to **All Method** to include all methods in the filter criteria.

Result Code

- Determine whether to **Include** or **Exclude** events that have a particular result code.
- Check the box next to each result code you want to include specifically in the filter criteria, or check the box next to **All Result Code**. Result codes can be one of Error, Success or Warning.

SQL Code

- Determine whether to **Include** or **Exclude** events that have a particular SQL Code.
- Check the box next to each SQL code you want to include specifically in the filter criteria, or check the box next to **All SQL Code** to include all SQL codes in the filter criteria. The list of SQL codes is generated from events collected in the project database.
- To manually add an SQL code to the list, enter its name and click **Add**. To remove an SQL code from the list, check the box next to it and click **Delete**. Only SQL codes that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

SQL codes used by the JDBC application vary based on the underlying database vendor.

SQL State

- Determine whether to **Include** or **Exclude** events that have a particular SQL state.
- Check the box next to each SQL state you want to include specifically in the filter criteria, or check the box next to **All SQL State** to include all SQL states in the filter criteria. The list of SQL states is generated from events collected in the project database.
- To manually add an SQL state to the list, enter its name and click **Add**. To remove an SQL state from the list, check the box next to it and click **Delete**. Only SQL states that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Similar to SQL Code, SQL state values used by the JDBC application vary based on the underlying database vendor.

SQL Statement

- Determine whether to **Include** or **Exclude** events for a particular SQL statement type.
- Check the box next to the statement type (a Select, Insert, update, and so on) you want to include specifically in the filter criteria, or check the box next to **All SQL Statement** to include all classes in the filter criteria.

Web Application

- Determine whether to **Include** or **Exclude** events for selected web applications.
- Check the box next to each web application you want to include specifically in the filter criteria, or check the box next to **All Web Application** to include all web applications in the filter criteria. The list of web applications is generated from events collected in the project database.
- To manually add a web application to the list, enter its name and click **Add**. To remove a web application from the list, check the box next to it and click **Delete**. Only web applications that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

 **JMS Sensor Criteria**

Following are the JMS criteria to filter JMS events:

Web Application

- 1** Select whether to Include or Exclude events for selected web applications.
- 2** Check the box next to each web application you want to include specifically in the filter criteria, or check the box next to **All Web Application** to include all web applications in the filter criteria. The list of web applications is generated from events collected in the project database.
- 3** To manually add a web application to the list, enter its name and click **Add**. To remove a web application from the list, check the box next to it and click **Delete**. Only web applications that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Class

- Determine whether to **Include** or **Exclude** events for selected JMS classes.
- Check the box next to each JMS class you want to include specifically in the filter criteria, or check the box next to **All JMS Class** to include all classes in the filter criteria.

Method

- Determine whether to **Include** or **Exclude** events for selected JMS methods.
- Check the box next to each JMS method you want to include specifically in the filter criteria, or check the box next to **All JMS Method** to include all methods in the filter criteria.
- Check **Do not collect events when no message is available** to avoid collecting JMS Consumer events with no message.

Topic

- Determine whether to **Include** or **Exclude** events for selected topics.
- Check the box next to each topic you want to include specifically in the filter criteria, or check the box next to **All Topic** to include all topics in the filter criteria. The list of topics is generated from events collected in the project database.
- To manually add a topic to the list, enter its name and click **Add**. To remove a topic from the list, check the box next to it and click **Delete**. Only topics that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

JMS Queue

- Determine whether to **Include** or **Exclude** events for selected queues.
- Check the box next to each queue you want to include specifically in the filter criteria, or check the box next to **All Queue** to include all queues in the filter criteria. The list of queues is generated from events collected in the project database.
- To manually add a queue to the list, enter its name and click **Add**. To remove a queue from the list, check the box next to it and click **Delete**. Only queues that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Connection Name

- Determine whether to **Include** or **Exclude** events for selected connection names.
- Check the box next to each connection you want to include specifically in the filter criteria, or check the box next to **All Connection Name** to include all connections in the filter criteria. The list of connections is generated from events collected in the project database.
- To manually add a connection name to the list, enter its name and click **Add**. To remove a connection name from the list, check the box next to it and click **Delete**. Only connection names that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Exception Setting

Select whether to include both normal and exception events, normal events only, or exception events only.



Servlet Sensor Criteria

Following are the Servlet criteria to filter Servlet events:

Application Server

- Determine whether to **Include** or **Exclude** events for selected application servers.
- Check the box next to each application server you want to include specifically in the filter criteria, or check the box next to **All Application Server** to include all application servers in the filter criteria. The list of application servers is generated from events collected in the project database.
- To manually add an application server to the list, enter its name and click **Add**. To remove an application server from the list, check the box next to it and click **Delete**. Only application servers that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Web Application

- Determine whether to **Include** or **Exclude** events for selected web applications.
- Check the box next to each web application you want to include specifically in the filter criteria, or check the box next to **All Web Application** to include all web applications in the filter criteria. The list of web applications is generated from events collected in the project database.
- To manually add a web application to the list, enter its name and click **Add**. To remove a web application from the list, check the box next to it and click **Delete**. Only web applications that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Note: If you are collecting from 5.0 Java sensors, you need to be aware of the following limitation: If you wish your 5.0 TransactionVision Java Sensors to filter on an application server, you must manually specify in the data collection filter an entry that only contains the server name. That is, for a 5.0 Sensor, if you wished to limit filtering to a WAS instance WASNode02Cell/WASNode02/server1, you would enter a custom entry of server1 in the filter. While this would not allow you to filter on the cell or node name, it would still allow you to match the server name component of the application server's full name.

Servlet Method

- Select whether to Include or Exclude events for selected servlet methods.
- Check the box next to each servlet method you want to include specifically in the filter criteria, or check the box next to **All Servlet Method** to include all servlet methods in the filter criteria.

Client Host/IP

- Determine whether to **Include** or **Exclude** events for selected client host/IPs.
- Check the box next to each client you want to include specifically in the filter criteria, or check the box next to **All Client Host/IP** to include all clients in the filter criteria. The list of client host/IPs is generated from events collected in the project database.
- To manually add a client/host/IP to the list, enter its name and click **Add**. To remove a client host/IP from the list, check the box next to it and click **Delete**. Only client host/IPs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Note: The Client Host/IP setting does not affect the **init** method.

URI

- Determine whether to **Include** or **Exclude** events for selected URIs.
- Check the box next to each URI you want to include specifically in the filter criteria, or check the box next to **All URI** to include all URIs in the filter criteria. The list of URIs is generated from events collected in the project database.
- To manually add a URI to the list, enter its name and click **Add**. To remove a URI from the list, check the box next to it and click **Delete**. Only URIs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Status Code

- Determine whether to **Include** or **Exclude** events for selected status codes.
- Check the box next to each status code you want to include specifically in the filter criteria, or check the box next to **All Status Code** to include all status codes in the filter criteria.

Exception Setting

Determine whether to include both normal and exception events, normal events only, or exception events only.

User Event Sensor Criteria

The following sections describe the user event criteria that can be used to filter the User Event Sensor events, and specific information about User Event Sensors provided with TransactionVision: BEA Tuxedo, .NET, NonStop TME, and Java User Event.

- “User Event Criteria Descriptions” on page 240
- “BEA Tuxedo Sensor” on page 244
- “.NET Sensor” on page 245
- “NonStop TME Sensor” on page 246
- “Java User Event Sensor” on page 247

The following sections describe the user event criteria for each User Event Sensor that can be used to filter the User Event Sensor events:

User Event Criteria Descriptions

Following are the User Event criteria to filter user events:

Technology

Specifies the type of User Event Sensor: .NET, Tuxedo, NonStop-TMF, or Java.

Class

Note: For the BEA Tuxedo Sensor, the Class criterion is not evaluated

Specifies which classes are filtered.

- Determine whether to **Include** or **Exclude** events for selected User Event classes.
- Check the box next to each User Event class you want to include specifically in the filter criteria, or check the box next to **All Class** to include all classes in the filter criteria. The list of classes is generated from events collected in the project database.
- To manually add a class to the list, enter its name and click **Add**. To remove a class from the list, check the box next to it and click **Delete**. Only classes that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Method

Specify which methods are filtered:

- Determine whether to **Include** or **Exclude** events for selected User Event methods.

- Check the box next to each User Event method you want to include specifically in the filter criteria, or check the box next to **All Method** to include all methods in the filter criteria. The list of methods is generated from events collected in the project database.
- To manually add a method to the list, enter its name and click **Add**. To remove a method from the list, check the box next to it and click **Delete**. Only methods that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Note: For BEA Tuxedo, **Method** must match **tpenqueue**, **tpdequeue** or **tpcall**, as appropriate per operation.

Completion Code

Note: Not available for BEA Tuxedo, .NET, NonStop TME, or Java User Event Sensors.

- Determine whether to **Include** or **Exclude** events for selected User Event method completion codes.
- Check the box next to each method completion code you want to include specifically in the filter criteria, or check the box next to **All Completion Codes** to include all possible User Event completion codes in the filter criteria.

Status

Note: Not available for .NET, BEA Tuxedo, or NonStop TMF Sensors.

- Determine whether to **Include** or **Exclude** events for selected User Event status strings.
- Check the box next to each status string you want to include specifically in the filter criteria, or check the box next to **All Status** to include all status strings in the filter criteria. The list of status strings is generated from events collected in the project database.
- To manually add a status string to the list, enter its name and click **Add**. To remove a Status string from the list, check the box next to it and click **Delete**. Only status strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Queue

Note: Not available for .NET, NonStop TMF, or Java User Event Sensors.

For the BEA Tuxedo Sensor, **Queue** must match the name of the queue for which the **tpenqueue/tpdequeue** operation is intended.

- Determine whether to **Include** or **Exclude** events for selected User Event queue strings.

- Check the box next to each queue string you want to include specifically in the filter criteria, or check the box next to **All Queue** to include all queue strings in the filter criteria. The list of queue strings is generated from events collected in the project database.
- To manually add a queue string to the list, enter its name and click **Add**. To remove a queue string from the list, check the box next to it and click **Delete**. Only queue strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Queue Space

Note: Not available for .NET, NonStop TME, or Java User Event Sensors.

For the BEA Tuxedo Sensor, **Queue Space** must match the queue space for which the **tpenqueue/tpdequeue** operation is intended.

- Determine whether to **Include** or **Exclude** events for selected User Event queue space strings.
- Check the box next to each queue space string you want to include specifically in the filter criteria, or check the box next to **All Queue Space** to include all queue space strings in the filter criteria. The list of queue space strings is generated from events collected in the project database.
- To manually add a queue space string to the list, enter its name and click **Add**. To remove a queue space string from the list, check the box next to it and click **Delete**. Only queue space strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Service

Note: For the BEA Tuxedo Sensor, **Service** must match the service to which the **tpcall** request is sent.

- Determine whether to **Include** or **Exclude** events for selected User Event service strings.
- Check the box next to each service string you want to include specifically in the filter criteria, or check the box next to **All Service** to include all service strings in the filter criteria. The list of service strings is generated from events collected in the project database.
- To manually add a service string to the list, enter its name and click **Add**. To remove a service string from the list, check the box next to it and click **Delete**. Only service strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

BEA Tuxedo Sensor

Note: Only one data collection filter is applied by the BEA Tuxedo Sensor. The BEA Tuxedo Sensor applies the first data collection filter it finds that matches the Technology and Program/Servlet/EJB criteria. It is recommended that you enable only one data collection filter that specifies the desired criteria. This is because it is difficult to determine the order in which the BEA Tuxedo Sensor receives and processes multiple data collection filters.

The BEA Tuxedo Sensor uses a subset of the criteria for User Event Sensors. The following criteria are used to filter BEA Tuxedo events. Each may be matched either explicitly, by including the exact value for a criterion, or implicitly, by selecting **All** for a criterion:

- **Host.** Must match the host name on which the BEA Tuxedo Sensor is running. See “Host” on page 220 for instructions on setting this criterion.
- **Program/Servlet/EJB.** Must be set to include the name of the executable process that made the Tuxedo call. See “Program/Servlet/EJB” on page 220 for instructions on setting this criterion.
- **Technology.** See “Technology” on page 240 for instructions on setting this criterion.
- **Method.** See “Method” on page 240 for instructions on setting this criterion.
- **Queue.** See “Queue” on page 242 for instructions on setting this criterion.
- **Queue Space.** See “Queue Space” on page 243 for instructions on setting this criterion.
- **Service.** See “Service” on page 244 for instructions on setting this criterion.

Note: All User Event and Common Options criteria not mentioned above are not evaluated for BEA Tuxedo Sensors.

.NET Sensor

The .NET Sensor uses a subset of the criteria for User Event Sensors. The following criteria are used to filter .NET events:

- **Host.** See “Host” on page 220 for instructions on setting this criterion.
- **Program/Servlet/EJB.** See “Program/Servlet/EJB” on page 220 for instructions on setting this criterion.
- **Data Range.** See “Data Range” on page 222 for instructions on setting this criterion.

- **Event Packaging.** See “Event Packaging” on page 223 for instructions on setting this criterion.
- **Technology.** See “Technology” on page 240 for instructions on setting this criterion.
- **Class.** See “Class” on page 240 for instructions on setting this criterion.
- **Method.** See “Method” on page 240 for instructions on setting this criterion.

Note: All User Event and Common Options criteria not mentioned above are not evaluated for .NET Sensors.

NonStop TMF Sensor

Note: Only one data collection filter is applied by the NonStop TMF Sensor. The NonStop TMF Sensor applies the last data collection filter it finds that matches the criteria. It is recommended that you enable only one data collection filter that specifies the desired criteria. This is because it is difficult to determine the order in which the NonStop TMF Sensor receives and processes multiple data collection filters.

The NonStop TMF Sensor uses a subset of the criteria for User Event Sensors. The following criteria are used to filter NonStop TMF events.

- **Host.** See “Host” on page 220 for instructions on setting this criterion.
- **Program/Servlet/EJB.** See “Program/Servlet/EJB” on page 220 for instructions on setting this criterion.
- **Technology.** See “Technology” on page 240 for instructions on setting this criterion.

- **Class.** Must be set to include the name of the Enscribe files to monitor. This must be done by explicitly selecting the files. This Sensor does not collect data for the **All Classes** selection. See “Class” on page 240 for instructions on setting this criterion.
- **Method.** May be set to filter by Enscribe file **update**, **insert**, or **delete**.

Note: All User Event and Common Options criteria not mentioned above are not evaluated for NonStop TMF Sensors.

Java User Event Sensor

The Java User Event Sensor uses a subset of the criteria for User Event Sensors. The following criteria are used to filter Java User events:

- **Host.** See “Host” on page 220 for instructions on setting this criterion.
- **Program/Servlet/EJB.** See “Program/Servlet/EJB” on page 220 for instructions on setting this criterion.
- **Data Range.** See “Data Range” on page 222 for instructions on setting this criterion.
- **Event Packaging.** See “Event Packaging” on page 223 for instructions on setting this criterion.
- **Technology.** See “Technology” on page 240 for instructions on setting this criterion.
- **Class.** See “Class” on page 240 for instructions on setting this criterion.
- **Method.** See “Method” on page 240 for instructions on setting this criterion.
- **Completion Code.** See “Completion Code” on page 241 for instructions on setting this criterion.

- **Status.** See “Status” on page 242 for instructions on setting this criterion.

Note: All User Event and Common Options criteria not mentioned above are not evaluated for Java Sensors.



WebSphere MQ Sensor Criteria

Following are the WebSphere MQ criteria to filter WebSphere MQ events:

WebSphere MQ API Name

- Determine whether to **Include** or **Exclude** selected WebSphere MQ APIs in the filter criteria.
- Check the box next to each API you want to include specifically in the filter criteria, or check the box next to **All API NAME** to include all WebSphere MQ APIs in the filter criteria.
- Check **Discard MQCMIT in empty transactions** to avoid collecting empty transactions. Empty transactions occur in applications that pool a queue by making **MQGET** calls in a loop. When there are no inputs, the **MQGET** calls fail with a **MQRC_NO_MSG_AVAILABLE** and are typically filtered out by customers in TransactionVision. However, the **MQGET** may be followed by an **MQCMIT**. This option provides an easy way to filter out these **MQCMIT** events.
- Check **Do not send MQDISC exit event** to report only the entry information of each **MQDISC** call. This improves performance and reduces overhead for the Sensor. Otherwise, the Sensor must reconnect and reopen the event queue in order to send **MQDISC** exit information, and this information is typically not necessary.
- Check **Do not send browsing MQGET** if you do not wish to collect **MQGET** calls used to browse messages on the queue.

WebSphere MQ Completion Code

- Determine whether to **Include** or **Exclude** events for selected WebSphere MQ API completion codes.
- Check the box next to each API completion code you want to include specifically in the filter criteria, or check the box next to **All Completion Code** to include all WebSphere MQ completion codes in the filter criteria.

Reason Code

- Determine whether to **Include** or **Exclude** events for selected WebSphere MQ API reason codes.
- Check the box next to each API reason code you want to include specifically in the filter criteria, or check the box next to **All Reason Code** to include all reason codes in the filter criteria.

Note: MQGET calls with a **MQRC_NO_MSG_AVAILABLE** reason code are filtered out by default. When monitoring WebSphere MQ Brokers, make sure to filter out MQGET calls to WebSphere MQ Broker system queues with a reason code of **MQRC_NO_MSG_AVAILABLE**.

Object

- Determine whether to **Include** or **Exclude** selected queue managers and objects in the filter criteria.

- Check the box next to each queue manager and/or object you want to include specifically in the filter criteria, or check the box next to **All Object** to include all objects in the filter criteria. Selecting objects not listed under a queue manager is useful if the same object name is present on multiple queue managers (for example, as in clustered queues). The object list is generated from events collected in the project database.
- To manually add a new queue manager/object pair to the list, enter the names and click **Add**. To remove a queue manager or object name from the list, check the box next to it and click **Delete**. Only queue managers and objects that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard. The queue manager or object field may be set to an asterisk to represent all queue managers or objects.

ReplyTo Object

- Determine whether to **Include** or **Exclude** selected ReplyTo queue managers and objects in the filter criteria.
- Check the box next to each queue manager and/or object you want to include specifically in the filter criteria, or check the box next to **All ReplyTo Object** to include all objects in the filter criteria. Selecting objects not listed under a queue manager is useful if the same object name is present on multiple queue managers (for example, as in clustered queues). The object list is generated from events collected in the project database.
- To manually add a new ReplyTo queue manager/object pair to the list, enter the names and click **Add**. To remove a queue manager or object name from the list, check the box next to it and click **Delete**. Only queue managers and objects that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard. The queue manager or object field may be set to an asterisk to represent all queue managers or objects.
- To collect events that do not have the MQMD ReplyTo fields, select **Collect MQPUT, MQPUT1, and MQGET events if both "ReplyTo" fields are not set**.

WBI Message Broker

- Determine whether to **Include** or **Exclude** events for selected WebSphere Business Integrator (WBI) brokers.
- Check the box next to each broker you want to include specifically in the query criteria, or check the box next to **All WBI Message Broker** to include all brokers in the query criteria. The list of brokers is generated from events collected in the project database.
- To manually add a new message flow to the list, enter the name and click **Add**. To remove a message flow name from the list, check the box next to it and click **Delete**. Only WBI message brokers that were manually added and do not exist in the project may be deleted.

WBI Message Flow

- Determine whether to **Include** or **Exclude** selected WebSphere Business Integrator (WBI) message flows in the filter criteria.
- Check the box next to each message flow you want to include specifically in the filter criteria, or check the box next to **All WBI Message Flow** to include all flows in the filter criteria. The message flow list is generated from events collected in the project database.
- To manually add a new message flow to the list, enter the name and click **Add**. To remove a message flow name from the list, check the box next to it and click **Delete**. Only WBI message flows that were manually added and do not exist in the project may be deleted.

MQIMS Bridge API

- Determine whether to **Include** or **Exclude** events for selected MQIMS bridge APIs.
- Check the box next to the APIs you want to include specifically in the filter criteria, or check the box next to **All MQIMS Bridge API** to include all MQIMS bridge APIs in the filter criteria.

For more information about the MQSeries-IMS Bridge Sensor, see the *HP TransactionVision Deployment Guide* PDF.

z/OS Jobs and Steps

- Determine whether to **Include** or **Exclude** selected z/OS job steps in the filter criteria.
- Check the box next to each job name or step you want to include specifically in the filter criteria, or check the box next to **All Jobs and Steps** to include all job steps in the filter criteria. The job name and job step list is generated from events already collected in the project database.
- To manually add a new job/step pair to the list, enter the names and click **Add**. To remove a job or step from the list, check the box next to it and click **Delete**. Only jobs and steps that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard. The job or step field may be set to an asterisk to represent all jobs or steps.
- Only events generated on the z/OS platform have job names and job steps. To collect events that do not have a CICS job and step (non-z/OS events), choose **Collect events regardless of whether the field "Jobs and Steps" is available**.
- Choose **Collect events if and only if the field "Jobs and Steps" is available** to collect events that only have the CICS job names and steps.

z/OS CICS Task

- Determine whether to **Include** or **Exclude** events for selected CICS tasks.
- Check the box next to each CICS task you want to include specifically in the filter criteria, or check the box next to **All CICS Task** to include all CICS tasks in the filter criteria. The list of CICS tasks is generated from events collected in the project database.
- To manually add a new task to the list, enter the name and click **Add**. To remove a task from the list, check the box next to it and click **Delete**. Only tasks that were manually added and do not exist in the project may be deleted.

- Only events generated on the z/OS CICS platform have a CICS task. To collect events that do not have a CICS task (non-z/OS CICS events), choose **Collect events regardless of whether the field "CICS Task" is available**.
- Choose **Collect events if and only if the field "CICS Task" is available** to collect events that only have CICS tasks.

z/OS CICS SYSID

- Determine whether to **Include** or **Exclude** events for selected CICS SYSIDs.
- Check the box next to each CICS SYSID you want to include specifically in the filter criteria, or check the box next to **All CICS SYSID** to include all CICS SYSIDs in the filter criteria. The list of CICS SYSIDs is generated from events collected in the project database.
- To manually add a new SYSID to the list, enter the name and click **Add**. To remove a SYSID from the list, check the box next to it and click **Delete**. Only SYSIDs that were manually added and do not exist in the project may be deleted.
- Only events generated on the z/OS CICS platform have a CICS SYSID. To collect events that do not have a CICS SYSID (non-z/OS CICS events), choose **Collect events regardless of whether the field "CICS SYSID" is available**.
- Choose **Collect events if and only if the field "CICS SYSID" is available** to collect events that only have CICS SYSIDs.

z/OS CICS Transaction

- Determine whether to **Include** or **Exclude** events for selected z/OS CICS transactions.
- Check the box next to each CICS transaction you want to include specifically in the filter criteria, or check the box next to **All CICS Transaction** to include all CICS transactions in the filter criteria. The list of CICS transactions is generated from events collected in the project database.
- To manually add a new transaction to the list, enter the name and click **Add**. To remove a transaction from the list, check the box next to it and click **Delete**. Only transactions that were manually added and do not exist in the project may be deleted.

- Only events generated on the z/OS CICS platform have a CICS transaction. To collect events that do not have a CICS transaction (non-z/OS CICS events), choose **Collect events regardless of whether the field "CICS Transaction" is available**.
- Choose **Collect events if and only if the field "CICS Transaction" is available** to collect events that only have CICS transactions.

z/OS IMS Region Type

- Determine whether to **Include** or **Exclude** events for selected IMS region types.
- Check the box next to each IMS region type you want to include specifically in the filter criteria, or check the box next to **All IMS Region Type** to include all IMS region types in the filter criteria. The list of IMS region types is generated from events collected in the project database.
- To manually add a new region type to the list, enter the name and click **Add**. To remove a region type from the list, check the box next to it and click **Delete**. Only region types that were manually added and do not exist in the project may be deleted.
- Only events generated on the z/OS IMS platform have an IMS region type. To collect events that do not have an IMS region type (non-z/OS IMS events), choose **Collect events regardless of whether the field "IMS Region Type" is available**.
- Choose **Collect events if and only if the field "IMS Region Type" is available** to collect events that only have IMS region types.

z/OS IMS Region Identifier

- Determine whether to **Include** or **Exclude** events for selected IMS region identifiers.
- Check the box next to each IMS region identifier you want to include specifically in the filter criteria, or check the box next to **All IMS Region Identifier** to include all IMS region identifiers in the filter criteria. The list of IMS region identifiers is generated from events collected in the project database.

- To manually add a new region identifier to the list, enter the name and click **Add**. To remove a region identifier from the list, check the box next to it and click **Delete**. Only region identifiers that were manually added and do not exist in the project may be deleted.
- Only events generated on the z/OS IMS platform have an IMS region identifier. To collect events that do not have an IMS region identifier (non-z/OS IMS events), choose **Collect events regardless of whether the field "IMS Region Identifier" is available**.
- Choose **Collect events if and only if the field "IMS Region Identifier" is available** to collect events that only have IMS region identifiers.

z/O IMS Identifier

- Determine whether to **Include** or **Exclude** events for selected IMS identifiers.
- Check the box next to each IMS identifier you want to include specifically in the filter criteria, or check the box next to **All IMS Identifier** to include all IMS identifiers in the filter criteria. The list of IMS identifiers is generated from events collected in the project database.
- To manually add a new identifier to the list, enter the name and click **Add**. To remove an identifier from the list, check the box next to it and click **Delete**. Only identifiers that were manually added and do not exist in the project may be deleted.
- Only events generated on the z/OS IMS platform have an IMS identifier. To collect events that do not have an IMS identifier (non-z/OS IMS events), choose **Collect events regardless of whether the field "IMS Identifier" is available**.
- Choose **Collect events if and only if the field "IMS Identifier" is available** to collect events that only have IMS identifiers.

z/OS IMS Transaction

- Determine whether to **Include** or **Exclude** events for selected IMS transactions.
- Check the box next to each IMS transaction you want to include specifically in the filter criteria, or check the box next to **All IMS Transaction** to include all IMS transactions in the filter criteria. The list of IMS transactions is generated from events collected in the project database.

- To manually add a new transaction to the list, enter the name and click Add. To remove a transaction from the list, check the box next to it and click Delete. Only transactions that were manually added and do not exist in the project may be deleted.
- Only events generated on the z/OS IMS platform have an IMS transaction. To collect events that do not have an IMS transaction (non-z/OS IMS events), choose **Collect events regardless of whether the field "IMS transaction" is available**.
- Choose **Collect events if and only if the field "IMS Transaction" is available** to collect events that only have IMS transactions.

z/OS IMS PSB

- Determine whether to **Include** or **Exclude** events for selected IMS PSBs.
- Check the box next to each IMS PSB you want to include specifically in the filter criteria, or check the box next to **All IMS PSB** to include all IMS PSBs in the filter criteria. The list of IMS PSBs is generated from events collected in the project database.
- To manually add a new PSB to the list, enter the name and click **Add**. To remove a PSB from the list, check the box next to it and click **Delete**. Only PSBs that were manually added and do not exist in the project may be deleted.
- Only events generated on the z/OS IMS platform have an IMS PSB. To collect events that do not have an IMS PSB (non-z/OS IMS events), choose **Collect events regardless of whether the field "IMS PSB" is available**.
- Choose **Collect events if and only if the field "IMS PSB" is available** to collect events that only have IMS PSBs.

i5/OS Job Name

- Determine whether to **Include** or **Exclude** events for selected job names.
- Check the box next to each job name you want to include specifically in the filter criteria, or check the box next to **All Job Name** to include all job names in the filter criteria. The list of job names is generated from events collected in the project database.
- To manually add a new job name to the list, enter the name and click **Add**. To remove a job name from the list, check the box next to it and click **Delete**. Only job names that were manually added and do not exist in the project may be deleted.

BTTRACE API

- Determine whether to **Include** or **Exclude** events for selected BTTRACE APIs. Applications may use the BTTRACE function provided by the Sensor library to send user-defined trace messages to TransactionVision.
- Check the box next to each API you want to include specifically in the filter criteria, or check the box next to **All BTTRACE API** to include all BTTRACE APIs in the filter criteria.

BTTRACE Severity

- Determine whether to **Include** or **Exclude** events for selected BTTRACE severity levels. Applications may use the BTTRACE function provided by the Sensor library to send user-defined trace messages to TransactionVision.
- Check the box next to each severity you want to include specifically in the filter criteria, or check the box next to **All Severity** to include all BTTRACE severity levels in the filter criteria.

Create a New Data Collection Filter

This task includes the following steps:

- “Prerequisites” on page 258
- “Open the Add Filter Page” on page 258
- “Enter the Data Collection Filter Name” on page 258
- “Select One or More Technologies to Include in Your Filter” on page 258
- “Specify the Collection Criteria” on page 259
- “Change the Default Values for Filter Conditions You Want to Change for One or More Technologies (Sensors)” on page 260
- “Save Your Settings for the Current Category” on page 260
- “Click Finish to Save All Filter Criteria Settings” on page 260

1 Prerequisites

- The Analyzer has been installed and configured. See “Analyzers” on page 71 for details.
- A project has been selected. See “Projects” on page 43 for details.

2 Open the Add Filter Page

Click **New Filter** on the Project Data Collection filters page.

3 Enter the Data Collection Filter Name

This name is used on the Project List, Project Status, and Edit Project pages.

4 Select One or More Technologies to Include in Your Filter

Check each technology you want to include in your filter, or check All to include events for all technologies.

For events from .NET, BEA Tuxedo, or NonStop TMF Sensors, specify the User Event technology.

5 Specify the Collection Criteria

Click **Next** to specify the collection criteria on the Edit Data Collection Filter page.

» Edit Data Collection Filter

Project: Trade
Filter: Collect All

☒ Criteria for

All Agents

Common Option

All Agents

Host

CICS Agent

Program/Servlet

EJB Agent

Time

JDBC Agent

User

JMS Agent

Data Range

Servlet Agent

Event Packaging

User Event Agents

WebSphere MQ Agent

Criterion Setting for Program/Servlet/EJB

Include

qualified events where the Program/Servlet/EJB is equal to

☒ All Program/Servlet/EJB

☐ /bond.jsp

☐ /fx.jsp

☐ /main.jsp

☐ /simulation.jsp

☐ /status.jsp

☐ /stock.jsp

☐ AppServerTestServletComponent

☐ AppServerXAPTTestServletComponent

☐ AppServerXAPubSubTestServletComponent

☐ AppServerXATestServletComponent

☐ BondTradeProcessor

☐ ...

Edit Program/Servlet/EJB

Name:

Add

Delete

Summary

Filter criteria for CICS Agent

Collect All CICS events

Filter criteria for EJB Agent

Collect All EJB events

Filter criteria for JDBC Agent

Finish

Cancel

Set Criteria

Clear Criteria

Reset All Criteria

The left side of the page lists all categories of filter criteria that you can change. The right side shows the settings for the selected criteria. Sensors only collect and send events that meet their respective filter criteria.

6 Change the Default Values for Filter Conditions You Want to Change for One or More Technologies (Sensors)

Click the criteria category name on the left side for the condition that you wish to change and make desired changes on the right side of the page. See “Filter Conditions” on page 219 for information on setting criteria for each condition.

To specify filter criteria, perform the following steps:

- a Select the Technology (Sensor) for which you wish to specify filter criteria from the drop-down list. You may specify criteria that apply to all technologies, or you may specify criteria for a specific technology.
- b Check the box to the left of the Sensor to include events for the selected technology (Sensor) in your filter results, or clear the box to exclude events for the selected technology (Sensor) from your filter results.

7 Save Your Settings for the Current Category

Click **Set Criteria** to save your settings. To clear changes for the current category, click **Clear Criteria**. To reset all categories to the default criteria, click **Reset All Criteria**.

8 Click Finish to Save All Filter Criteria Settings

The new filter appears in the Filter List on the Project Data Collection Filters page.

Edit a Data Collection Filter

This task includes the following steps:

- “Open the Edit Data Collection Filter Page” on page 261
- “Edit the Criteria Category(s) and Filter Conditions You Wish to Change” on page 261
- “Save Your Settings for the Current Category” on page 261
- “Click Finish to Save All Filter Criteria Settings” on page 261

1 Open the Edit Data Collection Filter Page

On the Project Data Collection Filters page, select the filter you wish to edit and click **Edit Filter** (shown in step 5 of “Create a New Data Collection Filter” on page 258).

2 Edit the Criteria Category(s) and Filter Conditions You Wish to Change

Click **New Filter** on the Project Data Collection filters page.

3 Save Your Settings for the Current Category

Click **Set Criteria** to save your settings. To clear changes for the current category, click **Clear Criteria**. To reset all categories to the default criteria, click **Reset All Criteria**.

4 Click Finish to Save All Filter Criteria Settings

The edited filter settings appear on the Project Data Collection Filters page.

Data Collection Filters User Interface

This section describes:

- Add Filter Page on page 262
- Edit Data Collection Filter Page on page 263
- Project Data Collection Filters Page on page 265

Add Filter Page

Description	Allows you to create a new data collection filter for a project. To access: On the Administration > TransactionVision > Current Projects > Data Collection Filter > Project Data Collection Filters page, click New Filter
Included in Tasks	➤ “Create a New Data Collection Filter” on page 258
Useful Links	➤ “Filter Conditions” on page 219 ➤ “Projects” on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to return to the Project Data Collection Filter page without creating a new filter.
Next >	Click to specify criteria for the new filter. Displays the “Edit a Data Collection Filter” on page 260.
Filter Name	Enter a unique name for the data collection filter. This name is used on the Project List, Project Status, and Edit Project pages.
Technology	Check each technology you want to include in your filter, or check All to include events for all technologies. NOTE: For events from .NET, BEA Tuxedo, or NonStop TMF Sensors, specify the User Event technology.

Edit Data Collection Filter Page

Description	<p>Allows you to specify collection criteria.</p> <p>To access:</p> <ul style="list-style-type: none"> ➤ On the Administration > TransactionVision > Current Projects > Data Collection Filter > Project Data Collection Filters page, click Edit Filter ➤ On the Add Filter page, click Next >
Important Information	<p>The left side of the page lists all categories of filter criteria that you can change. The right side shows the setting for the selected criteria.</p> <p>Click the criteria category name on the left side for the condition that you wish to change from the default value and make desired changes on the right side of the page. See “Filter Conditions” on page 219 for details on setting criteria for each condition.</p>
Included in Tasks	<ul style="list-style-type: none"> ➤ “Create a New Data Collection Filter” on page 258 ➤ “Edit a Data Collection Filter” on page 260
Useful Links	<ul style="list-style-type: none"> ➤ “Filter Conditions” on page 219 ➤ “Projects” on page 43

Edit Data Collection Filter Buttons

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to return to the previous filter page without modifying the default filter criteria.
Clear Criteria	Click to clear changes for the current category.
Finish	Click when all filter criteria are set.
Reset All Criteria	Click to reset all categories to the default criteria.
Set Criteria	Click to save your settings for the current category.

Edit Data Collection Filter Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Common Options	<p>Criteria to filter events for all technologies. See “Common Options” on page 219.</p> <ul style="list-style-type: none"> ➤ “Host” on page 220 ➤ “Program/Servlet/EJB” on page 220 ➤ “Time” on page 221 ➤ “User” on page 222 ➤ “Data Range” on page 222 ➤ “Event Packaging” on page 223
Criteria for drop-down	<p>Select the Technology (Sensor) for which you wish to specify filter criteria from the drop-down list. You can specify criteria from the Common Options that apply to all technologies (All Sensors), or you may specify criteria for a specific technology:</p> <ul style="list-style-type: none"> ➤ All Sensors, see “Common Options” on page 219 ➤ CICS Sensor, see “CICS Sensor Criteria” on page 225 ➤ EJB Sensor, see “EJB Sensor Criteria” on page 229 ➤ JDBC Sensor, see “JDBC Sensor Criteria” on page 231 ➤ JMS Sensor, see “JMS Sensor Criteria” on page 235 ➤ Servlet Sensor, see “Servlet Sensor Criteria” on page 237 ➤ User Event Sensors, see “User Event Sensor Criteria” on page 239 ➤ WebSphere MQ Sensor, see “WebSphere MQ Sensor Criteria” on page 248
Criterion Setting for <criteria>	Dialog box for setting the filter criterion.
Summary	Summary of the data collection filter criteria that have been set.

 **Project Data Collection Filters Page**

Description	Displays a list of data collection filters available for the current project. To access: Select Administration > TransactionVision > Current Projects > Data Collection Filters
Included in Tasks	<ul style="list-style-type: none">➤ “Create a New Data Collection Filter” on page 258➤ “Edit a Data Collection Filter” on page 260
Useful Links	<ul style="list-style-type: none">➤ “Filter Conditions” on page 219➤ “Projects” on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Delete Filter	Click to delete the selected filter.
Edit Filter	Click to modify a selected filter’s criteria. Displays the “Edit a Data Collection Filter” on page 260.
Filter Detail	Displays the selected filter’s criteria settings.
Filter List	Displays existing data collection filters.
New Filter	Click to create a new data collection filter. Displays the “Add Filter Page” on page 262.

8

Jobs

This chapter includes:

Concepts

- Jobs Overview on page 267

Tasks

- Create a New Job – Workflow on page 272
- Make a Job Accessible in the UI/Job Server on page 274

Reference

- Jobs User Interface on page 275

Jobs Overview

A job is a bean that implements a particular task to run 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 complex, time-consuming queries to present report data. Instead, they use already calculated data that is periodically updated by a job running in the background.

A job instance is defined for a particular project; that instance can only be run on that project. However, multiple instances of the job bean may be run concurrently, each configured for a different project.

Only users with administrator permissions may manage jobs. To do so go to **Admin > TransactionVision > Administration > Jobs**. Users may see job status for any projects to which they have view permissions. To see job status for a project, go to **Administration > TransactionVision > Current Projects > Jobs**.

For information about creating job beans, see the *HP TransactionVision Advanced Customization Guide* PDF.

TransactionVision provides the following jobs:

deleteEvents Job

The deleteEvents job deletes all events more than 48 hours old from your project database. This job is optional and configurable. By default it is turned off when you create a project. The deleteEvents job uses the deleteEvents utility; for detailed information about this utility, see “DeleteEvents” on page 392.

Note: TransactionVision does not support running multiple instances of this job simultaneously. Be sure to schedule this job so that it does not run for multiple projects at the same time.

Transaction Statistics Job

The Transaction Statistics job calculates business transaction statistics. This job is required if you wish to collect transaction statistics; these statistics are used by a number of different reports. If this job is not running, those reports do not show any data.

Note: TransactionVision does not support running multiple instances of this job simultaneously. Be sure to schedule this job so that it does not run for multiple projects at the same time.

RunStats Job

The RunStats job updates statistics about the physical characteristics of a project's schema table, and the associated indexes. These characteristics include number of records, number of pages, and average record length. The database optimizer uses these statistics when determining access paths to the data.

This job should be run when a table has had many updates, such as when data is continuously collected into the database by the TransactionVision Analyzer. It could result in large performance gains in queries made by TransactionVision views and reports, as well as queries made internally by the Transaction Vision Analyzer to correlate events. However the RunStats job is very resource intensive. Therefore, schedule it to run at a time when its impact on overall performance can best be tolerated.

IBM DB2 v9 has a built-in job scheduler to automatically collect statistics. You do not need to use the RunStats job if you are using DB2 v9.

If you use an Oracle database, the RunStats job uses the OracleRunstats script provided with Transaction Vision. For more information about this script, see the *HP TransactionVision Deployment Guide* PDF. Note that Oracle 10g has a built-in job scheduler to automatically collect statistics. You do not need to use the RunStats job if you are using Oracle 10g.

Parameter	Description
database (d)	Database name on which to perform the job.
user (u)	User name for accessing the database.
trace (t)	Log trace information.

CMDB Population Job

The CMDB Population Job publishes CIs corresponding to the infrastructure elements discovered by TransactionVision sensors, to the uCMDB. This job is set to run by default every 15 minutes. During each run, this job queries for Business Transaction Classes and the infrastructure elements they are dependant on, from the TransactionVision database. It then publishes any infrastructure dependencies of the business transactions to the uCMDB. This job currently publishes to the uCMDB, J2EE applications, J2EE servers on which the J2EE applications reside, and any machines on which J2EE servers are hosted. The job also creates a dependency relationship between business transactions and corresponding J2EE applications. These dependencies can then be viewed from the Business Transactions view in the Business Availability Center dashboard.

Parameter	Description
-debug	Set this parameter to true to enable debug statements.
-domain <domain>	Set this parameter to set the corresponding value as the probe_domain attribute for populating hosts in uCMDB (default value is DefaultDomain).

Note: If the -domain parameter is used, it should be set to the same value as the CMDB Update job.

CMDB Update Job

The CMDB Update Job updates infrastructure CIs corresponding to business transactions in the uCMDB, to prevent the aging out of these CIs.

Parameter	Description
-debug	Set this parameter to true to enable debug statements.
domain <domain>	Set this parameter to set the corresponding value as the probe_domain attribute for populating hosts in uCMDB (default value is DefaultDomain).

Note: If the -domain parameter is used, it should be set to the same value as the CMDB Population job.

TransactionVision CIs Deleted From the uCMDB

If the CIs populated by TransactionVision are deleted from the uCMDB (for example if the uCMDB database is re-created), then the following steps must be performed.

To re-create the CIs in the uCMDB:

- 1 Invoke the **Sync CMDB** option under **Admin > Business Availability Settings**, to create the business transaction CIs.
- 2 Run the CMDB Update Job, using the **Run Now** option.
- 3 Run the CMDB Population job to force the re-creation of the CIs in the uCMDB.

BAC Aggregate TV Sample Data

Business transaction samples are sent to Business Availability Center through a TransactionVision scheduled job running in the TransactionVision UI/Job server. BAC Aggregate TV Sample Data job creates and delivers aggregate business transaction data to Business Availability Center. It queries the business transaction table, and delivers aggregated samples of monitored transactions every 5 minutes (corresponding to the sample interval). The sample then delivers pre-defined metrics to Business Availability Center via HTTP Post.

For more information on this job, see “BAC Aggregate TV Sample Data” on page 364.

Create a New Job – Workflow

This task describes how to create a new job for a project.

This task includes the following steps:

- “Prerequisites” on page 272
- “Navigate to the Job Status page” on page 273
- “Enter a job name and description.” on page 273
- “Select a startup mode:” on page 273
- “Enter the name of the java class in the Class Name field.” on page 273
- “Enter any startup parameters.” on page 273.
- “Enter the desired job frequency.” on page 273.
- “Enter a next scheduled job time.” on page 273
- “Select projects for this job.” on page 274

1 Prerequisites

Make the job accessible to the UI/Job Server. See “Make a Job Accessible in the UI/Job Server” on page 274.

2 Navigate to the **Job Status** page

Go to **Admin > TransactionVision > Current Projects > Jobs**, click **New** The **Create Job** page displays.

3 Enter a job name and description.**4** Select a startup mode:

- **Automatic.** The job is automatically initialized and started when the web server starts or when a user clicks Start All.
- **Manual.** The job is only started when the job is selected and started with the Start Job or Run Now button.
- **Disabled.** The job cannot be run until its startup mode is changed to Automatic or Manual. Disabling a job is useful if you have an automatically starting job that you want to temporarily turn off so it doesn't start when the system comes up or when starting all jobs.

5 Enter the name of the java class in the **Class Name** field.

This class must implement the UI/Job interface (see the *HP TransactionVision Advanced Customization Guide* PDF for details).

6 Enter any startup parameters.

These parameters are passed to the job bean's initialization method when it starts. For the **deleteEvents** job, keep the following in mind:

- By default, the value for the **-commit** option for the **deleteEvents** job is 0, meaning that the commit is done after the entire deletion. If you are deleting a large number of events, explicitly specify the parameter **-commit n** to commit after **n** event deletions. Otherwise, the database transaction log may overflow or the table may be locked for a long period of time.
- By default, the **deleteEvents** job and script do not allow deleting events from the last 48 hours. To override the default of 48 hours, use the **-minage** number-of-hours option. For example, **-minage 10** prevents the job from deleting events from the last 10 hours.

7 Enter the desired job frequency.

This interval may be in minutes, hours, days, or months.

8 Enter a next scheduled job time.

If you do not enter a next scheduled time, the next scheduled time is calculated based on the information in the **Repeat Every** field.

- 9 Select projects for this job.

If you want to propagate this job to any existing projects, select the projects.

Click **Finish** and return to the **Job Status** page.



Make a Job Accessible in the UI/Job Server

This task describes how to make a job accessible in the UI/Job Server. The UI/Job Server includes the Tomcat servlet container. To make a job accessible, you add its classes to the Tomcat WAR file.

To make a Job accessible in the UI/Job Server:

- 1 Stop the Tomcat server.

See the the *HP TransactionVision Deployment Guide* PDF for information on stopping Tomcat.

- 2 Make a backup copy of the **tvision.war** file in **<TVISION_HOME>/apache-tomcat/webapps**.
- 3 Copy the **tvision.war** file to a temporary location and extract it using the following command:

```
jar -xf tvision.war
```

This command extracts all TransactionVision related files. Among the files are two directories: **WEB-INF/lib** and **WEB-INF/classes**.

- 4 Copy any required jar files to the **WEB-INF/lib** directory and any required class files to the **WEB-INF/classes** directory.

Note: If class files are part of a package, they must be placed in that corresponding directory. For example, a Java class named `com.mycompany.TestClass` would need to be placed in the `com/mycompany` directory within **WEB-INF/classes**.

- 5 Repack the **tvision.war** file. First, remove the old **tvision.war** file from the directory and then recreate it, packing all the files in the directory by using the following command:

```
jar -cf tvision.war
```

- 6 Copy this new EAR file to **<TVISION_HOME>/ apache-tomcat/webapps**.
- 7 Delete the following directory: **<TVISION_HOME>/ apache-tomcat/webapps/tv**.
- 8 Restart Tomcat.

See the the *HP TransactionVision Deployment Guide* PDF for information on restarting Tomcat.

See “Create a New Job – Workflow” on page 272 to add the job to the TransactionVision user interface.

Jobs User Interface

This section describes:

- Jobs Page on page 276
- Jobs for All Projects Page on page 279
- New Job Page on page 280
- Edit a Job Page on page 281

 **Jobs Page**

Description	Enables you to view jobs for the current project, start or stop jobs, run waiting jobs, create new jobs, and delete jobs. To access: Admin > TransactionVision > Current Projects > Jobs
Important Information	To view jobs for all projects, see Admin > TransactionVision > Administration > Jobs.
Included in Tasks	“Create a New Job – Workflow” on page 272
Useful Links	“Jobs for All Projects Page” on page 279

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Delete	To delete a job, click the radio button to the left of the job name to select it and click Delete . You may only delete a job if its status is stopped.
Edit	To modify a job configuration, click Edit Job to open the Edit Job page. For more information, see “Edit a Job Page” on page 281.
Force Stop	To force a job in the Quiescing state to stop immediately without completing pending cleanup or shutdown tasks, click the radio button to the left of the job name to select it, then click Force Stop . This button is only enabled if the selected job is in Quiescing mode.
New	To create a new job configuration, click New to open the Create Job page. For more information, see “New Job Page” on page 280.

GUI Element	Description
Run Now/Cancel Run	To force a waiting job to run immediately rather than continuing to wait for the scheduled time, click the radio button to the left of the job name to select it and click Run Now . To cancel the current run of a running job without stopping the job, select it and click Cancel Run . The job changes to a waiting state until the next scheduled run time.
Start All Jobs	To start all automatic jobs. Starting a job initializes and activates it. If it's next scheduled time has already passed, it runs immediately; otherwise, it is placed in a waiting state until its scheduled time arrives.
Start/Stop Job	To start or stop a single job, click the radio button to the left of the job name to select it, then click Start Job to start a stopped job or Stop Job to stop a running or waiting job.
Stop All Jobs	To shut down all currently active jobs. If the job is currently running, a cancel command is first issued, followed by commands to stop and shut it down.
<History Info>	For the selected job, the Job Status page also displays history information showing when the job has been initialized, run, and completed.
Project	Current active project.
Database Schema	The assigned database schema.
Active Jobs in Current Project	The number of active jobs there are in the current project.
Job Name	Click the job name link to edit an existing job configuration.
Job Description (optional)	The description given for the job.

GUI Element	Description
Job Startup Mode	<p>Automatic: the job is automatically initialized and started when the web server starts, or when a user clicks Start All.</p> <p>Manual: the job is only started when the job is selected and started</p> <p>Disabled: the job cannot be run until its startup mode is changed to Automatic or Manual. Disabling a job is useful if you have an automatically starting job that you want to temporarily turn off so it doesn't start when the system comes up or when starting all jobs.</p>
Job Interval	Specifies the time interval at which the job is configured to run.
Job Status	<p>Running: the job is in the process of executing the task it is programmed to do.</p> <p>Waiting: the job is active and waiting for its scheduled run time to arrive.</p> <p>Quiescing: the job is in the process of shutting down, performing any specific cleanup or shutdown processes applicable. A job initially transfers to this state when it is stopped, then automatically moves from this state to Stopped when the shutdown or cleanup process is complete. For example, the DeleteEvents job finishes making database updates for its current progress in this state. When in Quiescing mode, use the Force Stop button to end a job immediately without completing pending tasks.</p> <p>Stopped: the job is not active, is not loaded into memory, and does not perform any tasks until started with the Start Job or Start All Jobs button.</p>
Time of Next Activation	Specifies the time a waiting or running job next runs.



Jobs for All Projects Page

Description	Displays all jobs for all projects. To access: Admin > TransactionVision > Administration > Jobs
Important Information	To view job status for the current project only, see “Jobs Page” on page 276.
Useful Links	After selecting the project, see “Jobs Page” on page 276, for details on the jobs in this project.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Project Name	Select the project for which you wish to see the jobs.
Active Jobs	The number of active jobs in the selected project.



New Job Page

Description	Enables you to create a new Job. To access: Admin > TransactionVision > Current Projects > Jobs, click New button.
Included in Tasks	“Create a New Job – Workflow” on page 272.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Class Name	Enter the name of the java class, this class must implement the UI/Job interface.
Description	Enter a description for this new job.
Job Name	Enter a name for the new job.
Next Scheduled Time (Optional)	Enter in the following format: MM/dd/yyyy HH:mm.
Projects	Select the projects that you wish to create an instance of this job in.
Repeat Every	Enter the frequency with which the job should be repeated. The interval may be in minutes, hours, days, or months.
Startup Mode	Select Automatic , Manual , or Disabled .
Startup Parameters (Optional)	Enter any startup parameters needed for this job.

Edit a Job Page

Description	<p>Enables you to edit job parameters.</p> <p>To access: Admin > TransactionVision > Current Projects > Jobs, open a Job then click Edit</p>
--------------------	--

The following elements of the job are displayed (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Class Name	Can be edited, see “New Job Page” on page 280 for field details.
Job Description (Optional)	Can be edited, see “New Job Page” on page 280 for field details.
Job Name	Displays the Job name.
Job Usage	Displays the description given for the Job.
Next Scheduled Time (Optional)	Can be edited, see “New Job Page” on page 280 for field details.
Repeat Every	Can be edited, see “New Job Page” on page 280 for field details.
Startup Mode	Can be edited, see “New Job Page” on page 280 for field details.
Startup Parameters (Optional)	Can be edited, see “New Job Page” on page 280 for field details.

9

Transaction Definition Editor

This chapter includes the main concepts, tasks, and reference information for classifying transactions using the Transaction Definition Editor.

This chapter includes:

Concepts

- Transaction Definition Editor Overview on page 284

Reference

- Transaction Definition Editor User Interface on page 288

Troubleshooting and Limitations on page 309

Transaction Definition Editor Overview

TransactionVision provides a structured editor to assist managing transaction classes, rules, and their associated attributes. Defining transaction rules using the editor maps events to transaction classes, defines how data is mapped between Transaction Vision events and its business transactions, and defines any special actions to take when particular criteria are met.

Committing Changes

When editing components within the user interface, changes are being made to a temporary copy of the rules. Until you click **Finish** on the main page, the changes are not persisted to the database, nor is the analyzer notified of any changes. When you click **Finish**, any changes made in the current session are written to the database. The analyzer is then instructed to temporarily pause collection, while it reconfigures itself and recreates any database tables that are impacted by the changes made. If you logout before committing your changes, they are lost. Clicking **Cancel** on the main page also discards all current modifications, and reloads the transaction rules from the existing settings.

Transaction Rules

Transaction rules are evaluated in the order in which the rules appear, from top to bottom. When the analyzer evaluates rules, it stops evaluation and performs a classification on an event on the first set of conditions that completely match the given rules. Keep this in mind when designing your classification rules.

If a schema contains both common sections that are assigned directly to a schema along with common sections that are assigned to **All Schemas**, the object assigned specifically to the one schema has precedence and is evaluated first. It is not possible for an object assigned to **All Schemas** to be evaluated first, before any rules specific to that schema.

Classify/Evaluate Sections

A *Classify object* used by a transaction class, and an *Evaluate object* used by a common section, are very similar. They both comprise of:

- a Match Condition
- attribute rule/s
- action/s to perform when a Match Condition evaluates to true.

The additional significance of the **Classify** section is that if its Match Condition evaluates to true, the event then becomes classified as that transaction class. See the “Classify Section Page” on page 295 for field details when creating a new section.

Attribute Definitions

Attributes define:

- a mapping from data in the event document to data in the transaction document - OR -
- the setting of transaction document values to constant values.

Each attribute definition can also have Match Conditions and actions that control when and if it is processed.

Attribute Rules defined within a classify section only apply if the Match Condition to the classify section match, see “New Attribute Page” on page 300.

Attributes defined within the main body of the transaction class apply to any event classified as belonging to the transaction class, see “Transaction Class Attributes Page” on page 302.

Actions

Actions provide a mechanism to perform custom operations when certain rules with your classification system are fired. These actions represent java objects that the analyzer framework loads and calls into when fired.

A BPI action bean is provided, for sending TransactionVision events to BPI. Other action beans can be created.

Match Conditions

Match Conditions specify a set of conditions to evaluate, that if true, allow further execution of the rules active in this section. A Match Condition contains one or more stanzas that give an XML Path (XPath) into the event documents, a conditional operator (equals, greater than, etc), and a value to compare against. Values in Match criteria may contain one wildcard character, for example, *FlowEngine, DataFlow*, or Data*Engine.

A Match Condition is a statement that evaluates to true or false, indicating whether a set of rules should be implemented or not. The request is: when the data in my event matches a particular criteria, then perform the specified actions. A Match Condition consists of:

- an XPath
- a conditional operator
- a value.

An XPath is the named location of a piece of data, within the event payload that TransactionVision collects from its sensors. Each event the sensor sends is an XML document, and the XPath indicates which node in that document you wish to refer to.

A number of **conditional operators** are available for selection from the drop-down list under the XPath field, such as equal, greater than, etc. A **Value** is what you wish to compare against.

See “Edit Match Condition Page” on page 298, for details on creating and editing Match Conditions

Display of XPaths

By default a set of fixed XPaths is presented, which comprises of all the standard locations of TransactionVision event data. If your event data contains custom fields, such as XML from the payload data or data fields specific to a monitored application, TransactionVision does not know such permutations and these XPaths do not display in the Edit Match Condition page.

To address this, populate the tree control in the following way:

- a** Select a transaction, choose the project, and then click **Select**.
- b** The Transaction Tracking report opens, see “Transaction Tracking Report” on page 562 for report details.
- c** A prompt asks you to select from the list of available transactions, the tracking report then displays the Edit Match Condition page, and your selected transaction is saved as belonging to this transaction class. It can be changed to a different transaction class later.
- d** Next time the Edit Match Condition page is loaded, it first looks up your specified transaction and parses from all the events within that transaction all XPaths used, including any custom payload data Xpaths, and displays them in the tree control.

Additional benefits of this method are

- only XPaths relevant to the selected transaction are shown
- all the values from the selected events are populated into the second window, which allows you to select directly from the values instead of going to the event detail view to look up what types of values this xpath may contain.

If you wish to look at this transaction, a link is provided to directly open the event list of this transaction.

Edit Match Conditions

A Match Condition is a statement that evaluates whether a set of rules should be implemented or not. It consists of an XPath, a conditional operator and a value. See “Match Conditions” on page 286 for more details.

An XPath is the named location of a piece of data, within the event payload that TransactionVision collects from its sensors. Each event the sensor sends is an XML document, and the XPath indicates which node in that document you wish to refer to.

Remembering the full XPath selector for data in an event is not easy. This page has a number of features to help determine the XPath, and select the desired value.

Transaction Definition Editor User Interface

This section describes:

- Main Transaction Definition Editor Page on page 289
- Transaction Class Definition Page on page 292
- New Common Section Page on page 294
- Classify Section Page on page 295
- Evaluate Section Page on page 297
- Edit Match Condition Page on page 298
- New Attribute Page on page 300
- Edit Action Properties Page on page 301
- Transaction Class Attributes Page on page 302
- New Transaction Class Attribute Page on page 304
- New BPI Action Page on page 305
- Edit Value Properties Page on page 306
- Edit Rule Page on page 308

Main Transaction Definition Editor Page

Description	<p>Enables you to manage transaction rules and associated attributes.</p> <p>To access: Select Admin > TransactionVision > Transaction Definition Editor.</p>
Important Information	<ul style="list-style-type: none"> ➤ Be sure to click Finish to commit changes, see “Committing Changes” on page 284. ➤ Whenever a transaction class is updated in TransactionVision, the update is propagated to the corresponding CIs in the CMDB. Creating a new transaction class triggers the automatic creation of the CIs for the new business transaction in the CMDB. Deleting a transaction class causes the deletion of the corresponding business transaction CIs. ➤ Common section objects assigned to a specific schema have precedence over those assigned to all schemas, see “Common Sections” on page 291.

Transaction Definitions for Schemas combo box

The transaction editor page opens with the currently active schema selected in the combo box, and all the objects for that schema displayed.

GUI Element	Description
Synchronize CMDB	<p>Click to synchronize transaction classes with the CMDB. A results screen lets you know the status.</p> <p>In the normal course of operation using this button is not needed. The synchronization occurs automatically whenever you commit new changes. If TransactionVision and CMDB have become out of sync, for example due to the CMDB being unreachable when an update occurs, this button provides a mechanism to perform a synchronization.</p>
Transaction Definitions for Schema	<p>Select a schema from the drop-down list, or select All Schemas, Not Assigned, or All Classes.</p>

GUI Element	Description
Finish	Click to save the configuration.
Cancel	Click to cancel the configuration.

Transaction Class Definitions

A transaction class definition is a set of specific rules for defining data for an individual transaction class.

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Copy	To make a duplicate copy, select an object and click Copy .
Delete	Select an object and click Delete .
Move Down	Select a transaction class and click Move Down to change the order it is evaluated and displayed.
Move Up	Select a transaction class and click Move Up to change the order it is evaluated and displayed.
New	Click New to create a new object.
Summary	Select an object and click Summary to open a summary page in a new window. This enables viewing the whole transaction class while you drill down into a component for further editing.
<Collapse/Expand Tree>	Click the + icon next to each object to expand and show additional levels of detail.
<html links>	Use the html links as shortcuts to jump directly to editing pages
Description	Displays the description assigned to the transaction class.
Schema	The schema to which the transaction class is assigned.
Transaction Class	Transaction classes for this schema are listed. Click on the + icon to display details of a transaction class, or select the transaction class and click one of the buttons.

Common Sections

Common sections are rules that apply across all events, regardless of how they are classified.

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Copy	To make a duplicate copy, select an object and click Copy .
Delete	Select an object and click Delete .
Move Down	Select a transaction class and click Move Down to change the order it is evaluated and displayed.
Move Up	Select a transaction class and click Move Up to change the order it is evaluated and displayed.
New	Click New to create a new object.
Summary	Select an object and click Summary to open a summary page in a new window. This enables viewing the whole transaction class while you drill down into a component for further editing.
<Collapse/Expand Tree>	Click the + icon next to each object to expand and show additional levels of detail.
<html links>	Use the html links as shortcuts to jump directly to editing pages
Common Sections	The common section name.
Description	The description of the common section.
Schema	A common section can be assigned to one or more, or all schemas. Common section objects assigned to a specific schema have precedence over those assigned to all schemas.



Transaction Class Definition Page

Description	<p>Enables you to set the criteria to assign transactions to transaction classes.</p> <p>A transaction class is a set of rules comprising of classify sections, class attributes and attribute rules that determine the criteria to use to determine if a TransactionVision transaction should be a certain class.</p> <p>To access: Select Admin > TransactionVision > Transaction Definition Editor.</p>
Important Information	<p>If you are importing transaction classes from a previous release of TransactionVision (before 8.0), it is best to manually assign schemas.</p>

General Properties

A transaction class definition is a set of specific rules for defining data for an individual transaction class.

The following elements are included:

GUI Element	Description
View Class Attributes	Click to view the attributes of the selected transaction class.
Assigned Schemas	<p>The schema/s to which this transaction class is assigned.</p> <p>It is recommended that any transaction class is only assigned to one schema. If you have a transaction class that is assigned to multiple schemas this can lead to inconsistent results being reported in your Business Availability Center and TransactionVision reports and their drilldowns into TransactionVision.</p>
Description	The description given to the transaction class.
Transaction Class Name	The transaction class name

Classify Sections

The following elements are included:

GUI Element	Description
Copy	To make a duplicate copy, select an object and click Copy .
Delete	Select an object and click Delete .
Move Down	Select a transaction class and click Move Down to change the order it is evaluated and displayed.
Move Up	Select a transaction class and click Move Up to change the order it is evaluated and displayed.
New	Click New to create a new object.

Attribute Rules

See “New Attribute Page” on page 300 for details on creating new attribute rules. The following elements are included:

GUI Element	Description
New	Click to create a new attribute rule.
Delete	Select an attribute rule and click Delete to delete it.



New Common Section Page

Description	<p>Enables you to create a new common section. Common sections are used to define generic rules and behaviors that are not specific to a certain transaction class.</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor > Common Section > New</p>
Important Information	<ul style="list-style-type: none"> ➤ Common section objects assigned to a specific schema have precedence over those assigned to all schemas. ➤ A common section does not have an attribute within its main body like a transaction class. Nor does it have any class attributes (for example, SLA), since the events it operates on is not necessarily part of a classified transaction.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Copy	To make a duplicate copy, select an object and click Copy .
Delete	Select an object and click Delete .
Move Down	Select a transaction class and click Move Down to change the order it is evaluated and displayed.
Move Up	Select a transaction class and click Move Up to change the order it is evaluated and displayed.
New	Click New to create a new object.
Name	Enter a name for the common section.
Description	Enter a description for the common section.
Assign to Schema	A common section can be assigned to one or more, or all schemas.

Evaluate Sections

A list of any evaluate sections displays. See “Classify Section Page” on page 295, for details on creating new evaluate sections.

Classify Section Page

Description	Enables defining a classify section for a transaction class. To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor , in the Transaction Class section click New , see the Classify Section .
Important Information	If a Classify section's Match Condition evaluates to true, the event becomes classified as belonging to that transaction class.
Useful Links	<ul style="list-style-type: none"> ➤ “Edit Match Condition Page” on page 298 ➤ “New Attribute Page” on page 300 ➤ “Edit Action Properties Page” on page 301

The following elements are included:

GUI Element	Description
Classify Section Name	Enter a name for this new classify section.
Description	Enter a description for this new classify section.
New	Click to create a new Match Condition, or attribute. For field definitions see “Edit Match Condition Page” on page 298, or “New Attribute Page” on page 300.
Delete	Click to delete the selected attribute or Match Condition.

GUI Element	Description
Attributes	Displays the attributes created for this classify section. See “Attribute Definitions” on page 285 for more information.
Match Conditions	Displays the Match Conditions set for this classify section. See “Match Conditions” on page 286 for more information.

Actions

Actions provide a mechanism to perform custom operations when certain rules within your classification system are fired. These actions represent java objects that the analyzer framework loads and calls into when fired.

A BPI action bean is provided for sending TransactionVision events to BPI. Other action beans can be created.

The following elements are included:

GUI Element	Description
Delete	Select the action and click Delete to delete the action.
New	Click to create a new action. See “Edit Action Properties Page” on page 301 for field definitions.
New BPI Action	Click to create a new BPI action. See “Classify Events for Business Process Insight Integration” on page 372 for information on BPI Actions.
Class	The Action Class defined for the action.
Reason	The reason defined for the action.

Evaluate Section Page

Description	<p>Enables defining an Evaluate section for a Common section.</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor, in the Common Section click New, see the Evaluate Section.</p>
Useful Links	<ul style="list-style-type: none"> ➤ “Edit Match Condition Page” on page 298 ➤ “New Attribute Page” on page 300 ➤ “Edit Action Properties Page” on page 301

The following elements are included:

GUI Element	Description
Evaluate Section Name	Enter a name for this new evaluate section.
Description	Enter a description for this new evaluate section.
New	Click to create a new Match Condition, or attribute. For field definitions see “Edit Match Condition Page” on page 298, or “New Attribute Page” on page 300.
Delete	Click to delete the selected attribute or Match Condition.
Attributes	Displays the attributes defined for this evaluate section. See “Attribute Definitions” on page 285 for more information.
Match Conditions	Displays the Match Conditions set for this evaluate section. See “Match Conditions” on page 286 for more information.

Actions

Actions provide a mechanism to perform custom operations when certain rules within your classification system are fired. These actions represent java objects that the analyzer framework loads and calls into when fired.

A BPI action bean is provided for sending TransactionVision events to BPI. Other action beans can be created.

The following elements are included:

GUI Element	Description
Delete	Select the action and click Delete to delete the action.
New	Click to create a new action. See “Edit Action Properties Page” on page 301 for field definitions.
New BPI Action	Click to create a new BPI action. See “Classify Events for Business Process Insight Integration” on page 372 for information on BPI Actions.
Class	The Action Class defined for the action.
Reason	The reason defined for the action.

Edit Match Condition Page

Description	<p>Enables defining under what conditions a set of rules should be implemented.</p> <p>To access:</p> <p>Select Admin > TransactionVision > Administration > Transaction Definition Editor, click New in the Transaction Class section, then click New in the Classify Section, and click New in the Match Conditions section.</p> <p>OR -</p> <p>Select Admin > TransactionVision > Administration > Transaction Definition Editor, click New in the Common Section, then New in the Evaluate Section, and click New in the Match Conditions section.</p>
Important Information	The collapsible tree control is populated by default with known XPath's. All the XPath's may not display, see “Display of XPath's” on page 286 for further explanation.

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Cancel	Click to cancel any changes made.
OK	Click to save changes.
<conditional operators>	<p>Select from the drop-down list: equal to, not equal to, greater than, less than, greater than or equal to, less than or equal to, exists, does not exist, contains or matches regular expression.</p> <p>The request is: when the data in my event matches a particular criteria, then perform the specified actions.</p>
Complex XPath	Select true or false .
Value	Enter what you wish to compare against.
XPath	The path to the data, either entered automatically when the XPath is selected from the known list above, or entered manually.
XPath Data Populated From	<p>Select All Fixed XPaths, or Select a Transaction from project and then select the project from the drop-down list.</p> <p>Or click Select New to open the Transaction Tracking report from which you can make a selection. See “Display of XPaths” on page 286.</p>



New Attribute Page

Description	Enables you to define attribute rules for a transaction class that only apply if the Match Condition to the classify section match. To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor . In the Classify Section of a transaction class, click New . and in the Attribute section click New .
Important Information	See the <i>HP TransactionVision Advanced Customization Guide</i> PDF for more details on attributes and rules.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Apply	Click to apply any additions or edits made.
Cancel	Click to cancel any changes made.
OK	Click to return to the previous page.
Transaction Document Path	Enter the path, or click the breadcrumbs to open the Select Transaction Document Path list and select from the list.
Final	Select true or false .
Precedence	Select true or false .
New Rule	Click to open the Edit Rule page and create a new rule, see “Edit Rule Page” on page 308.
Rules	The rule, if any, created for this attribute.

Edit Action Properties Page

Description	<p>Enables you to define an Action which introduces a specified action when certain conditions occur.</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor > Classify Section > Edit Action Properties</p>
Important Information	<p>Defining the action requires at minimum to specify the full Action Class Name of the Java class that implements the action callback.</p>
Useful Link	<p>Action beans are custom written java beans, however TransactionVision provides one standard action for communicating events to Business Process Insight (BPI). See “New BPI Action Page” on page 305 for details on creating Actions to integrate with Business Process Insight.</p>

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to cancel any changes made.
OK	Click to return to the previous page.
Action Class Name	Specify the full Action Class Name of the Java class that implements the action callback.
Code	Provides a means to communicate the corresponding value to the <code>IAnalyzerAction.action</code> method in this Java class.
Domain	Domain is an optional attribute, currently used to indicate that this action is an action that BPI is interested in.
Reason	Provides a means to communicate the corresponding value to the <code>IAnalyzerAction.action</code> method in this Java class.



Transaction Class Attributes Page

Description	<p>Enables you to set parameters to specify settings common to all instances of a class definition, such as SLA value, cost per transaction, etc.</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor, select a Transaction Class, and click View Class Attribute</p>
Important Information	<p>Transaction Class Attributes define common metrics for all instances of a transaction class.</p>
Useful Links	<ul style="list-style-type: none">➤ See “New Transaction Class Attribute Page” on page 304 for creating a new transaction class attribute.➤ See “New Attribute Page” on page 300 for creating a new common section attribute.➤ See the <i>HP TransactionVision Advanced Customization Guide</i> PDF for more details on attributes and rules.

The following elements are included:

GUI Element	Description
Apply	Click to apply any additions or edits made.
Assign	<p>This button is enabled only when the attribute selected is customizable. Click to assign this attribute to this transaction class.</p> <p>Once changes are applied to the analyzer, the transaction class database table is recreated with a column for this attribute.</p>
Cancel	Click to cancel any changes made.
Delete	Select a Customizable Attribute from the list and click Delete to delete the attribute from the transaction class.
New	Click to create a new transaction class attribute. See “New Transaction Class Attribute Page” on page 304.
OK	Click to return to the previous page.

GUI Element	Description
Unassign	<p>This button is enabled only when the attribute selected is customizable. Click to unassign this attribute from this transaction class.</p> <p>Once changes are applied to the analyzer, the transaction class database table is recreated without this attribute.</p>
Customizable Attributes	These attributes can be deleted, edited, assigned, or unassigned.
Description	The description of the transaction class attribute.
Fixed Attributes	When you create a transaction class by default it is associated with a number of attributes that are required by TransactionVision, such as SLA, cost per transaction and currency code. These attributes cannot be deleted or unassigned, and must have their values set.
Name	The name of the transaction class attribute.
Type	Select from the drop-down a type, for example INTEGER, DOUBLE, CHAR, VARCHAR, BIGINT, DATE, DECIMAL.
Value	One or more Value rule blocks to control under what circumstances which data should be transferred to the transaction attribute.
XPath	The XPath field indicates the name of the attribute and makes it possible that TransactionVision components can refer to it.



New Transaction Class Attribute Page

Description	<p>Enables you to create a new attribute for a transaction class, whereby TransactionVision defines a database column for it.</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor, select a Transaction Class, in the Attributes section click New</p>
Important Information	<p>When a new currency code attribute is created, a default currency code of USD is inserted. Change the code if desired.</p> <p>For more details on transaction class attributes and their uses see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.</p>
Useful Links	<ul style="list-style-type: none">➤ “Transaction Class Attributes Page” on page 302➤ “Transaction Class Definition Page” on page 292

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Extended Attribute	Creating an attribute leads to TransactionVision defining a new database column.This is an optional field that may be used to specify advanced options on this column.
Size	For VARCHAR this is the length of the varchar field, for DECIMAL it turns into two fields – precision and scale, otherwise N/A.
Type	Select from the drop-down a type, for example, INTEGER, DOUBLE, CHAR, VARCHAR, BIGINT, DATE, DECIMAL.
Description	Enter a description for the new transaction class attribute.
Name	Enter a name for the new transaction class attribute.

 **New BPI Action Page**

Description	<p>Enables you to create a new Action for communicating events to Business Process Insight (BPI).</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor, select a Transaction Class, in the Common Section of the Transaction Class Definitions page select an object or click New, then in the Evaluate Section of the Common Section page select an object and click New, finally in the Action section of the Evaluate Section page click New BPI Action.</p> <p>OR -</p> <p>Select Admin > TransactionVision > Administration > Transaction Definition Editor, select a Transaction Class, in the Classify Section select an object or click New, then in the Action section of the Classify Section page click New BPI Action.</p>
Important Information	<p>Create one or more Match Conditions to limit which events trigger a BPI event. If none are defined the BPI action is generated for every event class.</p> <p>BPI actions can occur in a transaction class either in a Common Evaluate Section, or as part of a Classify Section:</p> <ul style="list-style-type: none"> ➤ If the BPI Action is created within a Common Section the action is fired every time the matching event occurs. ➤ If the BPI Action is created within a Classify Section the action fires only once for a particular transaction class.
Included in Tasks	<p>“Classify Events for Business Process Insight Integration” on page 372</p>

The following elements are included:

GUI Element	Description
Cancel	Click to cancel the new action record.
Ok	Click OK to confirm the creation of the new BPI action. You return to the Evaluate or Classify Section page. Confirm the required Match Conditions are there, or create a new one.
Description	Enter a description of the BPI event.
Enter the BPI event group	Specify the BPI event group.
Enter the BPI event name	Specify the BPI event name. The TransactionVision wizard then automatically creates the appropriately configured action definition for triggering a BPI event.
Transaction Class this BPI action is associated with	Select from the drop-down the transaction class you wish to associate this action with.

 **Edit Value Properties Page**

Description	Enables you to define the conditions and data to pull into a selected transaction document. To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor , select a transaction class, in the Classify Section select Attribute Value , click New .
Important Information	See the <i>HP TransactionVision Advanced Customization Guide</i> PDF for more details on attributes and rules.

The following elements are included:

GUI Element	Description
Cancel	Click to cancel any changes made.
OK	Click to return to the previous page.
Type	Choose which transactional attribute you wish to set with this rule - Constant or XPath.
Value	<p>These control under what circumstances, which data is transferred to the transaction attribute.</p> <p>There are two types of value you can set the transaction attribute to: Constant or XPath.</p> <p>A Constant is where you assign the value a fixed string. For example, if a certain condition is met a result attribute is set to 'success'.</p> <p>When setting the attribute based on an XPath you set the transaction attribute to the value contained within a TransactionVision event when a Match Condition occurs.</p>



Edit Rule Page

Description	<p>Enables you to edit a transaction attribute rule. Attribute rules are the mechanism by which data is transferred to the transaction related attributes of each transaction instance.</p> <p>To access: Select Admin > TransactionVision > Administration > Transaction Definition Editor, select a transaction class, open its Classify Section and in the Attribute Rules section click on a Rule</p>
Important Information	<ul style="list-style-type: none">➤ Each transaction attribute can have one attribute rule associated with it.➤ If you have multiple conditions under which you wish an attribute set, create an additional Value rule section (see “Edit Value Properties Page” on page 306) for that attribute, do not create another attribute rule section.➤ Each value rule section contains a Match Condition (see “Match Conditions” on page 286).➤ Each value rule section can also be assigned an action bean, see “Actions” on page 285 for more information.➤ See the <i>HP TransactionVision Advanced Customization Guide</i> PDF for more details on attributes and rules.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Apply	Click to apply any additions or edits made.
Cancel	Click to cancel any changes made.
OK	Click to return to the previous page.
General Properties	The name displays, or if creating a new rule enter its name.
Match Conditions	Displays any set, or click New to create a new one. See “Edit Match Condition Page” on page 298 for details.

GUI Element	Description
Attribute Values	Displays the Type and the Value, or click New to create an attribute value.
Actions	Displays any actions defined for this rule, or click New to create one, see “Edit Action Properties Page” on page 301 for more information.

Troubleshooting and Limitations

This section includes the following topics:

Assign Transaction Classes to only one Schema

It is recommended that any transaction class is only assigned to one schema. If you have a transaction class that is assigned to multiple schemas this can lead to inconsistent results being reported in your Business Availability Center and TransactionVision reports and their drilldowns into TransactionVision.

Incorrect Sample Counts - Classify Rule Problems

If the classify rules for a transaction class are too general such that different parts of a transaction can be classified while the transaction is still being processed, it is possible that incorrect sample counts are sent to the Business Availability Center Dashboard. To avoid this, make the classify rules for a transaction class specific enough so that they do not match multiple events for a transaction.

10

Alias Management

This chapter includes:

Concepts

- Alias Management Overview on page 311

Reference

- Alias Management User Interface on page 312

Alias Management Overview

An alias is a custom name used to define a system model object.

- Only users with administrator privileges can perform alias management tasks.
- An alias definition maps a real object name to a user-defined alias. You may use wild cards in the mapping. For example, you could map all queues that begin with “TV” to “tvqueue”.
- When you define an object with an alias, the alias is the name used in the **Event Analysis** page. In the **Component Topology Analysis** page, you can specify whether to use the real object name, the alias, or both as the node label.
- You may have multiple alias lists.

- An alias list can contain any number of alias definitions.
- An alias list can be associated with multiple project schemas. The association between schema and alias list can be changed at any time, but any open pages must be refreshed for changes to take effect.

Alias Management User Interface

This section describes:

- Manage Alias Lists Page on page 312
- Edit Alias List Page on page 313
- Edit Schema Associations Page on page 314

Manage Alias Lists Page

Description	Enables management of alias lists. To access: Administration > TransactionVision > Administration > Alias Management.
Important Information	An alias list can contain any number of alias definitions and can be associated with multiple project schemas.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Alias list name	Alias list name.
Hide List/Show List	To hide list definitions, click the Hide List link. To show hidden list definitions, click the Show List link.
Used by Project(s)	Names of the projects using the alias.
Delete Alias List	Click to delete an alias list.

GUI Element	Description
Edit Alias List	Click to edit an alias list, see “Edit Alias List Page” on page 313.
Edit Schema Associations	Click to change project assignments for an alias list, see “Edit Schema Associations Page” on page 314.
New Alias List	Click to create a new alias list, see “Edit Alias List Page” on page 313.
Alias Name	Name of alias definition contained in this alias list.
Object	Object name that an alias maps to.



Edit Alias List Page

Description	<p>Enables you to add, edit, or delete an alias list.</p> <p>To access this page: Administration > TransactionVision > Administration > Alias Management > Manage Alias Lists, then click New Alias List, Edit Alias List, or Delete Alias List.</p>
--------------------	---

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to cancel the work done on this page.
Finish	Click when the alias list is associated with the project schema.
Alias List Name	Enter the name for this alias list. This name appears on the Manage Alias Lists page, and the Edit Project page.
Alias Name	This name appears in place of the real object name on the Component Topology Analysis (if selected), Event Analysis pages, and in all reports.

GUI Element	Description
Object	The object mapped to the alias.
Delete Selected Aliases	Select any alias you wish to delete and click this button.

Add a new alias

GUI Element	Description
Add Alias	Click this button to save this new alias definition to this alias list.
Alias Name	This name appears in place of the real object name on the Component Topology Analysis (if selected), Event Analysis pages, and in all reports.
Alias object type	The drop-down list offers all the available object types.
Object to assign alias	You can select from existing objects of the selected type, or enter an object name. If you enter an object name, you may use wildcard characters to map the alias to all objects that match the specified characters.

 **Edit Schema Associations Page**

Description	Enables you to associate an alias list with a project schema. To access: Administration > TransactionVision > Administration > Alias Management > Manage Alias Lists. Select an alias list, then click Edit Schema Associations .
Important Information	The association between schema and alias list can be changed at any time, but any open pages must be refreshed for changes to take effect.

For the selected alias list, the following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Cancel the work done on this page.
Finish	Click this button when the alias list is associated with the project schema.
Alias List Name	Displays the selected alias list.
Alias Name	List of all the aliases in this list.
Object	Object mapped to alias definition.

Assign this Alias List to a schema

GUI Element	Description
Assigned Alias List	Select the alias list to be assigned.
Schema Name	Select any number of schemas to assign this alias list to.
Used by Project(s)	The project the schema is used by.

11

Queries

This chapter includes:

Concepts

- Queries Overview on page 317
- Optimizing Query Performance on page 318

Tasks

- Add a New Query on page 319
- Edit a Query on page 321

Reference

- Query Conditions on page 322
- Queries User Interface on page 357

Queries Overview

Queries control which events in the project database are displayed in the TransactionVision views. Query conditions such as Hosts, Queue Managers, Programs, Time, and APIs enable you to restrict displayed events to only those you need to identify for analysis or to resolve a problem. Only events that meet the selected filter conditions are displayed by TransactionVision.

Queries are created within each project definition, and are not shared between projects. You can create any number of queries for a project, but only one query can be active at any time.

Standard Queries

When you create a new project, two queries are automatically created: one that includes all events and one that includes all events from the last 24 hours.

You may add any number of additional queries. However, only one query may be active for the project at any time.

Choose an Active Query

The Project Queries page lists the queries that exist in the current project. This list shows the query name and optional description. Click the radio button to the left of the query name to make it the active query for your project.

The next time you display the Component Topology Analysis or Event Analysis view, only events that match all query conditions are used to create the view.

Optimizing Query Performance

To improve the performance of queries and views, include only those technologies that your query uses. For example, if your query only includes WebSphere MQ criteria, say a WebSphere MQ object, then disable all other technologies except WebSphere MQ to make your query faster. To disable a technology, clear the check box, next to the combo box, on the left side of the query page.

For most criteria, TransactionVision uses a database lookup table, which stores a set of predefined key values along with event IDs. However, it uses a linear search for the User Buffer Data, WebSphere MQ WBI Broker, and WebSphere MQ WBI Message Flow Name criteria. Although linear searching enhances query capabilities, it reduces query performance. If you include these criteria in your query, query performance is degraded. To minimize the impact on query performance, try not to use these criteria together in the same query.

Another drawback of linear searches is that the resulting views would not know the number of events that match the query conditions. This results in limited page navigation capabilities in the TransactionVision web interface since the number of events and result pages is unknown. For example, after running a query that contains a query condition for event User Buffer data, the Event Analysis view would show page forward and backward options, but would not provide the number of events, the number of event pages, or controls to navigate to a particular page.

Add a New Query

This task includes the following steps:

- “Prerequisites” on page 319
- “Open the Add Query page” on page 319
- “Enter the query name” on page 320
- “Enter a description for the query” on page 320
- “Select one or more technologies to include in the query” on page 320
- “Specify the query criteria” on page 320
- “Change the criteria from their default values” on page 320
- “Click Finish to save all query criteria settings” on page 321

1 Prerequisites

The project that is to contain the new query has been selected. See “Projects” on page 43 for details.

2 Open the Add Query page

Click **New Query** on the Project Queries page.

3 Enter the query name

The name is used on the Component Topology Analysis, Transaction Analysis, and Event Analysis reports.

4 Enter a description for the query

The description is optional.

5 Select one or more technologies to include in the query

Check each technology you want to include in your query, or check All to include events for all technologies.

For events from .NET Agents, BEA Tuxedo and NonStop TMF Sensors, and generic Java Agents, specify the User Event technology.

6 Specify the query criteria

Click **Next** to specify the collection criteria on the Edit Query: Add/Edit Criteria page.

7 Change the criteria from their default values

- a** If you want to view only criteria for a particular technology, click the criteria technology name on the left.

Note that there is a check box on the left side of the box, which indicates whether querying on the given technology is enabled or not. Clicking on this check box toggles the status.

- b** Make desired changes on the right side of the page. See “Query Conditions” on page 322 for instructions on setting criteria for each condition.
- c** Click **Set Criteria** to save your settings for the current category.
- d** To clear changes for the current category, click **Clear Criteria**.
- e** To reset all categories to the default criteria, click **Reset All Criteria**.

8 Click Finish to save all query criteria settings

The new query appears in the Query List on the Project Query page.

The query that is selected on the main Queries Page is the active query that is used. The Component Topology and Event Analysis Views provide the ability to select a different **Current Query** in a combo box in the upper right corner.

Edit a Query

This task includes the following steps:

- “Open the Project Queries page” on page 321
- “Edit the query conditions that you want to change” on page 321
- “Change any query conditions” on page 322
- “Save the settings for each category” on page 322
- “Click Finish” on page 322

1 Open the Project Queries page

Navigate to **Admin > TransactionVision > Administration > Current Projects > Queries**.

2 Edit the query conditions that you want to change

From the Query List on the left, select the query you want to edit and click **Edit Query** (shown in step 5 of “Add a New Query” on page 319).

Alternatively, if you want to edit only specific criteria of the query, click the criteria category on the right.

The Add/Edit Criteria page is displayed. The left side of the page lists all filter conditions that you may change. The right side shows the setting for the currently selected criteria.

3 Change any query conditions

Click the criteria category name on the left side for the condition that you want to change and make desired changes on the right side of the page. See “Query Conditions” on page 322 for instructions on setting criteria for each condition.

On the top left corner, there is a technology drop-down box to switch between groups of different technologies. Note that there is a check box on the left side of the box, which indicates whether querying on the given technology is enabled or not. Clicking on this check box toggles the status.

4 Save the settings for each category

Click Set Criteria to save your settings for the current category.

To clear changes for the current category, click Clear Criteria. To reset all categories to the default criteria, click Reset All Criteria.

5 Click Finish

To save all query criteria settings, click **Finish**.

Query Conditions

This section describes the following query condition categories:

- “Common Options” on page 323
- “BTTrace Criteria” on page 326
- “CICS Criteria” on page 327
- “EJB Criteria” on page 330
- “JDBC Criteria” on page 332
- “JMS Criteria” on page 336
- “Servlet Criteria” on page 340
- “WebSphere MQ Criteria” on page 342

- “WebSphere MQ IMS Bridge Criteria” on page 351
- “User Event Criteria” on page 354



Common Options

The common options apply to all technologies. However, if you specify a criteria for All Technologies, you may specify a different criteria for an individual technology. The setting for the individual technology overrides the setting for all technologies. For example, suppose you select All Hosts for All Technologies, but on the Host page for the JMS technology, you may select only the host Host1. JMS events originating on Host1 and other technology events originating on any host all match the query criteria.

Event Time

When you create a new query, the default event time is events from the last 24 hours, relative to the last collected event. For example, if the time and date of last event in the project is 07:56 on 7-22-2005, the query includes events collected since 07:56 on 7-21-2005.

- Select the time mode. The Any Time mode means there is no condition on event time. Other modes require a time setting.
- For time modes other than Any Time, specify the time settings. Depending on the time mode, time settings may require you to specify a start and end time, a start time, an end time, or a rolling time period.

Host

- Select whether to Include or Exclude selected hosts in the query criteria.
- Check the box next to each host you want to include specifically in the query criteria, or check the box next to Select All Host to include all hosts in the query criteria. The host list is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Host list field and click **Add**. To remove a host from the list, check the box next to the host and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Program/Servlet/EJB

- Determine whether to **Include** or **Exclude** selected programs, servlets, and/or EJBs in the filter criteria.
- Check the box next to each program, servlet, or EJB you want to include in the filter criteria, or check the box next to **All Program/Servlet/EJB** to include all programs, servlets, and EJBs in the filter criteria. The program list is generated from events collected in the project database.
- To manually add a new program to the list, enter the program name and click **Add**. To remove a program from the list, check the box next to the program name and click **Delete**. Only programs that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

User Name

- Select whether to Include or Exclude selected user names in the query criteria.
- Check the box next to each user name you want to include specifically in the query criteria, or check the box next to Select All User Name to include all user names in the query criteria. The user name list is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize User Name list field and click **Add**. To remove a user name from the list, check the box next to the user name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

User Buffer Data

Note: The TransactionVision Analyzer can be configured to use an XML Event Compression bean, which can reduce the size of the resulting event database. However, if this bean is in use by the Analyzer, it is not possible to query on the user buffer data.

- Click the appropriate radio button to indicate whether the specified user data is in binary string or normal string format and enter the buffer data. For normal string format, select the correct code page. To query on a double-byte character set (DBCS) string, you may either type in the DBCS text string or enter the binary representation of the text.
- Specify where to search. If you have non-XML based user data in events, select Search in user data. If you have XML user data and have enabled the DefaultModifierBean in the Analyzer, select Search in event XML to search for strings in that XML data.

User Buffer Data is not supported for SQL Server.

Note: TransactionVision uses linear searching for this criteria. Although linear searching enhances query capabilities, it reduces query performance. If you include this criteria in your query, query performance is degraded. To minimize the impact on query performance, try not to use this criteria with the other linear search criteria (WebSphere MQ WBI Broker and WebSphere MQ WBI Message Flow Name) in the same query.

Transaction Attributes

The transaction related query criteria follow the same basic rules as other query criteria, allowing you to add criteria based on the transaction in which the event was a part.

For example, you can combine these two criteria:

- Ask for all events that have a certain program name.
- Ask only for events that are part of transactions that have been marked as failed.

The available options are:

- **Transaction Class** — query on selected transaction class.
- **Transaction Starttime** — query on when the transaction starts.
- **Transaction Endtime** — query on when the transaction ends.
- **Transaction Response Time** — query on the response time, in milliseconds, of the transaction.
- **Transaction State** — select events from transactions that are in a Completed/Processing/Unknown state.
- **Transaction Result** — select events from transactions that have a Failed/Success/Unknown result.
- **Transaction Label** — select events from transactions based on matching the transaction label.
- **Transaction SLA State** — query on the SLA state of Aged out/None/Violated.



BTTrace Criteria

The following options are only available when BTTrace is the current **Criteria For:** technology.

API Name

- Select whether to Include or Exclude events for selected BTTRACE APIs. Applications may use the BTTRACE function provided by the Sensor library to send user-defined trace messages to the Analyzer.
- Check the box next to each API you want to include specifically in the query criteria, or check the box next to Select All API NAME to include all BTTRACE APIs in the query criteria. For more information about BTTRACE APIs, see “BTTRACE API” on page 257.

Severity

- Select whether to Include or Exclude events for selected BTTRACE severity levels. Applications may use the BTTRACE function provided by the Sensor library to send user-defined trace messages to the Analyzer.
- Check the box next to each severity you want to include specifically in the query criteria, or check the box next to Select All Severity to include all BTTRACE severity levels in the query criteria. For more information about BTTRACE severity levels, see “BTTRACE Severity” on page 257.



CICS Criteria

The following options are only available when CICS is the current **Criteria for** technology.

API Name

- Select whether to Include or Exclude events for selected APIs.
- Check the box next to each API you want to include specifically in the query criteria, or check the box next to Select All API Name to include all APIs in the query criteria.

API Type

- Select whether to Include or Exclude events for selected API types.
- Check the box next to each API type you want to include specifically in the query criteria, or check the box next to Select All API Type to include all API types in the query criteria.

File Name

- Select whether to Include or Exclude events for selected file names.
- Check the box next to each file name you want to include specifically in the query criteria, or check the box next to Select All File Name to include all file names in the query criteria. The list of file names is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize File Name list field and click **Add**. To remove a file name from the list, check the box next to the file name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Response Code

- Select whether to Include or Exclude events for selected response codes.
- Check the box next to each response code you want to include specifically in the query criteria, or check the box next to Select All Response Code to include all response codes in the query criteria.

SYSID

- Select whether to Include or Exclude events for selected SYSID values.
- Check the box next to each SYSID you want to include specifically in the query criteria, or check the box next to Select All SYSID to include all SYSID values in the query criteria. The list of SYSID values is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize SYSID list field and click **Add**. To remove a SYSID value from the list, check the box next to the SYSID value and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

TD Queue Name

- Select whether to Include or Exclude events for selected transient data (TD) queue names.
- Check the box next to each TD queue name you want to include specifically in the query criteria, or check the box next to Select All TD Queue Name to include all TD queue names in the query criteria. The list of TD queue names is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize TD Queue Name list field and click **Add**. To remove a TD queue name from the list, check the box next to the TD queue name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

TS Queue ID

- Select whether to Include or Exclude events for selected temporary storage (TS) queue IDs.
- Check the box next to each TS queue ID you want to include specifically in the query criteria, or check the box next to Select All TS Queue ID to include all TS queue IDs in the query criteria. The list of TS queue IDs is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize TS Queue ID list field and click **Add**. To remove a TS queue ID from the list, check the box next to the TS queue ID and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Terminal ID

- Select whether to Include or Exclude events for selected terminal IDs.
- Check the box next to each terminal ID you want to include specifically in the query criteria, or check the box next to Select All Terminal ID to include all terminal IDs in the query criteria. The list of terminal IDs is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Terminal ID list field and click **Add**. To remove a Terminal ID from the list, check the box next to the Terminal ID and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Transaction ID

- Select whether to Include or Exclude events for selected transaction IDs.
- Check the box next to each transaction ID you want to include specifically in the query criteria, or check the box next to Select All Transaction ID to include all transaction IDs in the query criteria. The list of transaction IDs is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Transaction ID list field and click **Add**. To remove a transaction ID from the list, check the box next to the transaction ID and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.



EJB Criteria

The following options are only available when EJB is the current **Criteria for** technology.

Application Server

- Select whether to Include or Exclude events for selected application servers.
- Check the box next to each application server you want to include specifically in the query criteria, or check the box next to Select All Application Server to include all application servers in the query criteria. The list of application servers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Application Server list field and click **Add**. To remove an application server from the list, check the box next to the application server and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

EJB Application

- Select whether to Include or Exclude events for selected EJB applications.
- Check the box next to each EJB application you want to include specifically in the query criteria, or check the box next to Select All EJB Application to include all EJB applications in the query criteria. The list of EJB applications is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize EJB Application list field and click **Add**. To remove an EJB application from the list, check the box next to the EJB application and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

EJB Name

- Select whether to Include or Exclude events for selected EJB names.
- Check the box next to each EJB name you want to include specifically in the query criteria, or check the box next to Select All EJB Name to include all EJB names in the query criteria. The list of EJB names is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize EJB Name list field and click **Add**. To remove an EJB name from the list, check the box next to the EJB name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

EJB Method

- Select whether to Include or Exclude events for selected EJB methods.
- Check the box next to each EJB method you want to include specifically in the query criteria, or check the box next to Select All EJB Method to include all EJB methods in the query criteria. The list of EJB methods is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize EJB Method list field and click **Add**. To remove an EJB method from the list, check the box next to the EJB method and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

EJB Status

- Select whether to Include or Exclude events for the selected EJB status.
- Check Failure to include only failed EJB events, or check Success to include only successful EJB events. Check Select All EJB Status to include successful and failed EJB events in the query criteria.



JDBC Criteria

The following options are only available when JDBC is the current **Criteria** for technology.

Application Server

- Select whether to Include or Exclude events for selected JDBC application servers
- Check the box next to each application server you want to include specifically in the query criteria, or check the box next to Select All Application Server to include all application servers in the query criteria. The list of application servers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Application Server list field and click **Add**. To remove an application server from the list, check the box next to the application server and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Class

- Select whether to Include or Exclude events for selected JDBC classes.
- Check the box next to each JDBC class you want to include specifically in the filter criteria, or check the box next to All JDBC Class to include all classes in the query criteria.

Database

- Select whether to Include or Exclude events for selected JDBC database.
- Check the box next to each database you want to include specifically in the filter criteria, or check the box next to Select All Database to include all databases in the query criteria.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Database list field and click **Add**. To remove a database from the list, check the box next to the database and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Database Object

- Select whether to Include or Exclude events that make use of certain database objects
- Check the box next to each database object you want to include specifically in the filter criteria, or check the box next to All Entries to include all objects in the query criteria.

Database objects are grouped by the schema that they belong to. For example, selecting a database table 'X' under schema 'Y' only selects JDBC events that access table 'X', if table 'X' is being used under schema 'Y'. Additionally by selecting a schema, you can filter on all tables that might occur under that schema.

- Select whether to find events that access all selected objects, or find events that access any one of the selected objects.

- The list of database objects is generated from events collected in the project database. To add a custom database object to the list, select the schema you want this database object to be under, enter the database object name and click **Add**. To remove a custom database object from the list, check the box next to it and click **Delete**.

Method

- Select whether to Include or Exclude events for selected JMS methods.
- Check the box next to each JMS method you want to include specifically in the query criteria, or check the box next to Select All Method to include all methods in the query criteria. The list of methods is generated from events collected in the project database.

Result Code

- Select whether to Include or Exclude events that have a particular result code.
- Check the box next to each result code you want to include specifically in the filter criteria, or check the box next to Select All Result Code to include all result codes in the query criteria. Result codes can be: Failure, Success or Warning.

SQL Code

SQL Code does not provide pre-filled values. Any values you want to query on must be first entered.

- To add an SQL code to the list, enter its name and click **Add**. To remove an SQL code from the list, check the box next to it and click **Delete**.
- Select whether to Include or Exclude events that have a particular SQL Code.
- Check the box next to each SQL code you want to include specifically in the filter criteria, or check the box next to Select All SQL Code to include all SQL codes in the query criteria.

SQL Codes used by the JDBC application vary based on the underlying database vendor.

SQL State

SQL State does not provide pre-filled values. Any values you want to query on must be first entered.

- To add an SQL state to the list, enter its name in the Customize SQL State list field and click **Add**. To remove an SQL state from the list, check the box next to it and click **Delete**.
- Select whether to Include or Exclude events that have a particular SQL state.
- Check the box next to each SQL state you want to include specifically in the filter criteria, or check the box next to Select All SQL State to include all SQL states in the query criteria.

Similar to SQL Codes, SQL States used by the JDBC application vary based on the underlying database vendor.

SQL Statement

- Select whether to Include or Exclude events that have a particular SQL statement.
- Check the box next to each SQL statement you want to include specifically in the filter criteria, or check the box next to Select All SQL Statement to include all SQL statements in the query criteria.

SQL Statement Type

- Select whether to Include or Exclude events for a particular SQL statement type.
- Check the box next to the statement type (such as Select, Merge, Insert, and Update) you want to include specifically in the filter criteria, or check the box next to Select All SQL Statement Type to include all statement types in the query criteria.

Web Application

- Select whether to Include or Exclude events for selected web applications.
- Check the box next to each web application you want to include specifically in the filter criteria, or check the box next to Select All Web Application to include all web applications in the filter criteria. The list of web applications is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Web Application list field and click **Add**. To remove a web application from the list, check the box next to the web application and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

JMS Criteria

The following options are only available when JMS is the current **Criteria** for technology.

Web Application

- Select whether to Include or exclude Events for selected web applications.
- Check the box next to each web application you want to include specifically in the query criteria, or check the box next to Select All Web Application to include all web applications in the query criteria. The list of web applications is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Web Application list field and click **Add**. To remove a web application from the list, check the box next to the Web application and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Class

- Select whether to Include or Exclude events for selected JMS classes.
- Check the box next to each JMS class you want to include specifically in the query criteria, or check the box next to Select All Class to include all classes in the query criteria. The list of classes is generated from events collected in the project database.

Method

- Select whether to Include or Exclude events for selected JMS methods.
- Check the box next to each JMS method you want to include specifically in the query criteria, or check the box next to Select All Method to include all methods in the query criteria. The list of methods is generated from events collected in the project database.

Connection Name

- Select whether to Include or Exclude events for selected connection names.
- Check the box next to each connection name you want to include specifically in the query criteria, or check the box next to Select All Connection Name to include all connection names in the query criteria. The list of connection names is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Connection Name list field and click **Add**. To remove a connection name from the list, check the box next to the connection name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

JMS Topic

- Select whether to Include or Exclude events for selected JMS topics.
- Check the box next to each JMS topic you want to include specifically in the query criteria, or check the box next to All Entries to include all JMS topics in the query criteria. The list of JMS topics is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize JMS Topic list field and click **Add**. To remove a JMS topic from the list, check the box next to the topic and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

JMS Queue

- Select whether to Include or Exclude events for selected JMS queues.
- Check the box next to each queue you want to include specifically in the query criteria, or check the box next to All Entries to include all JMS queues in the query criteria. The list of JMS queues is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize JMS Queue list field and click **Add**. To remove a JMS queue from the list, check the box next to the JMS queue and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

WebSphere MQ Queue Manager

- Select whether to Include or Exclude events for selected WebSphere MQ queue managers.
- Check the box next to each queue you want to include specifically in the query criteria, or check the box next to Select All WebSphere MQ Queue Manager to include all queue managers in the query criteria. The list of queue managers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize WebSphere MQ Queue Manager list field and click **Add**. To remove a queue manager from the list, check the box next to the queue manager and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

WebSphere MQ Queue

- Select whether to Include or Exclude events for selected WebSphere MQ queues.
- Check the box next to each queue you want to include specifically in the query criteria, or check the box next to Select All WebSphere MQ Queue to include all queues in the query criteria. The list of queues is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, select the queue manager in the Owner list, enter the WebSphere MQ queue name in the Customize WebSphere MQ Queue list field and click **Add**. To remove a WebSphere MQ queue from the list, check the box next to the WebSphere MQ queue and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Data Size

- Select whether to Include or Exclude events for the selected data size.
- Select the matching criteria for the selected data size. Choices include less than, less than or equal to, equal to, not equal to, greater than or equal to, or greater than.
- Enter the desired data size in bytes.

Exception Code

- Select whether to Include or Exclude events for selected exception codes.
- Check the box next to each exception code you want to include specifically in the query criteria, or check the box next to Select All Exception Code to include all exception codes in the query criteria. The list of exception codes is generated from events collected in the project database.

Exception Class Name

- Select whether to Include or Exclude events for selected exception class names.
- Check the box next to each exception class name you want to include specifically in the query criteria, or check the box next to Select All Exception Class Name to include all exception class names in the query criteria. The list of exception class names is generated from events collected in the project database.



Servlet Criteria

The following options are only available when Servlet is the current **Criteria** for technology.

Web Application Server

- Select whether to Include or Exclude events for selected web application servers.
- Check the box next to each web application server you want to include specifically in the query criteria, or check the box next to Select All Web Application Server to include all web application servers in the query criteria. The list of web application servers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Web Application Server list field and click **Add**. To remove a web application server from the list, check the box next to the web application server and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Web Application

- Select whether to Include or Exclude events for selected web applications.
- Check the box next to each web application you want to include specifically in the query criteria, or check the box next to Select All Web Application to include all web applications in the query criteria. The list of web applications is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Web Application list field and click **Add**. To remove a web application from the list, check the box next to the web application and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Method

- Select whether to Include or Exclude events for selected servlet methods.
- Check the box next to each servlet method you want to include specifically in the query criteria, or check the box next to Select All Method to include all servlet methods in the query criteria. The list of servlet methods is generated from events collected in the project database.

Status Code

- Select whether to Include or Exclude events for selected status codes.
- Check the box next to each status code you want to include specifically in the query criteria, or check the box next to Select All Status Code to include all status codes in the query criteria. The list of status codes is generated from events collected in the project database.

Data Size

- Select whether to Include or Exclude events for the selected data size.
- Select the matching criteria for the selected data size. Choices include less than, less than or equal to, equal to, not equal to, greater than or equal to, or greater than.
- Enter the desired data size in bytes.

WebSphere MQ Criteria

The following options are only available when WebSphere MQ is the current **Criteria for** technology.

API Name

- Select whether to Include or Exclude events for selected WebSphere MQ APIs.
- Check the box next to each API you want to include specifically in the query criteria, or check the box next to Select All API Name to include all WebSphere MQ APIs in the query criteria. The API list is generated from events collected in the project database.

Connection Name

The connection name and queue manager are typically the same, except in the case where a queue on one queue manager is opened using a connection (specified with the **hConn** parameter) to a different queue manager. For example, suppose a program opens queue QM1.Q using a connection to queue manager QM1, then later opens QM1.Q using a connection to queue manager QM2. This event shows QM1.Q as the object name, QM1 as the queue manager, and QM2 as the connection name.

- Select whether to Include or Exclude events for the selected WebSphere MQ connections.
- Check the box next to each connection you want to include specifically in the query criteria, or check the box next to Select All Connection Name to include all connections in the query criteria. The list of connection names is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Connection Name list field and click **Add**. To remove a connection name from the list, check the box next to the connection name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Queue Manager

- Select whether to Include or Exclude events for selected WebSphere MQ queue managers.
- Check the box next to each queue manager you want to include specifically in the query criteria, or check the box next to Select All Queue Manager to include all WebSphere MQ queue managers in the query criteria. The list of queue managers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Queue Manager list field and click **Add**. To remove a queue manager from the list, check the box next to the queue manager and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Object

- Select whether to Include or Exclude events for the selected WebSphere MQ objects.
- Check the box next to each object you want to include specifically in the query criteria, or check the box next to All Entries to include all objects in the query criteria. The list of objects is generated from events collected in the project database. They are grouped according to the queue managers they are associated with.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, select the queue manager in the Owner list, enter the object name in the Customize Object list field and click **Add**. To remove an object from the list, check the box next to the object and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

ReplyTo Queue Manager

- Select whether to Include or Exclude events where the selected WebSphere MQ queue managers match the Reply To field.
- Check the box next to each queue manager you want to include specifically in the query criteria, or check the box next to Select All Reply To Queue Manager to include all WebSphere queue managers in the query criteria. The list of queue managers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Reply to Queue Manager list field and click **Add**. To remove a Reply To queue manager from the list, check the box next to the Reply To queue manager and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

ReplyTo Object

- Select whether to Include or Exclude events where the selected WebSphere MQ objects match the Reply To field.
- Check the box next to each object you want to include specifically in the query criteria, or check the box next to All Entries to include all WebSphere MQ objects in the query criteria. The list of objects is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, select the queue manager in the Owner list, enter the object in the Customize Reply To Object list field and click **Add**. To remove an object from the list, check the box next to the object and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Data Size

- Select whether to Include or Exclude events for the selected data size.
- Select the matching criteria for the selected data size. Choices include less than, less than or equal to, equal to, not equal to, greater than or equal to, or greater than.
- Enter the desired data size in bytes.

Completion Code

- Select whether to Include or Exclude events for selected WebSphere MQ API completion codes.
- Check the box next to each API completion code you want to include specifically in the query criteria, or check the box next to Select All Completion Code to include all API completion codes in the query criteria.

Reason Code

- Select whether to Include or Exclude events for selected WebSphere MQ API reason codes.
- Check the box next to each API reason code you want to include specifically in the query criteria, or check the box next to Select All Reason Code to include all reason codes in the query criteria.

Message ID

- Select whether the Message ID contains or does not contain either a binary string or normal string.
- Click the appropriate radio button to indicate whether the specified message ID is in string or binary format and enter the message ID. For string format, select the correct code page.

Correlation ID

- Select whether the Correlation ID contains or does not contain either a binary string or normal string.
- Click the appropriate radio button to indicate whether the specified correlation ID is in string or binary format and enter the correlation ID. For string format, select the correct code page.

z/OS Criteria

Job Name

- Select whether to Include or Exclude events for selected z/OS job names.
- Check the box next to each job name you want to include specifically in the query criteria, or check the box next to Select All Job Name to include all job names in the query criteria. The list of job names is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Job Name list field and click **Add**. To remove a job name from the list, check the box next to the job name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Job Step

- Select whether to Include or Exclude events for selected z/OS job steps.
- Check the box next to each job step you want to include specifically in the query criteria, or check the box next to Select All Job Step to include all job steps in the query criteria. The list of job steps is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Job Step list field and click **Add**. To remove a job step from the list, check the box next to the job step and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

CICS Task

- Select whether to Include or Exclude events for selected z/OS CICS tasks.
- Check the box next to each CICS task you want to include specifically in the query, or check the box next to Select All CICS Task to include all CICS tasks in the query criteria. The list of CICS tasks is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize CICS Task list field and click **Add**. To remove a CICS task from the list, check the box next to the CICS task and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

CICS SYSID

- Select whether to Include or Exclude events for selected z/OS CICS SYSIDs.
- Check the box next to each CICS SYSID you want to include specifically in the query, or check the box next to Select All CICS SYSID to include all CICS SYSIDs in the query criteria. The list of CICS SYSIDs is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize CICS SYSID list field and click **Add**. To remove a CICS SYSID from the list, check the box next to the CICS SYSID and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

CICS Transaction

- Select whether to Include or Exclude events for selected z/OS CICS transactions.
- Check the box next to each CICS transaction you want to include specifically in the criteria, or check the box next to Select All CICS Transaction to include all CICS transactions in the query criteria. The list of CICS transactions is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize CICS Transaction list field and click **Add**. To remove a CICS transaction from the list, check the box next to the CICS transaction and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

IMS Region Type

- Select whether to Include or Exclude events for selected z/OS IMS region types.
- Check the box next to each IMS region type you want to include specifically in the query criteria, or check the box next to Select All IMS Region Type to include all IMS region types in the query criteria. The list of IMS region types is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize IMS Region Type list field and click **Add**. To remove an IMS region type from the list, check the box next to the IMS region type and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

IMS Region Identifier

- Select whether to Include or Exclude events for selected z/OS IMS region identifiers.
- Check the box next to each IMS region identifier you want to include specifically in the query criteria, or check the box next to Select All IMS Region Identifier to include all IMS region identifiers in the query criteria. The list of IMS region identifiers is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize IMS Region Identifier list field and click **Add**. To remove an IMS region identifier from the list, check the box next to the IMS region identifier and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

IMS Identifier

- Select whether to Include or Exclude events for selected z/OS IMS identifiers.
- Check the box next to each IMS identifier you want to include specifically in the query criteria, or check the box next to Select All IMS Identifier to include all IMS identifiers in the query criteria. The list of IMS identifiers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize IMS Identifier list field and click **Add**. To remove an IMS identifier from the list, check the box next to the IMS identifier and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

IMS Transaction

- Select whether to Include or Exclude events for selected z/OS IMS transactions.
- Check the box next to each IMS transaction you want to include specifically in the query criteria, or check the box next to Select All IMS Transaction to include all IMS transactions in the query criteria. The list of IMS transactions is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize IMS Transaction list field and click **Add**. To remove an IMS transaction from the list, check the box next to the IMS transaction and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

IMS PSB

- Select whether to Include or Exclude events for selected z/OS IMS PSBs.
- Check the box next to each IMS PSB you want to include specifically in the query criteria, or check the box next to Select All IMS PSB to include all IMS PSBs in the query criteria. The list of IMS PSBs is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize IMS PSB list field and click **Add**. To remove an IMS PSB from the list, check the box next to the IMS PSB and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

i5/OS Criteria

Job Name

- Select whether to Include or Exclude events for selected i5/OS job names.
- Check the box next to each job name you want to include specifically in the query criteria, or check the box next to Select All Job Name to include all job names in the query criteria. The list of i5/OS job names is generated from events collected in the project database.

- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Job Name list field and click **Add**. To remove a job name from the list, check the box next to the job name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.



WebSphere MQ IMS Bridge Criteria

The following options are only available when WebSphere MQ IMS Bridge is the current **Criteria for** technology.

API Name

- Select whether to Include or Exclude events for selected WebSphere MQ IMS Bridge APIs.
- Check the box next to each API you want to include specifically in the query criteria, or check the box next to Select All API Name to include all APIs in the query criteria. The list of APIs is generated from events collected in the project database.

IMS Transaction

- Select whether to Include or Exclude events for selected IMS transactions.
- Check the box next to each transaction you want to include specifically in the query criteria, or check the box next to Select All IMS Transaction to include all transactions in the query criteria. The list of transactions is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize IMS Transaction list field and click **Add**. To remove an IMS transaction from the list, check the box next to the IMS transaction and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Connection Name

- Select whether to Include or Exclude events for selected connection names.
- Check the box next to each connection name you want to include specifically in the query criteria, or check the box next to Select All Connection Name to include all connection names in the query criteria. The list of connection names is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Connection Name list field and click **Add**. To remove a connection name from the list, check the box next to the connection name and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Queue Manager

- Select whether to Include or Exclude events for selected queue managers.
- Check the box next to each queue manager you want to include specifically in the query criteria, or check the box next to Select All Queue Manager to include all queue managers in the query criteria. The list of queue managers is generated from events collected in the project database.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, enter it in the Customize Queue Manager list field and click **Add**. To remove a queue manager from the list, check the box next to the queue manager and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Object

- Select whether to Include or Exclude events for selected objects.
- Check the box next to each object you want to include specifically in the query criteria, or check the box next to All Entries to include all objects in the query criteria. The list of objects is generated from events collected in the project database. The objects are grouped according to the queue managers they are associated with.
- In addition to selecting individual entries, you may specify a custom query string using wildcard characters. To specify a query using wildcards, select the queue manager in the Owner list, enter the object in the Customize Object list field and click **Add**. To remove an object from the list, check the box next to the object and click **Delete**.

In a customized query, the * character represents zero or more characters, and the double quote (") character represents an escape.

Message ID

- Select whether the Message ID contains or does not contain either a binary string or normal string.
- Click the appropriate radio button to indicate whether the specified message ID is in string or binary format and enter the message ID. For string format, select the correct code page.

Correlation ID

- Select whether the Correlation ID contains or does not contain either a binary string or normal string.
- Click the appropriate radio button to indicate whether the specified correlation ID is in string or binary format and enter the correlation ID. For string format, select the correct code page.

User Event Criteria

This section describes the user event criteria that can be used to query the User Event Sensor events.

Technology

Specifies the type of User Event Sensor: .NET, Tuxedo, NonStop-TMF, or Java. This is the main User Event criteria that allows you to distinguish between different technologies/Sensors that use User Events.

Class

Specifies which classes are filtered.

- Determine whether to **Include** or **Exclude** events for selected User Event classes.
- Check the box next to each User Event class you want to include specifically in the filter criteria, or check the box next to **All Class** to include all classes in the filter criteria. The list of classes is generated from events collected in the project database.
- To manually add a class to the list, enter its name and click **Add**. To remove a class from the list, check the box next to it and click **Delete**. Only classes that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Method

Specify which methods are filtered:

- Determine whether to **Include** or **Exclude** events for selected User Event methods.
- Check the box next to each User Event method you want to include specifically in the filter criteria, or check the box next to **All Method** to include all methods in the filter criteria. The list of methods is generated from events collected in the project database.
- To manually add a method to the list, enter its name and click **Add**. To remove a method from the list, check the box next to it and click **Delete**. Only methods that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Completion Code

- Determine whether to **Include** or **Exclude** events for selected User Event method completion codes.
- Check the box next to each method completion code you want to include specifically in the filter criteria, or check the box next to **All Completion Codes** to include all possible User Event completion codes in the filter criteria.

Status

- Determine whether to **Include** or **Exclude** events for selected User Event status strings.
- Check the box next to each status string you want to include specifically in the filter criteria, or check the box next to **All Status** to include all status strings in the filter criteria. The list of status strings is generated from events collected in the project database.
- To manually add a status string to the list, enter its name and click **Add**. To remove a Status string from the list, check the box next to it and click **Delete**. Only status strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Queue

- Determine whether to **Include** or **Exclude** events for selected User Event queue strings.
- Check the box next to each queue string you want to include specifically in the filter criteria, or check the box next to **All Queue** to include all queue strings in the filter criteria. The list of queue strings is generated from events collected in the project database.
- To manually add a queue string to the list, enter its name and click **Add**. To remove a queue string from the list, check the box next to it and click **Delete**. Only queue strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Queue Space

- Determine whether to **Include** or **Exclude** events for selected User Event queue space strings.
- Check the box next to each queue space string you want to include specifically in the filter criteria, or check the box next to **All Queue Space** to include all queue space strings in the filter criteria. The list of queue space strings is generated from events collected in the project database.
- To manually add a queue space string to the list, enter its name and click **Add**. To remove a queue space string from the list, check the box next to it and click **Delete**. Only queue space strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Service

- Determine whether to **Include** or **Exclude** events for selected User Event service strings.
- Check the box next to each service string you want to include specifically in the filter criteria, or check the box next to **All Service** to include all service strings in the filter criteria. The list of service strings is generated from events collected in the project database.
- To manually add a service string to the list, enter its name and click **Add**. To remove a service string from the list, check the box next to it and click **Delete**. Only service strings that were manually added and do not exist in the project may be deleted. An asterisk may be used at the beginning and/or end as a wildcard.

Queries User Interface

This section describes:

- Project Queries Page on page 357
- Add Query Page on page 358
- Edit Query: Add/Edit Criteria Page on page 359

Project Queries Page

Description	Displays the query conditions for each query. Click the radio button next to a query name to make it the active query. To access: Select Admin > TransactionVision > Administration > Current Projects > Queries
Included in Tasks	“Add a New Query” on page 319
Useful Links	“Projects” on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Delete Query	Click to delete the selected query.
Edit Query	Click to modify a selected query’s criteria. Displays the “Edit Query: Add/Edit Criteria Page” on page 359.
Query Detail	Displays the selected query’s criteria settings.
Query List	Displays existing data collection queries.
New Query	Click to create a new data collection query. Displays the “Add Query Page” on page 358.



Add Query Page

Description	Allows you to create a new query for a project. To access: Select Admin > TransactionVision > Administration > Current Projects > Queries , click New Query .
Included in Tasks	“Add a New Query” on page 319
Useful Links	“Projects” on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to return to the Project Queries page without creating a new query.
Next >	Click to specify criteria for the new query. Displays the “Edit Query: Add/Edit Criteria Page” on page 359.
Description	Enter a description for the query. (Optional)
Query Name	Enter a unique name for the query. This name is used on the on the Component Topology Analysis, Transaction Analysis, and Event Analysis views.
Technology	Check each technology you want to include in your query, or check All to include events for all technologies. NOTE: For events from .NET, BEA Tuxedo, or NonStop TMF Sensors, specify the User Event technology.

Edit Query: Add/Edit Criteria Page

Description	<p>Allows you to modify a query for a project.</p> <p>To access:</p> <ul style="list-style-type: none"> ➤ On the Admin > TransactionVision > Administration > Current Projects > Queries page, click New Query, fill in the fields and click Next. ➤ On the Admin > TransactionVision > Administration > Current Projects > Queries page, click Edit Query.
Important Information	<p>The left side of the page lists all query conditions that you may change. The right side shows the setting for the currently selected criteria.</p> <p>Click the criteria category name on the left side for the condition that you want to change from the default value, and make desired changes on the right side of the page. See “Query Conditions” on page 322 for instructions on setting criteria for each condition.</p>
Included in Tasks	“Add a New Query” on page 319
Useful Links	“Projects” on page 43

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Cancel	Click to return to the previous query page without modifying the default query criteria.
Clear Criteria	Click to clear changes for the current category.
Finish	Click when all query criteria are set.
Reset All Criteria	Click to reset all categories to the default criteria.
Set Criteria	Click to save your settings for the current category.
Common Options	Criteria to filter events for all technologies. See “Common Options” on page 323.
Criteria for drop-down	<p>Select the Technology (Sensor) for which you wish to specify query criteria from the drop-down list. You can specify criteria from the Common Options that apply to all technologies (All Sensors), or you may specify criteria for a specific technology:</p> <ul style="list-style-type: none"> ➤ “Common Options” on page 323 ➤ “BTTrace Criteria” on page 326 ➤ “CICS Criteria” on page 327 ➤ “EJB Criteria” on page 330 ➤ “JDBC Criteria” on page 332 ➤ “JMS Criteria” on page 336 ➤ “Servlet Criteria” on page 340 ➤ “WebSphere MQ Criteria” on page 342 ➤ “WebSphere MQ IMS Bridge Criteria” on page 351 ➤ “User Event Criteria” on page 354
Criterion Setting for <criteria>	Dialog box for setting the query criterion.
Summary	Lists the query criteria for each Sensor.

12

TransactionVision HP Business Settings

This chapter includes:

Concepts

- HP Business Availability Settings on page 362
- TransactionVision Data in Business Availability Center on page 362
- Drill downs from End User Management to TransactionVision on page 366
- HP Business Process Insight Settings on page 369

Tasks

- Classify Events for Business Process Insight Integration on page 372

Reference

- HP Business Settings User Interface on page 376

Troubleshooting and Limitations on page 381

HP Business Availability Settings

As part of the HP Business Transaction Management solution, you can customize how TransactionVision integrates with Business Availability Center, from **Admin > TransactionVision > Administration > HP Business Availability Settings**.

See “HP Business Settings User Interface” on page 376 for details on the various pages and dialogs involved in configuration, and the following sections:

- “TransactionVision Data in Business Availability Center” on page 362
- “Drill downs from End User Management to TransactionVision” on page 366

TransactionVision Data in Business Availability Center

The following sections describe the structure and mapping of data between TransactionVision and Business Availability Center.

- “Transaction Class/Business Transaction CI Mapping” on page 363
- “Publishing Transaction Classes to CMDB” on page 363
- “Attributes of the Transaction Monitor CI” on page 363
- “BAC Aggregate TV Sample Data” on page 364
- “Infrastructure CIs Populated by TransactionVision” on page 365

Transaction Class/Business Transaction CI Mapping

TransactionVision transaction classes correspond to *Business Transaction* and *TV Monitor* CIs in Business Availability Center. The Business Transaction object defines the class entity, and the TV Monitor holds a specific subset of the class attributes relevant to Business Availability Center.

Each transaction class in TransactionVision is mapped to a CI of type **Business Transaction** in the CMDB. Each Business Transaction CI has an attached child CI (using **monitored_by** link) of type **TV Monitor**. Both these CIs take their name from the name of the transaction class in TransactionVision.

Publishing Transaction Classes to CMDB

Whenever a transaction class is updated in TransactionVision, the update is propagated to the corresponding CIs in the CMDB. Creating a new transaction class triggers the automatic creation of the CIs for the new business transaction in the CMDB. Deleting a transaction class causes the deletion of the corresponding business transaction CIs.

Attributes of the Transaction Monitor CI

Aggregation delta	The time period allocated for the delay while the sample is created. The default is 360 seconds (defined in TransactionVision).
Collection interval	The time period, in seconds, covered by each sample received for this business transaction. The default is 300 seconds (defined in TransactionVision).
TransactionVision currency unit	The applicable currency for the monetary values received for this business transaction. This currency unit is displayed in the Dashboard tabs and reports for all monetary values associated with this business transaction.
TransactionVision Exception attribute rule	The check box for this property is selected if the Exception attribute rule is defined for the transaction class in TransactionVision. If the check box is selected, Business Availability Center assigns the Exceptions KPI to the business transaction.

TransactionVision Failures attribute rule	The check box for this property is selected if the Failures attribute rule is defined for the transaction class in TransactionVision. If the check box is selected, Business Availability Center assigns the Failures KPI to the business transaction.
TransactionVision ID	The internal ID for the business transaction in TransactionVision.
TransactionVision threshold	The response time threshold defined for the business transaction in TransactionVision. Instances of the transaction that take longer than this time are considered <i>late</i> .
TransactionVision Value attribute rule	The check box for this property is selected if the Value attribute rule is defined for the transaction class in TransactionVision. If the check box is selected, Business Availability Center assigns the Value KPI to the business transaction.

BAC Aggregate TV Sample Data

Business transaction samples are sent to Business Availability Center through a TransactionVision scheduled job running in the TransactionVision UI/Job server. BAC Aggregate TV Sample Data creates and delivers aggregate business transaction data to Business Availability Center. It queries the business transaction table, and delivers aggregated samples of monitored transactions every 5 minutes (300 seconds). The sample then delivers pre-defined metrics to Business Availability Center via HTTP Post.

Samples are sent for two modes of business transaction: **completed** and **in-process**:

- The **completed** metrics are based on the aggregation of all completed instances (reached their end criteria) during the time frame.
- The **in-process** metrics are based on a snapshot of all in-process instances (not yet reached their end criteria) at the time TransactionVision created the sample.

The metrics for both these transaction instance types are mapped to a single TV Monitor CI (based on the TransactionVision ID attached to the CI), but different KPI rules deal with each metric type.

Samples also contain corresponding end user data forwarded by Real User Monitor, if the RUM communicator link is enabled.

The samples contain numerical count and monetary value metrics for the completed and in-process transaction instances, broken down by various state categories.

For more details, see “Data Sample for TransactionVision” in *Reference Information*.

Infrastructure CIs Populated by TransactionVision

TransactionVision publishes CIs to the uCMDB, corresponding to the infrastructure elements discovered by TransactionVision sensors. The **CMDB Population** job publishes J2EE applications, J2EE servers on which the J2EE applications reside, and any machines on which J2EE servers are hosted, to the uCMDB. This job also creates a dependency relationship between business transactions and corresponding J2EE applications. These dependencies can then be viewed from the Business Transactions view in the Business Availability Center Dashboard.

Drill downs from End User Management to TransactionVision

The integration between End User Management and TransactionVision helps isolate problems occurring on the server side. In TransactionVision you can identify which parts of a server call are causing long server response times, and see all the events that participate in a server call plus the time taken for each sub-step of the server call.

The following sections outline the End User Management drilldowns available as a result of the integration of TransactionVision with Business Availability Center:

- “Drill downs from Real User Monitor (RUM)” on page 366
- “Drill downs from Business Process Monitor Reports” on page 368

Drill downs from Real User Monitor (RUM)

In the following Real User Monitor reports there are direct links to TransactionVision reports:

Note: A GUID is an identifier used to provide a unique reference number.

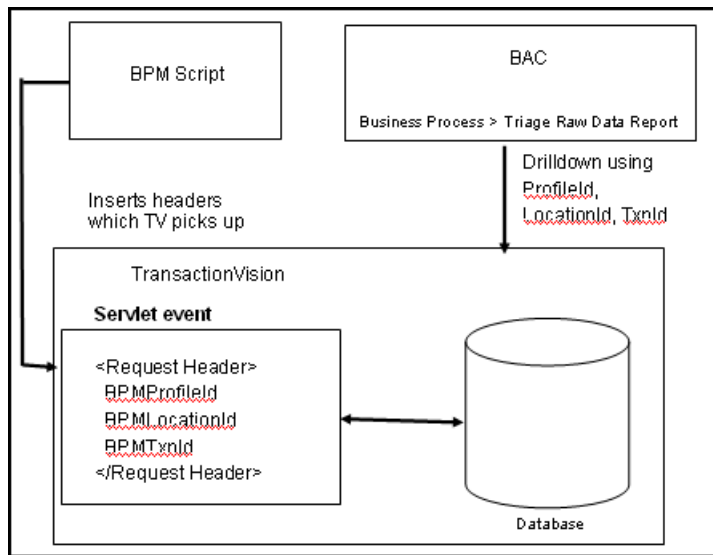
Real User Monitor Report	TransactionVision
Page Summary	Gets GUIDs from RUM for up to 10 pages, by worst server time, and drilldown to TransactionVision report.
Page Summary: End Users by Page	Gets GUIDs from RUM for up to 10 pages, by worst server time, and drilldown to TransactionVision report.
Page Summary: Page over time (only for the first page applied to filter)	Gets GUIDs from RUM for up to 10 pages, by worst server time, and drilldown to TransactionVision report.

Real User Monitor Report	TransactionVision
Page Summary: Pages by Server Summary	Gets GUIDs from RUM for up to 10 pages, by worst server time, and drilldown to TransactionVision report.
Page Summary: Servers by Page Summary	Gets GUIDs from RUM for up to 10 pages, by worst server time, and drilldown to TransactionVision report.
Session Analyzer	Drilldown based on session ID + GUID of one of the pages.
Session Details	Drilldown based on a specific page instance according to GUID and start time.

Drill downs from Business Process Monitor Reports

The user can drill down, in context, from the BPM Triage Raw Data report to the TransactionVision Transaction Tracking report. The drilldown is based on Business Process Monitor's BPM Profile, BPM Location, and BPM Transaction.

The following diagram illustrates the data flow between Business Process Monitor and TransactionVision:



Business Process Monitor inserts a header identifying the BPM Profile, BPM Location and BPM Transaction, which is picked up by the Java Agent which reports to the TransactionVision analyzer. The analyzer persists this information in the TransactionVision database and this is used for contextual drill down from the Business Process Monitor Triage Data report to the TransactionVision Transaction Tracking Report.

HP Business Process Insight Settings

As part of the HP Business Transaction Management solution, you can manage how transactions are sent to Business Process Insight from **Admin > TransactionVision > Administration > HP Business Process Insight Settings**.

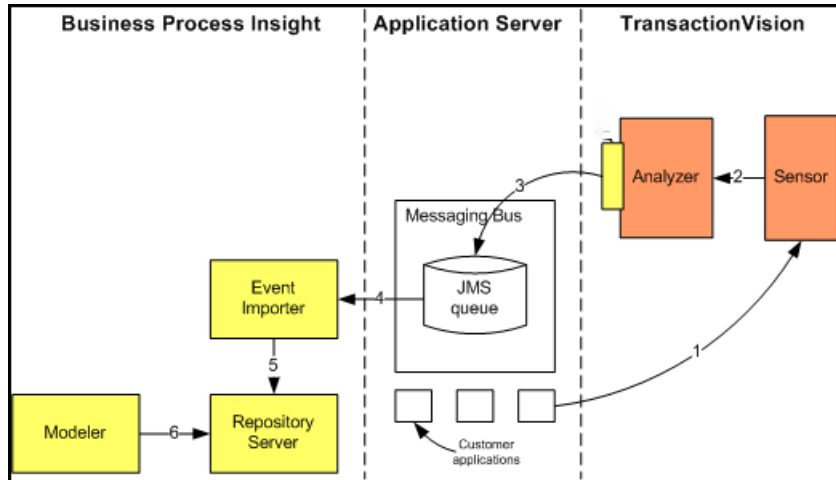
The integration between Business Process Insight and TransactionVision enables you to:

- Export TransactionVision business transaction event definitions into the Business Process Insight Model Repository. You can then use the Modeler to create processes which include Business Process Insight data.
- Progress to Business Process Insight business processes with data received from TransactionVision business transaction events.
- Link from details of TransactionVision instances, to the process instance details page within the Business Process Insight Business Process Dashboard.
- Provide the detail of TransactionVision transaction level problems to Business Process Insight.

You can define rules within TransactionVision which are specific to Business Process Insight. These rules define the business transaction events that you want made available to Business Process Insight. TransactionVision then collects these events and aggregates the data related to the events. The events and the aggregated data are placed on a JMS queue and are available for Business Process Insight to import into the Model Repository. See the task “Classify Events for Business Process Insight Integration” on page 372.

Design-Time Integration with Business Process Insight

When you design your TransactionVision processes and configure your systems to integrate with Business Process Insight, data flows in the following way:



The flow of configuration information, as shown in the diagram above, is described in the following steps:

Step 1: Transactions sourced from customer applications are imported into Transaction Vision.

Step 2: TransactionVision sensors are configured to collect transaction events from the customer applications.

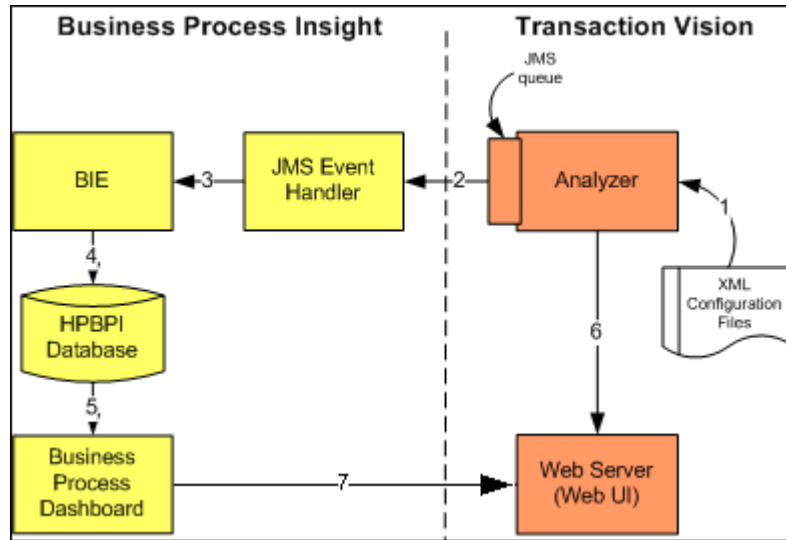
Step 3. Event and transaction rules are defined to filter the incoming transactions according to your requirements. The required transaction events are placed on the configured Business Process Insight JMS queue.

Step 4: The event importer component reads the transaction events from the JMS queue and imports them into the Model Repository, where they are available to the Modeler.

Step 5: You define your processes as usual and create Data definitions, which can then subscribe to the imported Event definitions.

Run-Time Integration with Business Process Insight

When your system is operational, data is processed between TransactionVision and Business Process Insight in the following way:



The flow of data, as shown in the diagram above, is described in the following steps:

Step 1: Using the rules defined for the business transaction events, TransactionVision filters the appropriate events and places these events on a JMS queue that you have defined. See “Classify Events for Business Process Insight Integration” on page 372.

Step 2: The JMS Business Event Handler is listening for new events that are arriving on the queue.

Step 3: The JMS Business Event Handler processes the transaction events and sends them to the Business Impact Engine for processing.

Step 4: Details of the process and its progression are held in the Business Process Insight database as usual.

Step 5: The Business Process Dashboard presents the results of the process progression.

Step 6: Details of the TransactionVision events are displayed within the TransactionVision user interface.

Step 7: The Business Process Dashboard provides links to the appropriate business transaction events within the TransactionVision user interface.

Classify Events for Business Process Insight Integration

This task describes how to classify TransactionVision events to deliver transaction data to Business Process Insight. Events must be classified in order to label classes, and rules must be defined for extracting the data from the events.

For details on the Transaction Definition Editor, see the chapter “Transaction Definition Editor” on page 283.

Open the Transaction Definition Editor

Select **Admin > TransactionVision > Administration > Transaction Definition Editor**.

Select the Transaction Class for the Event

The Transaction Definition Editor opens to display all the transaction rules defined on your system. Click on the transaction class for the event that you wish to send to Business Process Insight and the details for that class display.

TransactionVision reports to Business Process Insight all transactional attributes related to this class. For information on defining attributes and their content see the chapter “Transaction Definition Editor” on page 283.

Add New BPI Action

In order for a TransactionVision event to communicate with Business Process Insight you must define a BPI action. In the Action section click **New BPI Action**.

For every BPI action, an event is only sent if the transaction that the event is a part of, belongs to a matching transaction class. When creating a BPI action within a *classify section*, the transaction class is implied and cannot be changed. When Creating a BPI action within a *common section rule*, specifying the transaction class allows you to filter so that the action only fires when the transaction it is a part of is classified as a certain transaction class.

Within a Common Section, if the match condition for the BPI action occurs before the transaction is classified, TransactionVision delays sending the event to Business Process Insight until the transaction is classified. If the transaction is subsequently classified to the transaction class specified in the rule, an event will be sent to Business Process Insight. If it does not match that transaction class, the event is discarded. This allows the creation of rules for TransactionVision events that occur before transactions are classified, yet still report the events to Business Process Insight as the correct transactions, or to generate different BPI events depending on which transaction class the transaction ends up assigned to.

This guarantees that all BPI events generated by TransactionVision reference a valid transaction class. If it is necessary to send a BPI event regardless of its classification state, create the Action then manually edit its properties.

Format of the BPI Action Reason String

A transaction class is associated with every BPI action through the reason code of the action. The Class Name, BPI Event Group and BPI Event Name are stored in the **Reason** value of a BPI action in the form of
`<Classname>,<eventgroup>,<eventname>,....`

To specify *no transaction class* leave the `<Classname>` field empty. For example ,`<mygroup>,<myevent>,...`

If you change the name of a Transaction class you need to manually update the `<Classname>` field. Optionally you can list additional transaction document values for Business Process Insight to watch for, adding them after the first three fields.

By default Business Process Insight generates data definitions for all transaction fields. By specifying individual fields you filter out values that are of no interest. For example, using a reason code of `MyClass,BPIGroup,BPIName,Status,Value` causes Business Process Insight to only care about the *Status* and *Value* field values of the transaction when it consumes an event sent by TransactionVision.

Complete Create New BPI Action Properties Fields

When you create a new action, you need to specify the Business Process Insight event name and event group that you have defined within the BPI Modeler for the TransactionVision process.

You also need to specify the Transaction Class this BPI action is associated with. If you are creating the BPI action within a classify section, the transaction class is implied - that of the currently edited transaction class definition. If you create a BPI event within a common section rule, you need to specify which transaction class it is tied to. Complete the fields and click **OK**.

Handling of Date/Time Based Data sent to Business Process Insight

If you have a Date/Time value computed within TransactionVision, have assigned it to a custom attribute, and wish Business Process Insight to be able to read that attribute, the following steps are required:

- 1** When defining the attribute in the XDM column definition in TransactionVision, the field must be declared as a **TIMESTAMP** column.
- 2** The timestamp column must be populated by either:
 - a** setting the transaction attribute to a long value (using the standard milliseconds since January 1, 1970, 00:00:00 GMT),
 - b** or by parsing a date. In this case you need to specify the date format in the XDM definition.

Example of setting a long value

If no *dateFormat* *<Param>* is specified, the field must be set using a long value (milliseconds since Jan 1 1970).

```
<Column name="order_date" type="TIMESTAMP" description="OrderDate"
conversionType="Date">
```

```
    <Path>/Transaction/OrderDate</Path>
```

```
</Column>
```

Example of parsing a date

Setting this field in the XDM definition file only allows input in the specified *mm/dd/yyyy hh:mm* format.

```
<Column name="order_date" type="TIMESTAMP" description="OrderDate"
conversionType="Date">
```

```
    <Path>/Transaction/OrderDate</Path>
```

```
    <Param name="dateFormat" value="mm/dd/yyyy hh:mm"/>
```

```
</Column>
```

HP Business Settings User Interface

This section describes:

- HP Business Process Insight Configuration Page on page 376
- Export Business Process Insight Definitions on page 378
- HP Business Availability Settings Page on page 379

HP Business Process Insight Configuration Page

Description	Enables customizing settings for sending TransactionVision events to Business Process Insight (BPI) from this screen. To access: Select Admin > TransactionVision > Administration > HP Business Process Insight Settings .
Important Information	<ul style="list-style-type: none">➤ If using TransactionVision Sonic MQ to send events to BPI, the queue and corresponding JNDI mapping (TV2BPI.EVENT.QUEUE) are created by the TransactionVision install.➤ See the section “Transaction Definition Editor” on page 283, for more information about transactions, events, rules and actions.
Included in Tasks	“Classify Events for Business Process Insight Integration” on page 372

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Cancel	Click to cancel.
Export BPI Definitions	Click to generate these definitions within the BPI Modeler.

GUI Element	Description
Save Settings	Click to save the settings entered.
Select Analyzer to Update	From the drop-down list, select the analyzer you wish to update with BPI settings.

JMS Settings

GUI Element	Description
BPI Event Expiry (seconds)	Time that delayed BPI events are kept before being purged. BPI events are delayed if their transaction class is unclassified, or due to messaging failures.
BPI Event Thread Wait Interval (seconds)	Time that the analyzer periodically checks to clean up expired BPI events and retries any events that failed.
Enable Sending to BPI	Check to enable the sending of JMS settings to BPI.
JMS Connection Factory Name	JNDI name of the queue connection factory. Choose from BPIQueueFactory or BPISSLQueueFactory if using TransactionVision SoniqMQ.
JNDI Factory Class	Use com.sonicsw.jndi.mfcontext.MFContextFactory if using TransactionVision SoniqMQ. For other messaging providers, refer to provider documentation.
JNDI Provider Location	Location of the JNDI store. If using TransactionVision SoniqMQ use tcp://<analyzerhost>:21110
JNDI Queue Name	Enter the JNDI name of the destination to send BPI events. If using TransactionVision SoniqMQ use TV2BPI.EVENT.QUEUE.
Password	Enter password for authentication. Leave blank if not required.
Username	Enter username for authentication. Leave blank if not required.

 **Export Business Process Insight Definitions**

Description	Enables generating definitions within the BPI Modeler to correspond to the transaction rules and content defined within TransactionVision. To access: Select Admin > TransactionVision > Administration > HP Business Process Insight Settings > Export BPI Definitions .
Important Information	<ul style="list-style-type: none">➤ Any time you change your transaction definitions, either through the classification rules (adding new BPI actions), or by defining new elements in your transaction document (by changing the transaction xdm files), you should export this information to BPI to ensure that it is kept up to date with these definitions.➤ This exporting process can be done through the user interface, or by running the TVHOME/bin/bpitveventimporter utility.

The following elements are included:

GUI Element	Description
Select BPI Server to connect to	
BPI Server host	Enter the full domain name of the BPI server to connect to.
BPI Server port	Enter the port number of the RMI registry on the BPI server. The default is 44000.

HP Business Availability Settings Page

Description	<p>Enables you to customize settings for configuring TransactionVision to work with Business Availability Center.</p> <p>To access: Admin > TransactionVision > Administration > HP Business Availability Settings</p>
--------------------	--

The following elements are included:

GUI Element	Description
Cancel	Click to cancel the configuration.
Finish	Click to save the configuration.
Refresh	To verify that samples are being delivered to Business Availability Center, click this button, then check on the updated status and messages output by this job.
Run Now	Click to run the selected job immediately.
Start Job	Click to start the selected job.
Stop Job	Click to stop the job for this project.
Synchronize CMDB	Click to synchronize transaction classes with the CMDB. A results screen lets you know the status.

TransactionVision Data Sample Job Configuration

GUI Element	Description
Collect In-Process KPIs for last N seconds only	By default In-Process KPIs are collected without time constraints. Use this option to control what time range should apply. 0 means collect In-Process KPIs for all time ranges.
Customer for Sending Samples	Set the Business Availability Center customer name for whom the samples should be sent to Business Availability Center.

GUI Element	Description
Retry sending samples for last N seconds only	By default TransactionVision retries sending samples to Business Availability Center if there are network connection failures. Use this option to control the aging threshold for the number of samples sent to Business Availability Center once a network connection is restored. 0 means send all samples.
Sample Interval (in seconds)	This time interval is set to 300 seconds, which is the optimum value. It cannot be modified.
Time slice shift in job (in seconds)	Set the time slice shift. This is the interval by which the job that sends statistics, lags in calculating statistics. The default value is set to 360 seconds.

Drill-down Project

GUI Element	Description
Select Project for Drill-down from End User Monitor Reports	Select a project from the drop-down list for drilling down to TransactionVision from the Business Availability Center End User Monitor reports.

Gateway Server

GUI Element	Description
Gateway Host name	Enter the hostname for the Business Availability Center Gateway server.
Gateway Port	Enter the port for the Business Availability Center Gateway server.
Gateway Protocol	Specify whether HTTP or HTTPS is to be used to connect to this server.

Project Name

GUI Element	Description
Job Description	The description given for the job.
Job Interval	The sample interval set for the job.
Job Name	The name of the Job.
Job Startup Mode	What the startup mode of the Job is.
Job Status	The present status of the job (Running, Waiting, Stopped).
Time of next activation	Exact time of next activation recorded here.

Troubleshooting and Limitations

This section includes the following topics:

“Click Stream must be enabled” on page 382

“Time zone on target applications must be set correctly” on page 382

“Samples not delivered to Business Availability Center” on page 382

“Debugging the receiving of samples” on page 382

“Debugging the BAC Aggregate TV Sample Data” on page 383

“Sharing Transaction Classes across Schemas” on page 383

“Transaction End Time updated after Transaction Completed” on page 383

“Syncing Business Transactions with uCMDB” on page 384

“Incorrect Sample Counts - Classify Rule Problems” on page 384

“KPI Values Showing No Data” on page 384

Click Stream must be enabled

For Real User Monitor-TransactionVision drilldowns the 'click stream' must be enabled. For performance reasons this can be disabled and Business Availability Center can be configured to store only a very limited amount of data. Other Business Availability Center functionality does not work when this is disabled, so HP cautions against this.

Time zone on target applications must be set correctly

Make sure all components (for example TransactionVision, Real User Monitor, Business Process Monitor) are time synched with Business Availability Center, or the data might not be consistent across the different components.

Samples not delivered to Business Availability Center

When TransactionVision samples are not delivered to Business Availability Center check if:

- The Business Availability Center host name is valid.
- The HTTP post is receiving a status code other than 200.
- Business Availability Center is ignoring samples due to different formats etc.

If you still do not see TransactionVision samples in the Business Availability Center Dashboard, you can debug further by turning on tracing:

- Turn on Business Availability Center sample job's debug on the TransactionVision side, and make sure that samples are being sent successfully.
- Turn on Business Availability Center's trinity log, make sure no samples are ignored by Business Availability Center.

Debugging the receiving of samples

To debug when Business Availability Center is not receiving samples:

- Turn WDE debug on by going to:
`{HPBAC DIR}\conf\core\Tools\log4j\mercury_wde\wde.properties`

- Search loglevel and set loglevel=DEBUG.

See the logs in {HPBAC DIR}\log\mercury_wde:

- **wdeIgnoredSamples.log** (is filled in any mode)
- **wdePublishedSamples.log** (is filled only in DEBUG mode)
- **wde.log**
- **wde.all.log**

Debugging the BAC Aggregate TV Sample Data

To debug the BAC Aggregate TV Sample Data:

- 1 Stop the BAC Aggregate TV Sample Data job.
- 2 Edit the job by inputting -debug true parameter.
- 3 Restart the job.
- 4 Once the debug flag is enabled, debug information is written to **ui_job.log** file in {TVISION_HOME}/logs directory.
- 5 The debug information includes the Business Availability Center time server's time stamp, the time stamps used by queries, and the SQL themselves, etc. It also contains the HTTP status code.

Sharing Transaction Classes across Schemas

Transaction classes should not be shared across multiple schemas. If transaction classes are shared, the TransactionVision Aggregate and Instance Topologies in Business Availability Center, the Business Transaction reports, and the Business Availability Center drilldowns for business transactions, may not work reliably.

Transaction End Time updated after Transaction Completed

Whenever the end time of a transaction is updated after the transaction has been marked as completed, incorrect sample counts can be reported to Business Availability Center.

Syncing Business Transactions with uCMDB

This command line Diagnostic Utility is mainly used for debugging. Use the following commands:

-delete	Clears out all Business Transaction and Monitor CIs in CMDB.
-querycmdb	Displays a list of all Business Transaction and Monitor CIs:
-querytv	Displays a list of all Transaction Classes defined in TransactionVision.
-querydiff	Displays the data that is not in sync between TransactionVision and CMDB.
-sync	Synchronizes all Transaction Class data between TransactionVision and CMDB.

Example:

```
PublishTransactionsToCMDB.bat -h ovrntt124.ovrtest.adapps.hp.com -p 8080  
-u admin -p admin -a sync
```

Incorrect Sample Counts - Classify Rule Problems

If the classify rules for a transaction class are too general such that different parts of a transaction can be classified while the transaction is still being processed, it is possible that incorrect sample counts can be sent to the Business Availability Center Dashboard. To avoid this, make the classify rules for a transaction class specific enough so that they do not match multiple events for a transaction.

KPI Values Showing No Data

If using a custom class to set the value of a transaction, some KPI values may show a value of **No Data** in the Business Availability Center Dashboard. This can happen if the sample job checks to see if a value is defined before reporting the value field in the BAC sample data. As a workaround, the classification XML should be changed such that the /Transaction/Value attribute is defined.

13

Additional Reference Information

This chapter includes:

- Administration Utilities on page 385
- APIs Monitored by Sensors and Agents on page 413
- Technical References for Sensor and Agent Technologies on page 421
- Troubleshooting on page 421

Administration Utilities

The following utilities can be run from the command-line on the host on which the Analyzer is running.

CreateSqlScript

Location:

TVISION_HOME/bin/CreateSqlScript.[sh|bat]

Purpose:

To create and optionally execute an SQL script to create, drop, import or export a TransactionVision system table or project table.

Syntax:

```
CreateSqlScript
{-create(-c) | -drop(-d) | -import(-i) | -export(-ex) | -grant(-g) | -dbstats(-ds)}
{-system(-sys) | -schema(-s) SCHEMA | -table(-t) TABLE SCHEMA}
[[-noscript(-n)] -execute(-e)] [-noprompt(-np)]
[-noinsert(-ni)]
[-tablespace(-ts) TABLESPACE] [-dbMove(-m)]
[-fileType(-f) IXF|DEL] [-lobPath(-lp) PATH]
[-noLob(nl)] [-dbproperties(-db) FILE]
```

Options:

Option	Description
-drop (-d)	Drop tables
-create (-c)	Create tables
-execute (-e)	Execute script
-noscript (-n)	No script generation
-noinsert (-ni)	Do not insert initial table rows
-system (-sys)	Create script for system tables
-schema (-s) SCHEMA	Create/drop project tables in schema SCHEMA
-table (-t) TABLE SCHEMA	Create/drop table TABLE in schema SCHEMA
-dbMove (-m)	Use db2move to import/export data
-tablespace (-ts) TABLESPACE	Use tablespace TABLESPACE. See “Usage Notes: ” on page 388 for information about using this option with an Oracle database.
-dbproperties (-db) FILE	Use Database.properties file FILE.
-dbstats (-ds)	Generates a script for generating the database statistics used by the database optimizer.

Option	Description
-import (-i)	Generates a database import script. To run database scripts, use the command <code>db2 -n -t -f <sql script filename></code> . For a DB2 database, you can combine this option with the <code>-fileType</code> and <code>-lobPath</code> options to customize the data format and the location of LOB. You may NOT combine this option with the <code>-noscript</code> or <code>-execute</code> options.
-export (-ex)	Generates a database export script. To run database scripts, use the command <code>db2 -n -t -f <sql script filename></code> . For a DB2 database, you can combine this option with the <code>-fileType</code> and <code>-lobPath</code> options to customize the data format and the location of LOB. You may NOT combine this option with the <code>-noscript</code> or <code>-execute</code> options.
-fileType (-f) IXF DEL	Specify the DB2 data output file format used for importing/exporting data. The default type is IXF. For the IXF file type, the import/export script uses the LOBFILE option for rows that contain greater than 32K data. For the DEL type, the import/export script exports LOBFILES into a single file (requires FixPack 8).
-lobPath (-lp) PATH	Specifies the directory for DB2 LOBFILES. The default value is the current directory.
-grant (-g)	Generates a script for granting all database privileges required for accessing the tables of the schema to the database user specified in Database.properties (Oracle and DB2 only)
-noLob (-nl)	Do not generate DB2 export/import SQL with LOBINFILE option. This option truncates LOB data to the first 32K bytes.
-noprompt (-np)	Do not prompt for confirmation when dropping tables

Usage Notes:

Note: The TransactionVision Analyzer and UI/Job Server need to be stopped before performing any database imports or exports. This is to avoid causing the database import/exports to fail because of database locks held by the Analyzer or Web application. If a schema is dropped, make sure that there are no active projects or users logged in that are using that schema.

When the **-dbstats** option is specified, the generated SQL script updates statistics about the physical characteristics of a table and the associated indexes. These characteristics include metrics like number of records, number of pages, average record length, and so forth. The optimizer uses these statistics when determining access paths to the data.

This SQL script should get executed when a table has had many updates, such as when data is continuously collected into the database by the TransactionVision Analyzer.

It could result in large performance gains in queries made by TransactionVision views and reports, as well as queries made internally by the TransactionVision Analyzer to correlate events.

Alternatively, you can use the database vendor tools (Oracle Enterprise Manager, DB2 Control Center) to regenerate the database statistics.

Note: While the database statistics are being generated the TransactionVision Analyzer processing slows down.

SQL Server does not have the concept of “table spaces,” but the “file group” is a similar concept. Though project tables can be distributed to different file groups by using custom DDL scripts, the creation of project tables in specific file groups through the TransactionVision UI/Job Server or CreateSqlScript is not supported.

Examples:

Create system tables with schema as TVISION and execute the procedure without generating SQL script:

```
CreateSqlScript -e -n -c -sys
```

Generate SQL script for creating project tables with schema as PROJECT without executing the procedure:

```
CreateSqlScript -c -s PROJECT
```

Drop table EVENT in schema PROJECT and execute the procedure without generating SQL script:

```
CreateSqlScript -e -n -d -t EVENT PROJECT
```

DataUtil**Location**

TVISION_HOME/bin/DataUtil.[sh|bat]

Purpose

Offers various functionality to display and manipulate TransactionVision data stored in the database, and is mainly used for diagnostics purposes. This utility obtains database connection information from the TransactionVision **Database.properties** file.

Syntax

```
DataUtil -schema(-s) schema [options]
```

Options

Option	Description
-schema	The database schema.
-localtxn_list(-ltl) pii,seq btxnid	Prints data for all local transactions of a business transaction.
-localtxn_ids(-ltids) pii,seq btxnid	Prints the ids of all local transactions for a business transaction.
-businesstxn_ids(-btids) pii,seq ltxnid	Prints the business transaction id for a local transaction.
-localtxn_events(-ltev) pii,seq ltxnid	Prints data for all events of a local transaction.
-businesstxn_events(-btev) pii,seq btxnid	Prints data for all events of a business transaction.
-localtxn_data(-ltd) pii,seq ltxnid	Prints data for a local transaction.
-businesstxn_data(-btd) pii,seq btxnid	Prints data for a business transaction.
-correlated_events(-ce) pii,seq seqid	Prints out all correlated events for an event.
-eventdetail(-ed) pii,seq seqid {-conv}	Prints out the event XML document for an event (-conv translates object ids into object names).
-userdata(-ud) pii,seq seqid	Prints out the user data for an event.
-sysmodel_object(-smo) objid	Prints the data for a system model object.
-eventid(-eid) seqid	Prints proginst_id and sequence_no for a given seq_id.
-seqid(-sid) pii,seq	Prints seq_id for a given proginst_id and sequence_no.
-counts(-c)	Prints transaction counts.

Option	Description
-row_counts(-rc)	Prints the row count for all project tables.
-list_schemas(-ls)	Prints out all schema names registered within TransactionVision.
-show_ids(-si)	Prints schema and commlink IDs.
-sql QUERYFILE	Prints out the generated SQL for an XML query document.
-query QUERYFILE	Executes the query specified by the XML query document and returns the resulting <EventList> document.
-clear_txndef	Clears all transaction class and classification data in the database.
-import_txndef INFILE [CLASS_XDM_DIR]	Imports the classes and classification rules contained in the legacy classification rule file INFILE. If an optional XDMDIR is specified, the class attribute values from the XDM files are also imported.
-export_txndef [OUTFILE]	Exports the transaction classes and classification rules to a legacy classification rule file (all schemas).
-export_schema_txndef SCHEMA [OUTFILE]	Exports the transaction classes and classification rules for a specific schema to a legacy classification rule file.
-import_classxdm XDM_DIR	Imports all class attribute values from the transaction class XDM files located in XDM_DIR.
-export_classxdm SCHEMA	Exports all class attribute values for a specific schema to a legacy TransactionClass.xdm file.
-pii	Proginst_id a of an event.
-seq	Sequence_no (primary key) of an event.

Option	Description
-ltxnid	Local_trans_id of a local transaction.
-btxnid	Business_trans_id of a business transaction.

DeleteEvents

Location:

TVISION_HOME/bin/DeleteEvents.[sh|bat]

Purpose:

Allow the user to delete events in EVENT tables and clean up contents in other project tables related to deleted events.

The utility (by default) does only delete the event data. In order to clean up other related table data like system model objects and transaction data, use the options -deletePiis and -deleteTxns described below. If you have deleted event data without these options, you can clean up the remaining data by running the DeleteEvents utility again with the option -cleanup.

Syntax:

```

DeleteEvents
-query QUERY -project PROJECT
  [-keepPiis] [-keepTxns] [-interationcount EVENTS]
  [-commit COUNT] [-timeout MINUTES][-force]
  [loglevel LOGLEVEL]

-from TIME -to TIME | -older TIME | -newer TIME] [-minage HOURS]
  [-project PROJECT | -schema SCHEMA]
  [-timezone TZID][-businesstxns[-state STATE]][-result RESULT]]
  [-keepPiis] [-keepTxns] [-keepLocalTxns]
  [-deleteStats [topology][business]][-commit COUNT]
  [-timeout MINUTES] [-force] [-loglevel LOGLEVEL]
  [-threadcount THREADS] [-nosplit]

-eventCountLimit MAXCOUNT
  [-businesstxns(-bt) [-state STATE] [-result RESULT]]
  [-keepPiis(-kp)] [-keepTxns(-kt)] [-keepLocalTxns(-klt)]
  [-deleteStats(-ds) [topology(t)] [business(b)]]
  [-threadcount(-tc) THREADS] [-nosplit]

-event pii seqno
-btxn businessTxnId

-cleanup
  [-sysmodel][-transactions]
  [-commit COUNT][-timeout MINUTES][-force]
  [-loglevel LOGLEVEL]

```

Options:

Option	Description
-keepLocalTxns(-klt)	Do not delete local transactions from the transaction tables. (DEFAULT: keep only business transactions, only applicable if using -keepTxns).
-event PII SEQNO	Delete a single event.
-btxn TXNID	Delete a single business transaction.
-project (-p) PROJECT	Specifies the project name.
-query (-q) QUERY	Specifies the name of the query.

Option	Description
-commit(-c) COUNT	Commit database transactions after the specified COUNT. The default interval is 1000.
-from(-f) STARTTIME	Specifies the start time (Time format = mm/dd/yyyy hh:mm:ss.SSS, event time in the time zone you set (your local time zone by default) or, for relative time, n[hr min])
-to(-t) ENDTIME	Specifies the end time (Time format = mm/dd/yyyy hh:mm:ss.SSS, event time (your local time zone by default) or, for relative time, n[hr min])
-older (-o) TIME	Deletes events older than the specified time.
-newer (-n) TIME	Deletes events newer than the specified time.
-eventCountLimit (-ecl) MAXCOUNT	Deletes all but the latest MAXCOUNT events.
-minage (-m) HOURS	By default, the DeleteEvents utility does not allow deleting events from the last 48 hours. Use this option to override the default. No events from the last number of hours specified by the HOURS value is to be deleted.
-schema(-s) SCHEMA	Database schema for where events are stored.
-businesstxns (-bt) [-state STATE] [-result RESULT]	Deletes only events of transactions with the specified transaction state or result. State values: -1=UNKNOWN 0=PROCESSING 1=COMPLETED Result values: -1=UNKNOWN 0=FAILED 1=SUCCESS
-deleteStats (-ds) [topology(t)] [business(b)]	Deletes statistics data from the topology or business transaction statistics table. This is optional. The time period specifies the time slices for which statistics are deleted.
-timezone (-z) TZID	Timezone ID (for example, GMT, EST, etc.)

Option	Description
-help(-h)	Displays the usage message.
-keepPiis (-kp)	Do not delete program instance or CICS task objects from the system model table (DEFAULT: off; objects are deleted).
-keepTxns (-kt)	Do not delete transactions from the transaction tables (DEFAULT: off; transactions are deleted).
-cleanup (-cl) -sysmodel (-sys) -transactions (-txns)	Cleans up the system model tables (removes all program instance objects which are no longer referenced) or transaction tables (removes all transaction objects which are no longer referenced). You must specify either system model tables or transactions.
-iterationcount (-it) EVENTS	Maximum number of events to process before cleaning out the system model and transaction tables. The default is 1000.
-timeout (-ti) MINUTES	Force deletion to end after the specified time. The default is 60 minutes.
-force (-fc)	Force deletion to stop immediately after a cancel or timeout. This option is always in effect in timestamp deletion mode (for example, with the -older option).
-loglevel (-log) LOGLEVEL	Log output level. The default is 0.
-threadcount (-tc) THREADS	Number of deletion threads (only applicable if using -older). The default is 4.
-nosplit	Disables deletion of the same table by multiple concurrent threads.

Usage Notes:

This utility operates in three modes: deleting by query mode, deleting by time mode, and cleaning up tables. Therefore, the options are grouped into three sets.

- Deleting by time options include: -from, -to, -older, -newer, -timezone, -keepPiis, -keepTxns, -deleteStats, -businessstxns, -state, -result (and -threadcount, -nosplit if using -older). The -force option is always in effect when deleting by time.
- Deleting by query options include: -query, -keepPiis, -keepTxns
- The cleanup mode options include: -cleanup, -sysmodel, -transactions

Options that can be used in all three modes include: -c, -ti, -fc, -log. Otherwise, these sets of options may not be mixed. For example, the following commands are INVALID and result in errors:

```
DeleteEvents -project MYPROJECT -query MYQUERY -business
DeleteEvents -project MYPROJECT -query MYQUERY -older 05/15/2003
03:00:00.000
```

Execution of the DeleteEvents utility or the Deletion Job in the UI while the Analyzer is processing events is only allowed for the -older and -eventCountLimit modes in which the Analyzer executes specific code to safeguard against accidental deletion of data in process. But due to the negative impact of deletion on the Analyzer performance it is still not recommended to delete data while the Analyzer is processing. However, sometimes it may not be feasible to shut down the Analyzer while DeleteEvents is running. In that case, make sure that you do not delete transactions for which events are likely to come in.

On many systems, deleting events from 2 days to a week in the past can ensure this scenario does not occur. However, you need to consider the duration of your business transactions to determine a period beyond which events can be deleted. For example, if your business transactions can span a day, do not delete events in the last day, if your business transactions can span a week, do not delete events in the last week while the Analyzer is running. The utility (by default) does not delete any events that are younger than 48 hours. If you are sure that no transaction duration ever exceeds a certain value, you can override this behavior by specifying the option `-minage n`, where `n` is the time buffer to exclude in hours.

Deleting by time mode provides better performance results than deleting by query. The performance of the deletion options can be ordered in the following way:

- 1** Deletion by query (low performance)
- 2** Deletion by time with options `-from/to` or `-newer` (medium performance)
- 3** Deletion based on event count (medium-high performance)
- 4** Deletion by time with option `-older` (high performance)

The deletion by time with option `-older` is able to take advantage of timestamp columns present in the tables and achieves much higher performance than the other methods. Also, it can execute the deletion process with multiple threads (option `-threadcount`), where each thread deletes different portions of the tables at the same time. The default number of threads is 4. You should adjust this option to the number of CPUs (or higher) in your system.

The deletion based on event count (`-eventCountLimit`) has similar performance characteristics as the deletion by time. Internally, the deletion module determines the date from which to delete based on the event count parameter, and then performs a `-older` deletion.

The following conditions may also affect the performance of the `DeleteEvents` utility:

- Make sure your database log space is sufficient, and adjust the number of uncommitted database operations (option `-commit`) accordingly; otherwise, the deletion process fails running out of logspace. Usually a larger commit count yields better deletion performance.

- On DB2, you could encounter deadlocks if the deletion is using multiple threads and your database lock storage is low. In that case, please adjust the value of LOCKLIST and MAXLOCKS accordingly. On Oracle this is of no concern because Oracle never escalates row locks to table locks. Alternately, use the option -nosplit to DeleteEvents. This disables the concurrent deletion of a single table by multiple threads.

Time formats may be absolute format or relative format. Absolute times are specified as mm/dd/yyyy hh:mm:ss.SSS. Relative times are specified as n[hr|min]. Relative time is offset to the time you run the script. For example “12 hr” means “12 hours before the current time.” You can specify the time in any of the following ways. The descriptions assume that the script was run at 9:00 a.m.

Example	Description
-f 2 hr -t 1 hr	Deletes events from 7:00 a.m. to 8:00 a.m.
-o 1 hr	Deletes events that came into the system before 8:00 a.m.
-n 1 hr	Deletes events within the last hour.

Examples:

- 1 Delete everything (event and related data) from project PROJECT which is older than 72 hours using 2 threads and commit database transactions after 10000 operations:

DeleteEvents -project PROJECT -older 72 hr -kp -kt
-ds b t -tc 2 -c 10000
- 2 Delete events from project PROJECT using query QUERY and commit the database transaction after deletion of every 10 events:

DeleteEvents -project PROJECT -query QUERY -c 10
- 3 Delete events from GMT time 11:00AM, July 10, 2003 to GMT time 8:00PM, July 11, 2003:

DeleteEvents -from 07/10/2003 11:00:00.000 -to 07/11/2003 20:00:00.000 -z GMT

- 4 Delete events collected before 10 minutes together with statistic data from topology table:

DeleteEvents -older 1 hr -minage 1 -statistics t

- 5 Delete all but the latest 500,000 events in the project:

DeleteEvents -project PROJECT -eventCountLimit 500000

FlushStatusUtil

Location:

TVISION_HOME/bin/FlushStatusUtil.[sh|bat]

Purpose:

Shows and cleans the flush status of transaction analysis data and cached statistics for the Static Component Topology Analysis view. The status is stored in the database and determines whether recovery is needed when the Analyzer starts. This utility can be used to force Analyzer startup without performing recovery by cleaning the flush status before starting the Analyzer.

Syntax:

```
FlushStatusUtil -schema SCHEMA [-clean]
```

Options:

Option	Description
-schema	The name of the database schema to show or clean the flush status for.
-clean	Clean the current flush status of the analysis data and statistics (default action is 'show'). The Analyzer starts up in a "clean" state without recovery, even though data might be missing in the database. Use this option with extreme caution.

Usage Notes:

Shows and cleans the flush status of transaction analysis data and cached statistics for the Static Component Topology Analysis view. The status is stored in the database and determines whether recovery is needed when the Analyzer starts.

For cached statistics, the status contains the IDs of the events for which statistics and analysis data are missing in the database.

This utility obtains database connection information from the TransactionVision Database.properties file.

Example Output:

The following output indicates that IDs 7916-8084 and the set of additional IDs are missing statistics and analysis data and need to be recovered at the next startup time:

Last written sequential_id in EVENT_LOOKUP: 8084

DBCACHE flush status:

Last flushed sequential_id: 7915

Additional ids not flushed: 7912,7911,7901,7892,7870,7872

STATISTICS flush status:

Last flushed sequential_id: 7915

Additional ids not flushed: 7912,7911,7901,7892,7870,7872

The following output indicates that all statistics and analysis data has been flushed; no recovery is needed:

Last written sequential_id in EVENT_LOOKUP: 9640

DBCACHE flush status:

Last flushed sequential_id: 9640

STATISTICS flush status:

Last flushed sequential_id: 9640

ManageQueue**Location**

TVISION_HOME/bin/ManageQueue.[sh|bat]

Purpose:

Use this multipurpose tool to manipulate and maintain TransactionVision event queues. This utility allows you to copy, move, and delete messages (usually TransactionVision related, but not limited to) from a queue. The messages can be transferred from the queue to another queue or archived into a jar file that can later be extracted and copied into a queue.

This utility also provides a means to query project related database IDs and summarize the contents of an event queue (or other queues).

Syntax:

ManageQueue <Action Options> [Queue Connection Options] [Input/Output Options]
[Message Selection Options] [Miscellaneous Options]

Options:

Options	Description
Action Options	
-show_db_ids	Display all schema and communication link IDs from the database defined in Database.properties
-show_event_ids	Display all TransactionVision event IDs from a given Input
-copy [SCHEMA COMMLINK]	Copy messages from Input to Output using Message Selection Options. Optionally targets the messages to a new SCHEMA and COMMLINK. If no SCHEMA and COMMLINK are specified, all messages are copied using their original message IDs.
-move [SCHEMA COMMLINK]	Move messages from Input to Output using Message Selection Options. Optionally targets the messages to a new SCHEMA and COMMLINK. If no SCHEMA and COMMLINK are specified, all messages are moved using their original message IDs.
-delete	Deletes messages from Input using Message Selection Options.
Queue Connection Options	
-qm QUEUEMANAGER	Queue manager to connect to
-host (-h) HOST	Host name for client connection to queue manager
-port (-p) PORT	Listener port for client connection to queue manager
-channel (-c) CHANNEL	Client connection channel to queue manager
-username USERNAME	Username to connect with
-password PASSWORD	Password to connect with
Input/Output Options	

Options	Description
-input_file (-if) FILE1 [FILE2...]	Input files
-input_queue (-iq) QUEUE1 [QUEUE2...]	Input queues
-output_file (-of) FILE	Output file
-output_queue (-oq) QUEUE	Output queue
Message Selection Options	
-commlink ID NAME	All TransactionVision events targeted to a given link ID or name
-schema ID NAME	All TransactionVision events targeted to a given schema ID or name
-db_name NAME	All TransactionVision events targeted to a given database name
-db_host IP HOST	All TransactionVision events targeted to a given database host (IP or name)
-erroneous	All messages not recognized as TransactionVision events
-userevents	All messages recognized as TransactionVision User events
-alltv	All messages recognized as TransactionVision events
-otherdbs	All TransactionVision events targeted to databases other than the one defined in Database.properties
-orphaned	All TransactionVision events targeted to schema and/or communication link IDs which no longer exist
-valid	All TransactionVision events targeted to valid schema and communication link IDs
Miscellaneous Option	
-count COUNT	Process a specified limited number of messages.

Options	Description
-verbose (-v)	Print verbose output.
-help	Print usage.

Usage Notes:

The utility takes five different types of options, most of which are optional with the exception to the Action. In addition to the Actions, there are Queue Connection Options which provide information needed to connect to a given queue manager or broker if needed.

Input and Output options are used to copy, move, or delete from a queue or jar file. An output option is only needed if there is a copy or move action performed.

Input Message Selection Options are special options only used in conjunction with the -input_queue option. This provides a means for processing the Action on only a subset of the messages on the queue.

Miscellaneous Options provide additional functionality. The -verbose option shows additional information during message processing to show progress status, and displays a summary of what was processed at the end of execution. The -count option allows you to process only the first given number of events from the Input.

With the use of -show_event_ids, you can display a summary listing of all TransactionVision events on a given Input Queue. ManageQueue summarizes them by showing a count of each message and which database, schema, and communication link they are targeted for.

The -show_event_ids option also determines whether the messages are TransactionVision events with a valid destination (valid target schema and communication link on the defined database) and reports them as "Valid." Events which are targeted to a schema and/or communication link which no longer exist are shown as "Orphaned." And events targeted to another database which may or may not be valid are shown as "Other DB."

The `-show_event_ids` option also shows a count of events not recognized as TransactionVision events as “Erroneous Messages.”

Input and Output can be combinations of queues and/or jar files. The following limitations apply to Input Files:

- You cannot use `-show_event_ids` with Input Files.
- You cannot use `-move` with Input Files. Use `-copy` instead.
- You cannot use `-delete` with Input Files.

When using an Input Queue and/or Output Queue, you must provide Queue Connection Options. For WebSphere MQ, if only a queue manager is specified (`-qm`), server binding is used. For client connections, use `-host`, `-port`, and `-channel`. The `-port` option is optional as the default port is 1414 when using the `-host` and `-channel` options. For SonicMQ, specify the `-sonicmq` option along with `-host` and `-port` to the SonicMQ broker.

ManageQueue only supports connection to one queue manager or broker. If there is a need you transfer messages from a queue on one queue manager or broker to a queue on a different queue manager, use a remote queue or cluster for the destination queue or broker, or transfer the messages to an Output File. Then transfer them from the file to the destination queue.

When using an Input Queue, you can specify Message Selection Options to process only a subset of the messages on the specified Input Queue. There are two types of message selection:

➤ Discrete Selectors

- `-commlink`
- `-schema`
- `-db_name`
- `-db_host`

➤ Validity Selectors

- `-erroneous`
- `-alltv`
- `-otherdbs`
- `-orphaned`
- `-valid`

In you use any Discrete Selectors, all Validity Selectors are ignored, with the exception of `-erroneous` because it does not apply to TransactionVision events. Discrete Selectors are ANDed so all Discrete Selectors specified must match in order for the message to be processed.

If no Discrete Selectors are used, the specified Validity Selectors are ORed. All messages matching any of the specified Validity Selectors are processed.

If no selectors are used, all messages are processed.

Any option that requires the user to specify a schema and/or communication link can take either the database ID or name. Surround names containing spaces with quotation marks.

When copying or moving messages to a queue, it may be necessary to specify new targeted schema and communication link IDs or names. This is especially useful when the messages have been orphaned (the targeted schema and/or communication link does not exist anymore). If no schema and communication link is specified, each message's original event information is used.

When copying or moving messages to a jar file, new schema and communication link IDs or names are not required. If they are specified, they are ignored by `ManageQueue`.

Examples:

- Remove all messages on the event queue which are not being processed. This includes non-TransactionVision (`erroneous`) messages and events that are targeted to a schema and/or communication link which no longer exists (orphaned):

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -delete  
-erroneous -orphaned -verbose
```

- Show all messages on the event queue:

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -  
show_event_ids
```

This command results in output similar to the following:

Count	Database Name	Database Host IP	CommLink ID	Schema ID	Validity
284	TVISION	172.22.22.46	2	1	Other DB
22	TVDB	172.22.22.106	6	1	Orphaned
7	TVDB	172.22.22.106	1	1	Valid
17	TVDB	172.22.22.106	3	2	Valid

- Show all Database IDs:

```
ManageQueue.sh --show_db_ids
```

This command results in output similar to the following:

```
Schemas:
1: TVISION
2: TEST
Communication Links:
1: TVision - tv.qmgr
3: Test - TransactionVision.qmgr
```

- Connection to WebSphere MQ:

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -delete
-erroneous -orphaned -verbose
```

- Connection to TransactionVision SonicMQ:

```
ManageQueue.sh -sonicmq -host win2003 -port 21111 -input_queue
TVISION.EVENT.QUEUE -delete -erroneous -orphaned -verbose
```

- Move all TransactionVision events from the TransactionVision Exception Queue to the event queue. Target them to the TVISION schema and the TVision – TransactionVision.qmgr communication link.

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EXCEPTION.QUEUE -
output_queue TVISION.EVENT.QUEUE -move "TVISION" "TVISION -
tv.qmgr" -alltv
or
```

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EXCEPTION.QUEUE -
output_queue TVISION.EVENT.QUEUE -move 1 1 -alltv
```

- Retarget all events not being collected because they are currently targeted to a schema and/or communication link which no longer exists (orphaned):

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -  
output_queue TVISION.EVENT.QUEUE -move 1 1 -verbose  
or
```

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -  
output_queue TVISION.EVENT.QUEUE -move "TVISION" "TVision - tv.qmgr" -  
verbose
```

- Move all TransactionVision events from the TransactionVision Exception Queue to their respective event queue. This example assumes that the schema TVISION uses the communication link "TVision - tv.qmgr" which in turn uses TVISION.EVENT.QUEUE, and schema TVISION2 uses communication link "Test - tv.qmgr" which in turn uses TVISION.TEST.QUEUE.

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EXCEPTION.QUEUE -  
output_queue TVISION.EVENT.QUEUE -move -schema "TVISION" -commlink  
"TVision - tv.qmgr"
```

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EXCEPTION.QUEUE -  
output_queue TVISION.TEST.QUEUE -move -schema "TEST" -commlink "Test  
- tv.qmgr"
```

- Archive all messages on the event queue into a jar file while leaving them on the queue:

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -  
output_file events.jar -copy -verbose
```

- Replay all events stored in a jar file for the TVISION schema to the **TVision - tv.qmgr** communication link:

```
ManageQueue.sh -qm tv.qmgr -input_file events.jar -output_queue  
TVISION.EVENT.QUEUE -copy "TVISION" "TVision - tv.qmgr" -verbose  
or
```

```
ManageQueue.sh -qm tv.qmgr -input_file events.jar -output_queue  
TVISION.EVENT.QUEUE -copy 1 1 -verbose
```

ServicesManager

Location:

TVISION_HOME/bin/ServicesManager.[sh|bat]

Purpose:

Manage the TransactionVision Analyzer service. The Analyzer uses an embedded RMI registry so that Analyzers may be controlled by the TransactionVision UI/Job Server running on remote hosts.

Syntax:

```
ServicesManager
{(-start [-project (-proj) PROJECTNAME]) |
 (-stop [-quiesce][-project (-proj) PROJECTNAME]
 [-keepcollect]) |
 (-exit [-quiesce][-keepcollect]) |
 (-status [-project (-proj) PROJECTNAME]) |
 (-killserver) |
 (-reconfig {classification | logging | analyzer}) |
 (-versioninfo)
 ([-host HOST][-rmiregp PORTNUMBER][-debug])
```

Options:

Option	Description
-start	Starts the Analyzer process if it is not already running.
-stop	The Analyzer stops collecting event data.
-exit	The Analyzer stops collecting event data, then the process exits.
-quiesce	The default behavior of the Analyzer on a stop or exit is to immediately close down collection. If this flag is set, the Analyzer clears out any pending events in the event queue before stopping or exiting. Note that if there is a large event backlog, the Analyzer may take some time to stop or exit.

Option	Description
-keepcollect	Do not send stop message to Sensor. Used in combination with the -stop or -exit option so that the Sensor keeps collecting.
-status	Reports the current Analyzer status.
-reconfig	<p>Reloads Analyzer configuration settings (classification/logging/all settings) without stopping event collection.</p> <ul style="list-style-type: none"> ► logging. Reloads the log4j XML configuration files in the directory <TVISION_HOME>/config/logging. ► analyzer. Re-initializes the Analyzer by reloading most Analyzer settings in the Analyzer.properties file, all beans described in the Beans.xml file, re-initializing logging and all XML based rule files. <p>Settings in the Collection Properties section of the Analyzer.properties file, except the write_to_buffer_table property, are not reloaded by the Analyzer with the -reconfig option.</p>
-versioninfo	Returns Analyzer version information.
-killserver	Shuts down the Analyzer immediately. This flag is not recommended; -exit is the preferred method for shutting down the Analyzer cleanly.
-project (-proj) PROJECTNAME	Name of the project for the specified command. If the project name contains a space, enclose the project name in double quotation marks (for example, -proj "Project Name"). The project must have a communication link in order for the Analyzer to process events. Only the -start, -stop, and -status commands be performed on a single project. This option cannot be used in combination with the -host or -rmiregp options. When the -project option is used, the Analyzer host and port is looked up from the database.
-host HOST	Name of the host where the Analyzer runs. Can be either name or IP address. Local host is used if not specified.

Option	Description
-rmiregp PORTNUMBER	Port number on which the Analyzer listens for RMI connections. The default value is the value specified by the analyzer_port property in the Analyzer.properties file. This option only takes effect when communicating with an Analyzer that is already running; the Analyzer always uses the value specified in Analyzer.properties when starting up.
-debug	Start the Analyzer process in debug mode.

Examples:

- Start project PROJECT:
ServicesManager –start –proj PROJECT
- Stop Analyzer on host HOST:
ServicesManager –stop –host HOST

TimeServer**Location:**

TVISION_HOME/bin/TimeServer.[sh|bat]

Purpose:

Run the time server stand-alone.

Syntax:

```
TimeServer[.sh|.bat] [PORT]
```

Options:

Option	Description
PORT	Optionally, specify the listening port for the time server to use. The default is 21104.

Usage Notes:

In TIBCO and generic JMS environments, TransactionVision is unable to determine time skew across a communication link using a COA as it can with WebSphere MQ JMS. Instead, a time server reports the current time on the host where it is running. Any TransactionVision component can connect to it via TCP/IP and get the time skew from the current host to the host running the time server.

The Analyzer itself can act as a time server, or it can be run on a different host. To run the time server within the Analyzer, modify the `Analyzer.properties` file as follows:

- Set the `time_server` property to on
- Optionally, set `time_server_port` to the listening port you would like it to use. The default port is 21104.

Use the `TimeServer` utility to run the time server stand-alone.

The time server is configured per communication link found on the Time Server page in the Communication Link Wizard and under the Time Server Definition in the Edit Communication Link page. This definition tells the Analyzer and Sensors where the tie server is in order to calculate time skew.

ValidateXml

Location:

`TVISION_HOME/bin/ValidateXml.[sh|bat]`

Purpose:

Validates TransactionVision XML configuration files. Currently only the files `TransactionDefinition.xml`, `EventCorrelationDefinition.xml`, and `LocalTransactionDefinition.xml` can be validated. It is necessary to edit the XML file and remove the comments around the DTD declaration.

Syntax:

```
ValidateXML XML_filename
```

APIs Monitored by Sensors and Agents

CICS Sensor

The CICS Sensor collects events for five groups of CICS APIs: file control, temporary storage, transient data, interval control, and program control. In addition, the CICS Sensor collects Task Start, Task End and Program Start events as well.

- File Control APIs

APIs: WRITE, REWRITE, READ, STARTBR, RESETBR, READNEXT, READPREV, ENDBR, DELETE, UNLOCK

- Temporary Storage APIs

WRITEQ TS, READQ TS, DELETEQ TS

- Transient Data APIs

WRITEQ TD, READQ TD, DELETEQ TD

- Interval Control APIs

DELAY, START, ATTACH, CANCEL

- Program Control APIs

LINK, XCTL, RETURN

WebSphere MQ Sensor

The WebSphere MQ library, API exit and the z/OS WebSphere MQ Sensors collect events from the following MQ APIs.

- MQBACK

- MQBEGIN

- MQCLOSE

- MQCMIT
- MQCONN
- MQCONNX
- MQDISC
- MQGET
- MQINQ
- MQOPEN
- MQPUT
- MQPUT1
- MQSET

WBI (MQSI) Sensor

MQSI2TRACE. This is an event generated by the TransactionVision trace and failure nodes for WBI.

WebSphere MQ-IMS Bridge Sensor

- MQIMS_BRIDGE_ENTRY
- MQIMS_BRIDGE_EXIT

These are events generated by the WebSphere MQ-IMS Bridge Sensor when a MQ message is received by the MQ-IMS bridge and when a reply is generated by the WebSphere MQ-IMS bridge.

Servlet Sensor

The following servlet classes are instrumented:

- javax.servlet.http.HttpServlet
- org.apache.jasper.runtime.HttpJspBase
- com.ibm.servlet.PageListServlet

The servlet `Sensor` tracks the following methods:

```
void doGet(javax.servlet.http.HttpServletRequest,  
           javax.servlet.http.HttpServletResponse)
```

```
void doPost(javax.servlet.http.HttpServletRequest,  
            javax.servlet.http.HttpServletResponse)
```

```
void doPut(javax.servlet.http.HttpServletRequest,  
           javax.servlet.http.HttpServletResponse)
```

```
void service(javax.servlet.http.HttpServletRequest,  
             javax.servlet.http.HttpServletResponse)
```

```
void delete(javax.servlet.http.HttpServletRequest,  
            javax.servlet.http.HttpServletResponse)
```

```
void doHead(javax.servlet.http.HttpServletRequest,  
            javax.servlet.http.HttpServletResponse)
```

```
void doOptions(javax.servlet.http.HttpServletRequest,  
              javax.servlet.http.HttpServletResponse)
```

```
void doTrace(javax.servlet.http.HttpServletRequest,  
             javax.servlet.http.HttpServletResponse)
```

```
void init()
```

```
void init(javax.servlet.ServletConfig)
```

```
void destroy()
```

```
void _jspService(javax.servlet.http.HttpServletRequest,  
                 javax.servlet.http.HttpServletResponse)
```

```
void jspInit()
```

```
void jspDestroy()
```

Note: The TRACE and OPTIONS requests are different from the other requests because the servlet container has the following default implementation of these two requests:

- For TRACE requests, the servlet container returns the whole request string made by the client.
- For OPTIONS requests, the servlet container returns the list of HTTP requests handled by that particular servlet.

If you do NOT override the doTrace and doOptions methods, TransactionVision is unable to trace these requests, although the client still gets the response.

JMS Sensor

The JMS Sensor tracks the following methods:

Class:: ConnectionFactory
 createConnection
 createXAConnection

Class:: TopicConnectionFactory
 createTopicConnection
 createXATopicConnection

Class:: QueueConnectionFactory
 createQueueConnection

Class:: Connection
 createConnectionConsumer
 createDurableConnectionConsumer
 createSession
 createXASession
 close
 start
 stop
 setExceptionListener

Class:: QueueConnection
 createQueueSession
 createXAQueueSession

Class:: TopicConnection
 createTopicSession
 createXATopicSession

Class:: Session
 createProducer
 createConsumer

Class:: TopicSession
 createDurableSubscriber
 createPublisher
 createSubscriber
 unsubscribe
 commit
 recover
 rollback
 close

Class:: QueueSession
 createBrowser
 createReceiver
 createSender
 commit
 recover
 rollback
 createTemporaryQueue
 close

Class:: MessageProducer

Class:: MessageConsumer

Class:: QueueSender
 send
 close

Class:: TopicPublisher
publish
close

Class:: QueueBrowser
close
getEnumeration

Class:: QueueEnumeration
nextElement
hasMoreElements

Class:: TopicSubscriber
close
receive
setMessageListener

Class:: QueueReceiver
close
receive
setMessageListener

Class:: Message
acknowledge

EJB Sensor

The following EJB classes are instrumented:

- javax.ejb.EntityBean
- javax.ejb.SessionBean
- javax.ejb.MessageDrivenBean
- javax.ejb.EnterpriseBean (parent interface of all three above)

In addition to all business methods defined in public interfaces, the following methods are instrumented:

- ejbCreate()
- ejbPostCreate()
- ejbRemove() (session/entity/message driven bean)

- `ejbLoad()` (entity bean)
- `ejbStore()` (entity bean)
- `onMessage()` (message driven bean)

JDBC Sensor

The JDBC Sensor monitors the following interfaces and methods:

- Interface `java.sql.CallableStatement`
 - all execute methods
 - all set* methods
- Interface `java.sql.Connection`
 - `close`
 - `commit`
 - `releaseSavepoint`
 - `rollback`
 - `setSavePoint`
 - `prepareCall`
 - `prepareStatement`
- Interface `java.sql.Driver`
 - `connect`
- Interface `java.sql.PreparedStatement`
 - `execute`
 - `executeQuery`
 - `executeUpdate`
 - all set* methods

- Interface `java.sql.ResultSet`
 - `close`
 - `deleteRow`
 - `insertRow`
 - `refreshRow`
 - `updateRow`
 - all `get*` methods
 - cursor movement methods (`next`, `first`, etc.)
- Interface `java.sql.Statement`
 - `addBatch`
 - `cancel`
 - `clearBatch`
 - `close`
 - `execute`
 - `executeBatch`
 - `executeQuery`
 - `executeUpdate`

BEA Tuxedo Sensor

The BEA Tuxedo Sensor collects events from the following methods in the ATMI API:

- `tpcall`
- `tpenqueue`
- `tpdequeue`

NonStop TMF Sensor

The NonStop TMF Sensor monitors the following operations on audited Enscribe files:

- Insert
- Modify
- Delete

Technical References for Sensor and Agent Technologies

Following is a list of technical references for Sensor and Agent technologies:

- IBM z/OS or z/OS MVS Initialization and Tuning Reference
- MVS Initialization and Tuning Reference
- IBM DB2 Administration Guide
- WebSphere MQ Application Programming Guide (SC33-0807-08)
- Managing WebLogic Security

Troubleshooting

This section describes the following troubleshooting topics:

- “Enabling SHLIB_PATH on HP-UX” on page 422
- “Displaying Perl and Java Program Names” on page 423
- “Monitoring WebSphere MQ Triggered Programs” on page 423
- “Monitoring Client Applications” on page 424
- “Optimizing Performance” on page 424
- “Running TransactionVision Sensors and Analyzers on VMWare systems” on page 425
- “Missing Events” on page 425
- “Queue Manager or Event Broker Channel Limitations” on page 426

- “Servlet and EJB Sensor Problems” on page 426
- “JMS Sensor Problems” on page 427
- “UI/Job Server Troubleshooting” on page 427
- “Application Troubleshooting” on page 427
- “Event Queue Cleanup” on page 428
- “Problems Importing System Tables into a Database” on page 431
- “Connection to Database Is Lost” on page 434
- “Problems Using International Characters” on page 435

Enabling SHLIB_PATH on HP-UX

When you run a program on HP-UX, the loader may use the library paths embedded in the program instead of the library paths specified in the **SHLIB_PATH** environment variable. You may use the **chatr** command to check the current settings on a program:

```
chatr executable_name
```

The output of this command may be as follows:

```
shared executable
shared library dynamic path search:
  SHLIB_PATH    disabled second
  embedded path enabled  first /opt/mqm/lib
```

To enable run-time location of the WebSphere MQ Library Sensor, you may use the following command to turn on the loader’s use of **SHLIB_PATH** as the first paths to search for libraries:

```
chatr +s enable +b enable executable_name
```

Displaying Perl and Java Program Names

By default, the TransactionVision displays Perl or java.exe as the program name for Perl or Java applications. To display the program name, set the **TVISION_REPORT_ARGS** environment variable to any value. To disable this behavior, unset **TVISION_REPORT_ARGS**. This environment variable can be set on any platform except z/OS. On Windows, it must be set to a value other than an empty string.

Monitoring WebSphere MQ Triggered Programs

WebSphere MQ triggered programs on UNIX platforms can be monitored if two conditions are met:

- 1** The runmqtrm trigger monitor is not setuid/gid, or you are running runmqtrm as the mqm user. (This is because the programs started by the setuid/gid runmqtrm ignore any environment variables modifying the path if the user is not the setuid/gid user - the mqm user in this case.)
- 2** You actually trigger a script instead of your program. The script sets the appropriate environment variable (based on the OS platform) to point to the Sensor and then runs the actual program. For example, on Solaris this script might be:

```
#!/bin/sh
LD_LIBRARY_PATH=/opt/HP/TransactionVision/lib
export LD_LIBRARY_PATH
/usr/local/bin/trigproc "$@"
```

Where /usr/local/bin/trigproc was the program being triggered. The following table shows the appropriate environment variable and default installation directory for each platform:

Platform	Environment Variable	Default Directory
Windows	PATH	C:\Program Files\Hewlett-Packard\TransactionVision Sensor\lib
Sun Solaris	LD_LIBRARY_PATH	/opt/HP/TransactionVision/lib
HP-UX	SHLIB_PATH	/opt/HP/TransactionVision/lib

Platform	Environment Variable	Default Directory
Linux	LD_LIBRARY_PATH	/opt/HP/TransactionVision/lib
IBM AIX	LIBPATH	/usr/lpp/HP/TransactionVision/lib

Note: The **runmqtrm** trigger monitor program cannot itself be sensed by the Sensor, because it uses internal WebSphere MQ functions which are not compatible with the Sensor. To sense the trigger monitor itself, use either the **amqstrg0.c** trigger monitor or build your own using only public WebSphere MQ APIs.

Monitoring Client Applications

If you are monitoring client applications and notice that the WebSphere MQ listener does not generate TransactionVision events when invoked from **inetd**, check the user id and group id of the **mqm** user as follows:

► id mqm

uid=nnn(mqm) gid=nnn(mqm)

Both the user and the group should be **mqm**. If not, the user **mqm** should be added to the group **mqm**. Alternately, you may change the group id on the listener binary to belong to the group that the user **mqm** belongs to.

Optimizing Performance

To achieve optimum performance when using TransactionVision, follow these guidelines:

- Install the Analyzer on a computer that is fast enough to keep up with the rate of incoming messages from Sensors. For a slower Analyzer host, increase the maximum queue depth on the event queue used for sending event messages from Sensors to the Analyzer.

- Use data collection filters to restrict event collection to event data absolutely necessary to resolve a particular problem. Be sure to delete the default **Collect All** data collection filter if you define a more restrictive filter, since events that meet any filter are collected by Sensors.
- Use the **collect_dbstats.sql** script that can be generated by the CreateSqlScript utility to optimize database performance. See the *HP TransactionVision Deployment Guide* PDF for complete details.

Running TransactionVision Sensors and Analyzers on VMWare systems

Running TransactionVision Sensors and Analyzers on VMWare systems can result in inaccurate event times. Because virtual machines work by time-sharing physical hardware, a virtual machine cannot exactly duplicate the timing behavior of a physical machine. This can result in a drift in the time reported by the guest operating system to the TransactionVision Sensors. The effect of this drift is that for transactions spanning multiple machines, relative transaction times are not accurate. For a detailed description of this issue, see this document on the VMWare website: https://www.vmware.com/pdf/vmware_timekeeping.pdf.

Missing Events

If it appears that events are missing from a project, first make sure that no data collection filter or project query is preventing the events from being collected or displayed. In the data collection filter, if the event retry timeout is set to a value other than Retry Forever and the timeout is being met, events are lost.

If the events are still missing and they are from WebSphere MQ, make sure that the WebSphere MQ API call uses a valid hConn parameter. The Sensor uses the **hConn** value to determine which queue manager or event broker to connect to for configuration messages; if this value is invalid, the Sensor cannot check for configuration messages.

Queue Manager or Event Broker Channel Limitations

MaxChannels and **MaxActiveChannels** queue manager or event broker properties have to be set to a reasonable amount for TransactionVision. By default, these properties are not set and there is no limitation. For each active communication link, the TransactionVision Analyzer requires two channels plus one channel for each listening thread defined in the communication link.

If these properties are set too low, the user sees strange behavior such as a partial connection to the communication link. It may at first appear that the Analyzer is connected, but it reports a connection error and puts the communication link in a “Trying to Connect” state. This behavior usually points to a problem with these queue manager or event broker properties not being set properly.

Servlet and EJB Sensor Problems

If you are not able to collect servlet or EJB events from your web application, check the following:

- 1** You have run **SensorSetup.sh**(bat) to install TransactionVision Servlet Sensor on the application server your web application is running on. The Application server name is case sensitive. “Default Server” is different from “default server” or “DefaultServer” (space counts). Run **SensorSetup.sh**(bat) again if needed.
- 2** Check **config/sensor/SensorConfiguration.xml** from your TransactionVision home to make sure the information you provided there is correct and the queue manager or event broker is running. Run **SensorSetup.sh**(bat) to modify the setting if needed.
- 3** Check **sensor.log** in your log file directory to see if there are any errors.
- 4** Check WebSphere’s **system.out** and **system.err** log files to see if there are any errors.

JMS Sensor Problems

If you are not able to collect JMS events from stand-alone JMS application, check the following:

- 1** Make sure that file `java\lib\com.ibm.mqjms.jar` is in your TransactionVision home. Run **SensorSetup.sh**(bat) again if it is not there.
- 2** Your **ClassPath** points to `java\lib\com.ibm.mqjms.jar` in TransactionVision home AHEAD of the same file from WebSphere MQ directory.
- 3** Check `config/sensor/SensorConfiguration.xml` from your TransactionVision home to make sure the information you provided there is correct and the queue manager or event broker is running. Run **SensorSetup.sh**(bat) to modify the setting if needed.
- 4** If JMS functions are called from servlets, make sure you answer “YES” to the question “Do you wish to monitor JMS methods along with Servlet methods?” when running **SensorSetup.sh**(bat).

UI/Job Server Troubleshooting

Manually Deleting a Schema

If you manually delete a project’s schema (using **CreateSqlScript**), you may still login to TransactionVision, but most pages do not work because the database table does not exist. To remedy this situation, select the database schema associated with the project on the Database Schemas page and click **Create Tables**.

Application Troubleshooting

TransactionVision does not require changes to monitored applications. However, applications on the AIX platform may fail after installing a WebSphere MQ support pac. In WebSphere MQ support pacs, internal symbols exported from the TransactionVision WebSphere MQ Sensor on the AIX platform may change. When an internal symbol that has been exported from the Sensor library is no longer available in the WebSphere MQ library, the application cannot start and fails with various symbol resolution errors.

To work around this problem, run the `rebind_sensor` script whenever a WebSphere MQ support pac that modifies the WebSphere MQ libraries (**libmqm.a**, **libmqic.a**, **libmqm_r.a**, **libmqic_r.a**) are modified. For more information about this script, see the *HP TransactionVision Deployment Guide* PDF.

Event Queue Cleanup

There may be times when messages that are not being collected by a TransactionVision project accumulate on the event queue. These events may be targeted for invalid schemas or communication links, or they may be erroneous messages. In either case, you may use the `ManageQueue` utility to remove or replay these messages.

To view a snapshot of the messages on an event queue (TVISION.EVENT.QUEUE on queue manager QMGR1 in this example), enter the following command:

```
ManageQueue.sh -qm qmgr1 -input_queue TVISION.EVENT.QUEUE  
-show_event_ids
```

For complete information about **ManageQueue**, see “ManageQueue” on page 401.

Invalid Schema and Communication Link IDs

Events are targeted for a specific database, schema, and communication link identified within TransactionVision by a unique ID. If a schema is deleted or a communication link is changed, the unique ID no longer exists, and any events still on the event queue that targeted for that schema or communication link are no longer valid.

To display all schema and commlink IDs, enter the following command:

```
EventQueueCleanup -showids
```

To delete events targeted for database schemas or communication links which no longer exist from the queue TVISION.EVENT.QUEUE on the queue manager QMGR1, enter the following command:

```
ManageQueue.sh -show_db_ids
```

You can also use ManageQueue to replay events targeted for invalid schema or communication link IDs to a new schema and communication link so they can be processed by the Analyzer:

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -output_queue TVISION.EVENT.QUEUE -move 2 11 -verbose
```

In this example, the schema ID for the new schema is 2 and the communication link ID for the new communication link is 11.

Erroneous Messages

Erroneous messages are messages on the event queue that do not appear to be TransactionVision event messages. Erroneous messages could include messages generated from other applications or WebSphere MQ COA report messages used by communication link time skew tests.

To remove erroneous messages from TVISION.EVENT.QUEUE on queue manager TransactionVision.qmgr, enter the following command:

```
ManageQueue.sh -qm tv.qmgr -input_queue TVISION.EVENT.QUEUE -delete -erroneous -orphaned -verbose
```

Events Not Pulled from Event Queue

Event queues can be shared across Analyzers, projects, and communication links. Therefore, events have to be retrieved from the queue based on a specific message ID (correlation ID in the case of Java Sensor events).

The following diagram shows the layout of the message ID (correlation ID in the case of Java Sensor events) on TransactionVision events.

12 Bytes Database Name	4 Bytes Database Host IP	4 Bytes CommLink ID	4 Bytes Schema ID
---------------------------	-----------------------------	------------------------	----------------------

Given an assigned communication link, an Analyzer gets events only by a specific message ID/correlation ID. Therefore, it is possible to have events on the event queue which may never be removed from the queue. Different sceneries can cause this.

Dropped Database

If a database is dropped and never created again, events intended for projects that were previously on the dropped database may exist on the event queue.

Deleted Schema

If a project schema is deleted, events intended for the schema may exist on the event queue. Recreating the schema does not reassign these events on the event queue to the newly created schema because a new unique ID is always assigned to a created schema, even if it is given the same name as the previous.

Deleted or Changed Commlink

If a communication link is deleted from a project or deassigned from the Analyzer in a project, events intended for that communication link may exist on the event queue. If the communication link has just been deassigned, reassigning it to the Analyzer allows the Analyzer to retrieve the pending events on the event queue. However, if the communication link has been deleted from the project, recreating the communication link (even if created from a communication link template) does not provide the means to retrieve the events left on the event queue. This is because every time a communication link is created, it is assigned a unique ID. In addition, if a communication link is changed, a new unique id is assigned to it, and old events remain on the event queue.

WebSphere MQ Anomalies

TransactionVision events are not the only kinds of messages that may appear on an event queue.

For client connections to a host other than the one running the Sensor, TransactionVision uses WebSphere MQ Confirmation of Arrival (COA) messages to determine time skew across communication links. This means that a COA request message is sent to the configuration queue, and the COA report is sent back on the event queue.

If a remote (using multiple queue managers or event brokers) communication link is being used, the COA report has to travel across a WebSphere MQ channel to reach the final event queue destination. If there are problems in the WebSphere MQ environment (for example, the channel is down), the Analyzer may timeout waiting for a COA report message. If the WebSphere MQ problem is later fixed, the COA report message may finally arrive in the event queue. This COA report message is never removed from the event queue, because the Analyzer has given up on it.

There is also the possibility of other WebSphere MQ applications maliciously or accidentally putting messages on the TransactionVision event queue. The Analyzer ignores these because they are most likely not to have a matching message ID/correlation ID it is looking for, and therefore remain on the event queue indefinitely.

Problems Importing System Tables into a Database

Oracle

The referential constraints of TransactionVision system tables prevent Oracle data import utility from loading data into existing tables. The Oracle data import utility returns an error similar to the following example:

```
IMP-00019: row rejected due to ORACLE error 2291
IMP-00003: ORACLE error 2291 encountered
ORA-02291: integrity constraint (TVISION.FK1_FILTER) violated - parent key not found
Column 1 17
Column 2 Collect All
Column 3 16
Column 4 00540001020C0000000100000001000000179F0700008ADC00... 0 rows
imported
```

If you encounter this error, disable the constraints, import the data, and then re-enable the constraints.

For more information, see the following Oracle document:

http://download-west.oracle.com/docs/cd/A91202_01/901_doc/server.901/a90192/ch02.htm#1005094

Disabling Referential Constraints

To disable the referential constraints, use the following SQL commands:

```
alter table TVISION.ANALYZER_PROJ_COMMLINK DISABLE CONSTRAINT
FK2_ANLZ_PROJ_LINK;
alter table TVISION.ANALYZER_SCHEMA DISABLE CONSTRAINT
FK2_ANLZ_SCHEMA;
alter table TVISION.COMMLINK_FILTER DISABLE CONSTRAINT FK1_LINK_FILTER;
alter table TVISION.COMMLINK_FILTER DISABLE CONSTRAINT FK2_LINK_FILTER;
alter table TVISION.FILTER DISABLE CONSTRAINT FK1_FILTER;
alter table TVISION.PROJECT DISABLE CONSTRAINT FK1_PROJECT;
alter table TVISION.PROJECT_COMMLINK DISABLE CONSTRAINT
FK1_PROJ_LINK;
alter table TVISION.QUERY DISABLE CONSTRAINT FK1_QUERY;
```

Re-enabling Referential Constraints

To re-enable the referential constraints after importing the data, use the following SQL commands:

```
alter table TVISION.ANALYZER_PROJ_COMMLINK ENABLE CONSTRAINT
FK2_ANLZ_PROJ_LINK;
alter table TVISION.ANALYZER_SCHEMA ENABLE CONSTRAINT
FK2_ANLZ_SCHEMA;
alter table TVISION.COMMLINK_FILTER ENABLE CONSTRAINT FK1_LINK_FILTER;
alter table TVISION.COMMLINK_FILTER ENABLE CONSTRAINT FK2_LINK_FILTER;
alter table TVISION.FILTER ENABLE CONSTRAINT FK1_FILTER;
alter table TVISION.PROJECT ENABLE CONSTRAINT FK1_PROJECT;
alter table TVISION.PROJECT_COMMLINK ENABLE CONSTRAINT
FK1_PROJ_LINK;
alter table TVISION.QUERY ENABLE CONSTRAINT FK1_QUERY;
```

DB2

When importing system tables into a DB2 database, some tables may be placed in a Check Pending state. To resolve this problem, disable the constraints, import the data, and then re-enable the constraints.

Disabling Referential Constraints

To disable the referential constraints, use the following SQL commands:

```
alter table tvision.analyzer_proj_commlink alter FOREIGN KEY fk2_anlz_proj_link NOT
ENFORCED;
alter table TVISION.ANALYZER_SCHEMA alter FOREIGN KEY FK2_ANLZ_SCHEMA
NOT ENFORCED;
alter table TVISION.ANALYZER_SCHEMA alter FOREIGN KEY FK1_ANLZ_SCHEMA
NOT ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK1_LINK_FILTER
NOT ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK2_LINK_FILTER
NOT ENFORCED;
alter table TVISION.FILTER alter FOREIGN KEY FK1_FILTER NOT ENFORCED;
alter table TVISION.PROJECT alter FOREIGN KEY FK1_PROJECT NOT
ENFORCED;
alter table TVISION.PROJECT_COMMLINK alter FOREIGN KEY FK1_PROJ_LINK
NOT ENFORCED;
alter table TVISION.QUERY alter FOREIGN KEY FK1_QUERY NOT ENFORCED;
alter table TVISION.STORAGE alter FOREIGN KEY FK1_STORAGE NOT
ENFORCED;
alter table TVISION.REPORT_PARAMETERS alter FOREIGN KEY FK1_REP_PARMS
NOT ENFORCED;
alter table TVISION.JOB_STATUS alter FOREIGN KEY FK1_JOB_ID NOT
ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK1_LINK_FILTER
NOT ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK2_LINK_FILTER
NOT ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK3_LINK_FILTER
NOT ENFORCED;
alter table TVISION.ANALYZER_PROJ_COMMLINK alter FOREIGN KEY
FK1_ANLZ_PROJ_LINK NOT ENFORCED;
alter table TVISION.ANALYZER_PROJ_COMMLINK alter FOREIGN KEY
FK2_ANLZ_PROJ_LINK NOT ENFORCED;
```

Re-enabling Referential Constraints

To re-enable the referential constraints after importing the data, use the following SQL commands:

```
alter table TVISION.ANALYZER_PROJ_COMMLINK alter FOREIGN KEY
fk2_anlz_proj_link ENFORCED;
alter table TVISION.ANALYZER_SCHEMA alter FOREIGN KEY FK1_ANLZ_SCHEMA
ENFORCED;
alter table TVISION.ANALYZER_SCHEMA alter FOREIGN KEY FK2_ANLZ_SCHEMA
ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK1_LINK_FILTER
ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK2_LINK_FILTER
ENFORCED;
alter table TVISION.FILTER alter FOREIGN KEY FK1_FILTER ENFORCED;
alter table TVISION.PROJECT alter FOREIGN KEY FK1_PROJECT ENFORCED;
alter table TVISION.PROJECT_COMMLINK alter FOREIGN KEY FK1_PROJ_LINK
ENFORCED;
alter table TVISION.QUERY alter FOREIGN KEY FK1_QUERY ENFORCED;
alter table TVISION.STORAGE alter FOREIGN KEY FK1_STORAGE ENFORCED;
alter table TVISION.REPORT_PARAMETERS alter FOREIGN KEY FK1_REP_PARMS
ENFORCED;
alter table TVISION.JOB_STATUS alter FOREIGN KEY FK1_JOB_ID ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK1_LINK_FILTER
ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK2_LINK_FILTER
ENFORCED;
alter table TVISION.COMMLINK_FILTER alter FOREIGN KEY FK3_LINK_FILTER
ENFORCED;
alter table TVISION.ANALYZER_PROJ_COMMLINK alter FOREIGN KEY
FK1_ANLZ_PROJ_LINK ENFORCED;
alter table TVISION.ANALYZER_PROJ_COMMLINK alter FOREIGN KEY
FK2_ANLZ_PROJ_LINK ENFORCED;
```

Connection to Database Is Lost

If the connection to the TransactionVision database is lost, the following message is displayed:

"TransactionVision Error: The connection to the TransactionVision database has been lost. Please ensure that the database is accessible, and then re-login to BAC again."

In some cases, making sure the database is accessible and re-logging into Business Availability Center does not resolve the problem. In these cases, the problem can be resolved by first closing the browser window and then logging into Business Availability Center again.

Problems Using International Characters

See the *HP TransactionVision Deployment Guide* PDF for instructions on how to configure TransactionVision and your DBMS to store and display unicode data.

Garbage Characters in Database Column Values

The most likely cause of this problem is inserting Unicode characters into a database that has not been created as a UTF-8 database. If your user data contains any Unicode that gets mapped to database tables, you must create the database with the UTF-8 code set.

Value Too Large For Column Exceptions

Database exceptions during event processing reporting that the inserted value is too large for the column indicates that, for custom XDM columns, a corresponding table column might not have been generated large enough to accommodate the number of characters specified in the XDM definition. To resolve this problem, set the attribute `unicode=true` on the column.

Unable to Make a Client Connection to a Queue Manager or Event Broker Using an International Character Set

The problem indicates that the queue manager or event broker is using a coded character set identifier (CCSID) of mixed mode character set (for example, Shift-JIS or another double-byte character set) and the CCSID has not been set correctly.

Sensor Configuration Queue Manager or Event Broker

The CCSID of the queue manager or event broker should be set during SensorSetup. If you are unable to make a client connection to the Sensor configuration queue manager or event broker, you may see the following errors in the sensor.log file:

```
2006-03-10 10:17:23,218 [Thread-1] FATAL sensorLog - TransactionVision
Error(FailedToCreateTransportSessionForListener): Failed to create a transport session
for Sensor configuration listener.
```

```
2006-03-10 10:17:23,218 [Thread-1] FATAL sensorLog - TransactionVision
Error(TransportConnectFailed): Failed to connect ('jpsvr.tv1.manager' 'queue:///
TVISION.CONFIGURATION.QUEUE?priority=-1&persistence=-1&targetClient=1').
(2195) Unexpected error occurred.
```

```
2006-03-10 10:17:23,218 [Thread-1] FATAL sensorLog - MQJMS2005: failed to create
MQQueueManager for 'bristol-jpsvr:jpsvr.tv1.manager'
```

```
2006-03-10 10:17:23,218 [Thread-1] FATAL sensorLog -
javax.jms.JMSEException: MQJMS2005: failed to create MQQueueManager for
'bristol-jpsvr:jpsvr.tv1.manager'
```

```
at com.ibm.mq.jms.services.ConfigEnvironment.
    newException(ConfigEnvironment.java:546)
at com.ibm.mq.jms.MQConnection.createQM(MQConnection.
    java:1450)
at com.ibm.mq.jms.MQConnection.createQMNonXA(MQConnection.
    java:960)
at com.ibm.mq.jms.MQQueueConnection.<init>(MQQueueConnection.
    java:159)
at com.ibm.mq.jms.MQQueueConnection.<init>(MQQueueConnection.
    java:77)
at com.ibm.mq.jms.MQQueueConnectionFactory.createQueue
    Connection(MQQueueConnectionFactory.java:142)
at com.bristol.tvision.transport.MQJMSTransportSession.
    connect(MQJMSTransportSession.java:156)
at com.bristol.tvision.sensor.configurationmgr.Configuration
    ListenerThread.run(ConfigurationListenerThread.java:196)
```

To resolve this problem:

- 1 Run the following command to identify the CCSID of the target queue manager or event broker:

```
echo display qmgr ccsid | runmqsc <sensor_config_qmgr>
```

2 Re-run **SensorSetup** and enter the correct CCSID.

Analyzer Configuration Queue Manager or Event Broker

The CCSID of the queue manager or event broker should be set in the communication link. If you are unable to make a client connection to the Analyzer configuration queue manager or event broker, you may see the following errors in the analyzer.log file:

```
2006-03-10 10:44:44,718 [ConfigurationThread (id=10,action=timeskew)] FATAL
AppLog - TransactionVision Fatal(MQConnectionFailed): Failed to connect
('jpsvr.tv1.manager'): (2195) Unexpected error occurred.
2006-03-10 10:44:44,718 [ConfigurationThread (id=10,action=timeskew)] FATAL
AppLog -
com.ibm.mq.MQException: MQJE001: Completion Code 2, Reason 2195
  at com.ibm.mq.MQManagedConnectionJ11.<init>
    (MQManagedConnectionJ11.java:172)
  at com.ibm.mq.MQClientManagedConnectionFactoryJ11.
    _createManagedConnection(MQClientManagedConnection
    FactoryJ11.java:270)
  at com.ibm.mq.MQClientManagedConnectionFactoryJ11.
    createManagedConnection(MQClientManagedConnection
    FactoryJ11.java:290)
  at com.ibm.mq.StoredManagedConnection.<init>
    (StoredManagedConnection.java:80)
  at com.ibm.mq.MQSimpleConnectionManager.allocate
    Connection (MQSimpleConnectionManager.java:150)
  at com.ibm.mq.MQQueueManager.obtainBaseMQQueueManager
    (MQQueueManager.java:682)
  at com.ibm.mq.MQQueueManager.construct(MQQueueManager.
    java:620)
  at com.ibm.mq.MQQueueManager.<init>(MQQueueManager.
    java:393)
  at com.bristol.tvision.services.collection.
    MQConfigurationThread.createConnection(MQConfiguration
    Thread.java:356)
  at com.bristol.tvision.services.collection.
    MQConfigurationThread.getTimeSkew(MQConfiguration
    Thread.java:218)
  at com.bristol.tvision.services.collection.Configuration
    Thread.run(ConfigurationThread.java:141)
```

To resolve this problem:

- 1** Run the following command to identify the CCSID of the target queue manager or event broker:

 `echo display qmgr ccsid | runmqsc <analyzer_config_qmgr>`
- 2** Edit the communication link to set the correct CCSID.

Part III

Reports and Topologies

14

Introducing TransactionVision Reports and Topologies

This chapter introduces TransactionVision and the TransactionVision Reports and Topologies pages.

This chapter includes:

Concepts

- TransactionVision Reports and Topologies Overview on page 441
- User-created Reports on page 442

TransactionVision Reports and Topologies Overview

The TransactionVision Reports and Topologies pages enable users to display and set up transaction reports and topologies.

To access the TransactionVision Reports and Topologies pages, select **Applications > TransactionVision**.

TransactionVision Reports and Topologies include the following pages:

- Transaction Topology, see the following chapters for details:
 - “Aggregated Topology” on page 445
 - “Instance Topology” on page 453
 - “Transaction Topology Flow Map” on page 465
 - “Component Topology Analysis” on page 491

- Transaction Infrastructure, see “Transaction Infrastructure Reports” on page 569 for details.
- Transaction Detail, see “Transaction Detail Reports” on page 533 for details.
- Transaction Summary, see “Transaction Summary Reports” on page 587 for details.
- User-created Reports, see “User-created Reports” on page 442.

User-created Reports

It is possible to create your own TransactionVision reports, and have these reports appear on the User-created Reports page.

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. Use the standard TransactionVision reports as a reference when creating your own reports.

To create 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 framework, 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**.

See the *HP TransactionVision Advanced Customization Guide* PDF for further details.

15

Aggregated Topology

This chapter includes the main concepts and reference information for the Aggregated Topology.

This chapter includes:

Concepts

- Aggregated Topology Overview on page 445
- Response Time Calculation and Breakdown on page 447

Reference

- Aggregated Topology User Interface on page 448

Aggregated Topology Overview





The Aggregated Topology displays a topology flow map for each Business Transaction, along with aggregated response time and volume metrics. The data shown in this view is an aggregation of all the instances that occurred in the selected time range.

The Aggregated Topology table and flow map show the transaction execution from front-end real user experience through the back-end server level components (such as application servers, message middleware servers, database servers, mainframe transaction and batch jobs).

The data appears in both the Aggregated Transaction Topology table and flow map, based on one or more transaction instances of a Business Transaction, satisfying a given filter. Filtering conditions can be based on Business Transactions and date and time range. See Chapter 17, “Transaction Topology Flow Map,” for details about the Aggregated Transaction Topology flow map. For more information about setting the date and time, and filtering, see “Transaction Topology Flow Map for the Aggregated Topology” on page 476.

Drilldown to the Instance Topology Page From the Aggregated Topology Page

You can drill down from the Aggregated Topology to view specific instances of a Business Transaction on the Instance Topology page, through one of the following four icons at the top of the Aggregated Topology table (see “Aggregated Topology Page” on page 449 for details):

-  drill to the slowest instances
-  drill to failed instances
-  drill to delayed instances
-  drill to the greatest transaction value instances

Data Comparison Between the Aggregated and Instance Topology Pages

When you compare the Aggregated Topology page to the Instance Topology page for the same time period, there may be some differences between the data displayed on the two pages in the following cases:

- The default lag interval for the aggregated data on Business Transactions to be sent from TransactionVision to Business Availability Center is 6 minutes. Therefore, the Instance Topology page might show more recent data in comparison to the Aggregated Topology page. For example, when viewing the past 15 minutes of data, the Instance Topology page shows data for this entire period, while the Aggregated Topology page does not contain some of the data because of the lag associated with the aggregated data.

- The in-process transactions count is not displayed in the table on the Aggregated Topology page, while the Instance Topology page shows all instances including the ones that are currently in-process.

Response Time Calculation and Breakdown

Response time is calculated by the Event Analyzer and is based on a correlated set of events from the different instrumentation points (such as Real User Monitor and TransactionVision Sensors).

End-to-end response time when viewed as a breakdown of different response time measurements, is represented differently for synchronous and asynchronous Business Transactions.

The actual calculation of response time using the events is conceptually the same calculation for synchronous or asynchronous Business Transactions as follows:

Subtract the start time of the first event in the Business Transaction from the last event in the Business Transaction. This provides the end-to-end response time of all variations of synchronous and asynchronous Business Transactions.

Synchronous Business Transaction

In the case of a purely synchronous Business Transaction, the end-to-end response time is equal to the response time measured by the closest point to the end user. This is reported by either the Real User Monitor or by a TransactionVision server Sensor if the Real User Monitor is not present.

- If the page is being monitored by the Real User Monitor, the end-to-end response time is equivalent to the end user response time.
- If the Real User Monitor is not measuring the page, then the transaction begins at the server and the end-to-end response time does not include the end user time.

Asynchronous Business Transaction

In the case of an asynchronous Business Transaction, the end-to-end response time is the sum of the synchronous portion of the transaction (as described in “Synchronous Business Transaction” on page 447) and the asynchronous portion.

Aggregated Topology User Interface

This section describes:







- Aggregated Topology Page on page 449
- Active Filters Dialog Box for Aggregated Topology on page 452

Aggregated Topology Page

Description	<p>Enables you to view the Aggregated Transaction Topology flow map and aggregated response time and volume metrics for a Business Transaction. You can also filter which Business Transactions you want to see.</p> <p>To access: Select Applications > TransactionVision > Transaction Topology > Aggregated Topology</p>
Important Information	<p>The Aggregated Topology page is divided into two parts:</p> <ul style="list-style-type: none"> ➤ Aggregated Topology table — displays aggregated performance and volume metrics for transactions that ended in the selected time range. You can select a row in the table to see the Transaction Topology flow map for the selected Business Transaction. ➤ You can manipulate the table by adding and removing columns and sorting the table by column values. ➤ For information on how to modify the table, see “Common Report Elements” in <i>Reports</i>. ➤ For descriptions of the type of data that is available, see “Aggregated Topology Table” on page 451. ➤ Aggregated Transaction Topology flow map — displays a diagram representing the flow of a given Business Transaction. See “Transaction Topology Flow Map for the Aggregated Topology” on page 476 for details.
Useful Links	<ul style="list-style-type: none"> ➤ “TransactionVision Data” on page 20 ➤ Chapter 16, “Instance Topology” ➤ Chapter 17, “Transaction Topology Flow Map”

Aggregated Topology Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
<Common report elements>	See “Common Report Elements” in <i>Reports</i> . These settings describe how to set the date and time range, configure a favorite filter, publish a report of the Aggregated Topology page and annotate the Aggregated Topology flow map (see “Transaction Topology Flow Map for the Aggregated Topology” on page 476).
	Click to drill to the slowest instances on the Instance Topology page. See “Instance Topology Page” on page 456 for details.
	Click to drill to failed instances on the Instance Topology page. See “Instance Topology Page” on page 456 for details.
	Click to drill to delayed instances on the Instance Topology page. Delayed instances are transactions that violated their SLA. See “Instance Topology Page” on page 456 for details.
	Click to drill to the greatest transaction value instances on the Instance Topology page. See “Instance Topology Page” on page 456 for details.
	Click to drill to the Transaction Performance Summary Report. See “Transaction Performance Summary Report” on page 603 for details.
	Click to drill to the Transaction Over Time Report. See “Transaction Over Time Report” on page 598 for details.
Active Filters	Click to open the “Active Filters Dialog Box for Aggregated Topology” on page 452.

Aggregated Topology Table

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Transaction Name	Name of the Business Transaction.
Avg End to End Response Time (sec.)	Average of end-to-end transaction response time in seconds. See “Response Time Calculation and Breakdown” on page 447 for details on how this metric is calculated.
Avg Response Time of Back End (sec.)	Response time of back-end server level components, such as application servers, message middleware servers, database servers, mainframe transactions and batch jobs.
Avg End User Response Time (sec.)	Response time of end user as measured by Real User Monitor (RUM). If RUM is not configured, this value is blank.
Min Response Time (sec.)	Minimum transaction response time in seconds.
Max Response Time (sec.)	Maximum transaction response time in seconds.
Transaction Count	Count of total number of transaction instances in the selected time frame.
Failed Count	Count of transaction instances with Failed status.
Late Count	Count of transaction instances that violated their SLA.
Exception Count	Count of transaction instances that threw an exception.

 **Active Filters Dialog Box for Aggregated Topology**

Description	Enables you to filter the list of Business Transactions displayed in the page. To access: On the Applications > TransactionVision > Transaction Topology > Aggregated Topology > Aggregated Topology page, click Active Filters .
Important Information	You can filter the Aggregated Topology to change what you see in the table and flow map, by Business Transactions.
Useful Links	<ul style="list-style-type: none">➤ “Aggregated Topology Page” on page 449➤ “Transaction Topology Flow Map for the Aggregated Topology” on page 476

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Common report elements>	See “Common Report Elements” in <i>Reports</i> .
Browse	Mode in which you can filter the Business Transactions.
Cancel	Click to return to the Aggregate Topology page without making any changes.
OK	Click to save your settings.

16

Instance Topology

This chapter includes:

Concepts

- Instance Topology Overview on page 453
- Calculating Response Time on page 455

Reference

- Instance Topology User Interface on page 455

Troubleshooting and Limitations on page 463

Instance Topology Overview

The Instance Topology page allows you to see the performance details of Business Transaction instances that are executed in a specific time range. You can select a specific instance to see the Business Transaction's topology in the Instance Topology flow map.





The Instance Topology flow map shows the flow of a specific transaction instance overlaid on the Aggregated Topology for the Business Transaction. Like the Aggregated Topology flow, the Instance Topology flow is also end-to-end, front-end real user experience through the back-end server level components (such as application servers, message middleware servers, database servers, mainframe transaction and batch jobs).

The data appears in both the Instance Transaction Topology table and flow map, based on instances of a specific Business Transaction, which satisfy a given filter. Filtering conditions can be based on instances of the Business Transactions, date and time range, metrics, and end user groups.

For information about the Instance Transaction Topology flow map, see “Transaction Topology Flow Map for the Instance Topology” on page 486.

Accessing the Instance Topology Page From the Aggregated Topology Page

You can drill down from the Aggregated Topology to view specific instances of a Business Transaction on the Instance Topology page, through one of the following four icons at the top of the Aggregated Topology table (see “Aggregated Topology” on page 473 for details):

-  drill to the slowest instances
-  drill to failed instances
-  drill to delayed instances
-  drill to the greatest transaction value instances

Comparing Data in the Aggregated and Instance Topology Pages

When you compare the Aggregated Topology page to the Instance Topology page for the same time period, there may be some differences between the data displayed on the two pages in the following cases:

- The default lag interval for the aggregated data on Business Transactions to be sent from TransactionVision to Business Availability Center is 6 minutes. Therefore, the Instance Topology page might show more recent data in comparison to the Aggregated Topology page. For example, when viewing the past 15 minutes of data, the Instance Topology page shows data for this entire period, while the Aggregated Topology page does not contain some of the data because of the lag associated with the aggregated data.
- The in-process transactions count is not displayed in the table on the Aggregated Topology page, while the Instance Topology page shows all instances including the ones that are currently in-process.

Calculating Response Time

Response time is calculated by the Event Analyzer and is based on a correlated set of events from the different instrumentation points (such as Real User Monitor and TransactionVision Sensors). For details on calculating the response time, see “Response Time Calculation and Breakdown” on page 447.

Instance Topology User Interface

This section describes:



- Instance Topology Page on page 456
- Active Filters Dialog Box for Instance Topology on page 458

Instance Topology Page

Description	<p>Enables you to view data collected for instances of a particular Business Transaction in table and flow map format. You can filter on Business Transaction instances, date and time range, metrics, or user groups.</p> <p>To access:</p> <ul style="list-style-type: none"> ➤ Select Applications > TransactionVision > Transaction Topology > Instance Topology. Select a transaction instance and click Generate. ➤ On the Applications > TransactionVision > Transaction Topology > Aggregated Topology page, click any of the drill-down icons to the Instance Transaction Topology. See “Aggregated Topology Page” on page 449 for descriptions of the drill-down icons.
Important Information	<p>The Instance Topology page is divided into two parts:</p> <ul style="list-style-type: none"> ➤ Instance Topology table — displays performance metrics for instances that ended in a selected time frame. You can select a row in the table to see the Transaction Topology flow map for the selected Business Transaction instance. <p>You can manipulate the table by adding and removing columns and sorting the table by column values.</p> <ul style="list-style-type: none"> ➤ For information on how to modify the table, see “Common Report Elements” in <i>Reports</i>. ➤ For descriptions of the type of data that is available, see “Instance Topology Table” on page 457. ➤ Instance Transaction Topology flow map — displays a diagram representing the flow of a given instance of a Business Transaction. See “Transaction Topology Flow Map for the Instance Topology” on page 486 for details.
Useful Links	<ul style="list-style-type: none"> ➤ “TransactionVision Data” on page 20 ➤ “Calculating Response Time” on page 455 ➤ “Transaction Topology Flow Map for the Instance Topology” on page 486

Instance Topology Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
<Common report elements>	See “Common Report Elements” in <i>Reports</i> . These settings describe how to set the date and time range, configure a favorite filter, publish a report of the Instance Topology page and annotate the Instance Topology flow map (see “Transaction Topology Flow Map for the Instance Topology” on page 486).
	Click to drill to the Transaction Detail report. For information about this report, see Important Information in “Transaction Tracking Report” on page 562.
	Click to drill to the RUM End User Summary report. For information about this report, see “End User Summary Report” in <i>Using End User Management</i> .
Active Filters	Click to open the “Active Filters Dialog Box for Instance Topology” on page 458.

Instance Topology Table

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Transaction Name	Name of the Business Transaction instance.
Start Time	The time the transaction started.
End to End Response Time (sec.)	End-to-end transaction response time in seconds. See “Calculating Response Time” on page 455 for details on how this metric is calculated.
Response Time of Back End (sec.)	Response time of back-end server level components, such as application servers, message middleware servers, database servers, mainframe transactions and batch jobs.

GUI Element	Description
End User Response Time (sec.)	Response time of end user as measured by Real User Monitor (RUM). If RUM is not configured, this value is blank.
End User Group	Name of the end user group that initiated the transaction.
Value	Represents the monetary value associated with a Transaction.
Completion State	Indicates whether a transaction is currently in-process or completed.
Result State	Indicates whether the transaction failed or succeeded.
SLA State	Indicates whether a transaction violated its SLA.

 **Active Filters Dialog Box for Instance Topology**

Description	Enables you to filter the list of available Business Transaction Instances, metrics, and end user groups. To access: On the Applications > TransactionVision > Transaction Topology > Instance Topology > Instance Topology page, click Active Filters .
Important Information	You can filter the Instance Topology to change what you see in the table and flow map, by: <ul style="list-style-type: none">➤ “Business Transaction Instances” on page 459➤ “Metrics” on page 460➤ “End User Groups” on page 461
Useful Links	“Instance Topology Page” on page 456

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Common report elements>	See “Common Report Elements” in <i>Reports</i> .
Cancel	Click to return to the Instance Topology page without making any changes.
OK	Click to save your filter settings.

Business Transaction Instances

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Common report elements>	See “Common Report Elements” in <i>Reports</i> .
<instance of a Business Transaction>	Check to include all instances of a Business Transaction in the table and flow map.
Browse	Mode in which you can filter the Business Transaction instances.
Include Unclassified Transactions	Click to include unclassified instances of a Business Transaction in the table and flow map.

Metrics

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Completion State	Check one or more of the following options to filter the page by the Completion State of a transaction: <ul style="list-style-type: none">➤ Unknown➤ Processing➤ Completed
Result State:	Check one or more of the following options to filter the page by the Result State of a transaction: <ul style="list-style-type: none">➤ Unknown➤ Success➤ Failed
SLA State:	Check one or more of the following options to filter the page by the SLA State of a transaction: <ul style="list-style-type: none">➤ None➤ Violated

End User Groups

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Common report elements>	See “Common Report Elements” in <i>Reports</i> .
Browse	Mode in which you can filter the end user groups.
City	<p>You can further filter selected end-user groups by selecting a city from the drop-down list.</p> <ul style="list-style-type: none"> ➤ Only selected end-user groups (whether you chose specific end-user groups, or entered an end-user group name or an IP range) that are configured as being in the selected city, are displayed in the report. ➤ To select a city, you must first select a country and, if applicable, a state.
Country	<p>You can further filter selected end-user groups by selecting a country from the drop-down list.</p> <p>Only selected end-user groups (whether you chose specific end-user groups, or entered an end-user group name or an IP range) that are configured as being in the selected country, are displayed in the report.</p>
End user group name(*)	<p>Choose this option to specify a specific end-user group name to include in the filter. Any end-user group whose name includes the text you enter is added to the filter.</p> <p>Note: It is recommended to select the end-user group from the end-user group tree to enable the report to be generated faster.</p>
IP between: <IP address> and <IP address>	Choose this option to specify a range of IP addresses to include in the filter. Enter the starting and ending IP addresses of the required range.

GUI Element	Description
Specific end user groups	<p>Choose this option to select specific end-user groups from the end-user group tree.</p> <p>The end-user group tree displays the existing end-user CIs in the Universal CMDB. Click an entry in the tree to expand it.</p> <p>Select the check box for the end-user groups you want to include in the filter.</p> <p>Note:</p> <ul style="list-style-type: none"> ➤ If an end-user group you configured in End User Management Administration does not appear in the tree, synchronize the Real User Monitor engine in End User Management Administration, or wait until an automatic synchronization occurs, which can take up to an hour. For details, see “End User Management Administration” in <i>Using End User Management</i>. ➤ Descendants of selected items in the tree are automatically included in the filter as well, but do not appear as being selected in the tree.
State	<p>You can further filter selected end-user groups by selecting a state from the drop-down list.</p> <ul style="list-style-type: none"> ➤ Only selected end-user groups (whether you chose specific end-user groups, or entered an end-user group name or an IP range) that are configured as being in the selected state, are displayed in the report. ➤ To select a state, you must first select a country.

Troubleshooting and Limitations

There is a limit on how many instances can be retrieved from the server. If this limit is reached, the following message appears above the table:

The server timed out retrieving data. Only the most recent instances are included in the table below. Please try viewing the data in smaller time ranges by changing the View above.

See Chapter 11, “Queries” for information on changing the value of this limit.

You can see how much data was actually retrieved by sorting the table by Start Time and looking at the oldest instance. You can use this information to figure out what a more reasonable time frame would be. For example, if you had selected **Past Month** and the oldest transaction was 8 days old, then you could change the view to **Past Week** and use the back arrow button to view the past 4 weeks, one week at a time.

17

Transaction Topology Flow Map

This chapter includes the main concepts and reference information for the Transaction Topology flow map used by Dashboard and TransactionVision.

This chapter includes:

Concepts

- Transaction Topology Flow Map Overview on page 465
- Viewing the Transaction Topology Flow Map on page 466
- Time Frame for the Transaction Topology Flow Map on page 473

Reference

- Transaction Topology Flow Map User Interface on page 475

Transaction Topology Flow Map Overview

The Transaction Topology flow map provides a graphical representation of the path taken by a Business Transaction. This flow map includes metrics for the traffic between each flow component and associated KPI data.

Business Availability Center provides three transaction topologies, which are accessed from the following areas:

- Dashboard Flow Map page, see “Introducing Flow Map” in *Using Dashboard* for general information on the Dashboard flow map, and “Transaction Topology Flow Map for the Dashboard” on page 484 for details of the flow map contents.

- TransactionVision Aggregated Topology page, see “Aggregated Topology Page” on page 449 for information about filtering the Business Transactions to view the flow map, and “Transaction Topology Flow Map for the Aggregated Topology” on page 476 for details of the flow map contents.
- TransactionVision Instance Topology page, see “Instance Topology Page” on page 456 for information about filtering the Business Transaction instances to view the flow map, and “Transaction Topology Flow Map for the Instance Topology” on page 486 for details of the flow map contents.

The following sections describe the Transaction Topology flow map functionality and components.

- “Viewing the Transaction Topology Flow Map” on page 466
- “Time Frame for the Transaction Topology Flow Map” on page 473



Viewing the Transaction Topology Flow Map

The following sections describe how to view the Transaction Topology Flow Map:

- “Completed and Currently In-Process Tabs” on page 467
- “Transaction Topology Flow Map Tooltips” on page 467
- “Transaction Topology Flow Map Components” on page 469
- “Transaction Topology Flow Map Structure” on page 470
- “Transaction Topology Flow Map Metrics” on page 471

Completed and Currently In-Process Tabs

TransactionVision sends data samples for completed and in-process Business Transactions to Business Availability Center. The Transaction Topology flow map displays both **Completed** and **Currently In-process** Business Transactions in the respective tabs.

- **Completed:** Displays Business Transactions that completed in the given time frame.
- **Currently In-Process:** Displays Business Transactions that are currently in-process and have not yet completed.

Note: For the Aggregated and Instance Topology flow maps, the currently in-process transactions do not reflect the selected time range. Only transactions that are currently in-process are shown.

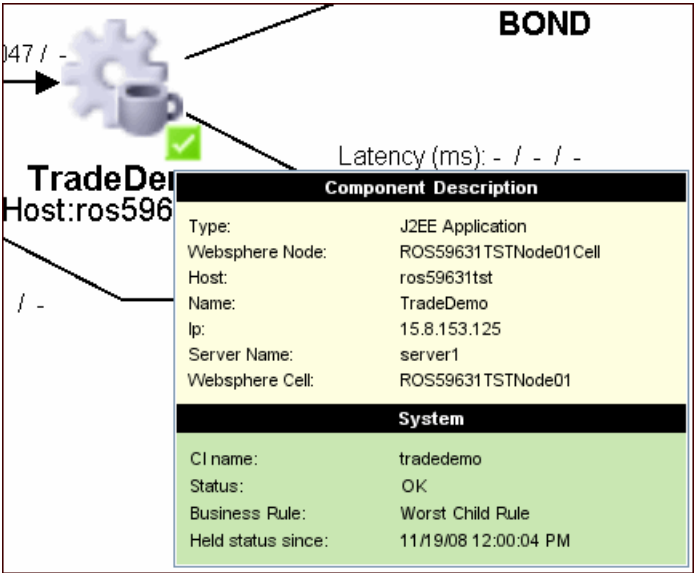
For information on how the time frame is calculated, see “Time Frame for the Transaction Topology Flow Map” on page 473.

For information on the flow map contents in Dashboard, see “Transaction Topology Flow Map for the Dashboard” on page 484.

Transaction Topology Flow Map Tooltips

The call to the TransactionVision machine requests information on the high-level components activated by the Business Transaction instances, for the relevant time frame. (For information on the time frame used for the request, see “Time Frame for the Transaction Topology Flow Map” on page 473.)

The Transaction Topology flow map displays System KPI status of J2EE applications.

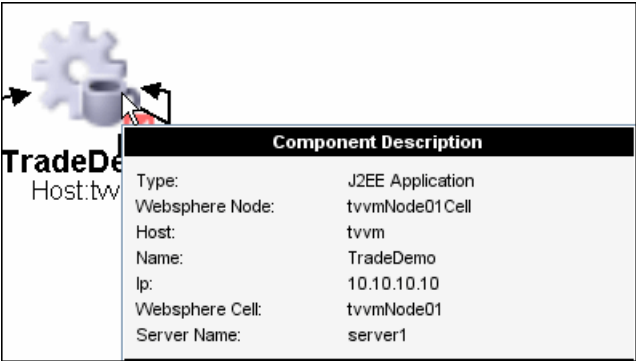


In addition, if you have mapped one or more End User Management transaction CIs to the Business Transaction CI (as described in “Map End-User Transactions to TransactionVision Business Transactions” in *Using End User Management*), then Business Availability Center displays information on the worst status Performance and Availability KPIs (of the KPIs attached to the End User Management transaction CIs), in the tooltip for the **Web Clients** and **End User Group** components.

Transaction Topology Flow Map Components

The Transaction Topology flow map shows the flow of a Business Transaction through the components of the network and target machines. The components are represented in the map by flow component icons. The map may also display a component icon labeled **Web Clients**, representing the Web client (or multiple Web clients) that originated from one or more instances of the transaction. In addition, the flow map may display a component icon labeled End User Group, if a Real User Monitor is deployed and the transaction originates from an end user group.

Each flow component represents an actual software or hardware element that one of the transaction instances passed through, and may include additional data from the software or hardware upper layers. For example, if a transaction instance passed through a J2EE application, the map shows a flow component for the J2EE application, and may include additional data for the application server and the host in the tooltip.



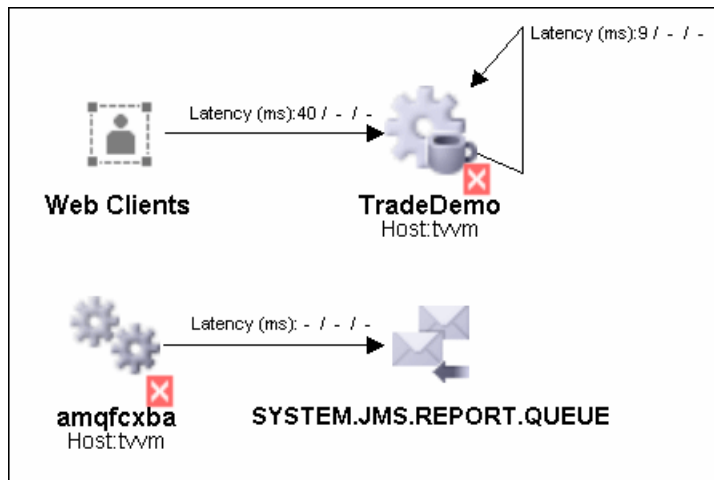
Each component is shown only once, regardless of how many instances of the transaction passed through that component.

The Transaction Topology flow map includes only specific component types (so you do not see all component types identified by TransactionVision). The icon for each component type is determined by the CMDB class type for that component. For a list of the component icons used in the Transaction Topology flow map, see “Transaction Topology Flow Map Component Icons” on page 482.

Transaction Topology Flow Map Structure

Communication between components in the map is represented by links (directed arrows). There is no particular order to the flow as represented in the graph, and the flow is a summarized one for all instances of the Business Transaction. Therefore, the following flow structures between the infrastructure elements are all possible:

- Two links between two components, each link going in the opposite direction.
- A component with a self-directed link.
- A flow with several end points.
- A circular flow.
- Unconnected flow paths (this can occur when there are elements that cannot be interpreted by TransactionVision, for example, if the technology is not supported).



Transaction Topology Flow Map Metrics

If the Transaction Topology flow map is set to show volume, value, or latency data (according to the selections made for the **Display Data** option in the Transaction Topology flow map toolbar), then the Transaction Topology flow map displays these metrics, broken down by transaction state categories, along the edge of each link arrow (the information can also be viewed in the tooltip for the link). For more information about the transaction state categories, see “Business Transaction State Categories” on page 21.

The displayed metrics are supplied by TransactionVision, and relate to the flow of transaction instances (either completed or currently in-process, as relevant) between the two components connected by the arrow.

The metrics provide the following information:

- **Volume.** Displays a count of the transaction instances that either passed between the two components or are currently in-process of being passed.
- **Value.** Displays the total monetary value associated with the transaction instances that either passed between the two components or are currently in-process of being passed.
- **Latency.** Displays the average latency (response time) for the transaction instances that either passed between the two components or are currently in-process of being passed (in milliseconds).

For each metric, data is displayed across three transaction state categories, Successful, Late, and Failed. If a metric is not applicable for a link, or if there is no criteria set in TransactionVision for late status or failed status, then the relevant metric shows a dash.

Following are two examples of information that may be displayed in tooltips:

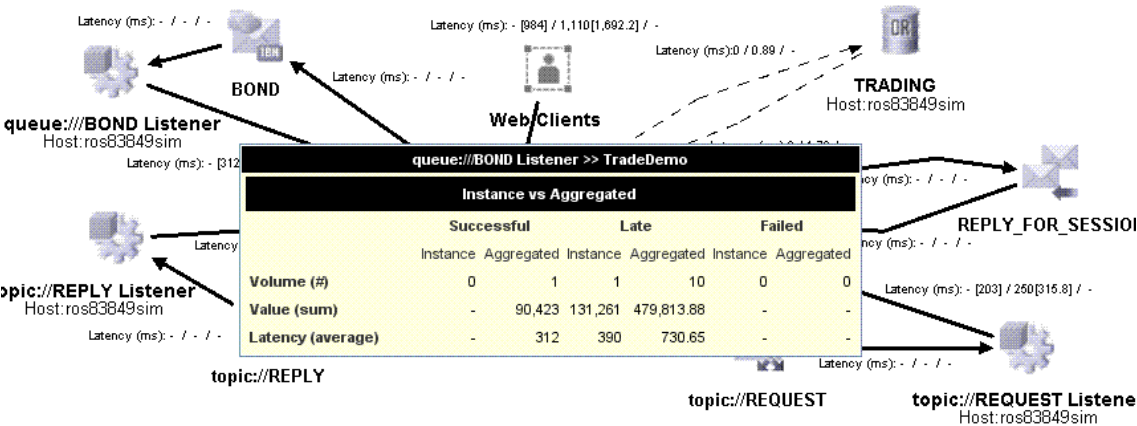
Example of Tooltip Metrics for an Aggregated Topology

This example shows a tooltip for an Aggregated Topology link showing the Late and Failed metrics with no criteria set. See “Transaction Topology Flow Map for the Aggregated Topology” on page 476 for more information.

topic://REQUEST Listener >> TradeDemo			
	Successful	Late	Failed
Volume (#)	2	0	0
Value (sum)	\$3,681.98	-	-
Latency (average)	187.5 ms	-	-

Example of Tooltip Metrics for an Instance Topology

This example shows a tooltip for an Instance Topology link that compares the metrics of the path taken by a Business Transaction instance with the path taken by other instances of that Business Transaction. The dotted line represents the Aggregated Topology of the Business Transaction with respect to the Instance Topology, which is a solid line, See “Transaction Topology Flow Map for the Instance Topology” on page 486 for more information.



Time Frame for the Transaction Topology Flow Map

The way the time frame is set for displaying data in the Transaction Topology flow map varies between the Dashboard, Aggregated Topology and Instance Topology.

Dashboard

The time range shown in the flow map aligns with the rest of the data shown in the Dashboard.

When you open a **Completed** Transaction Topology flow map, the query to TransactionVision requests all completed transactions over a predefined time period.

Note: If there is no data for the query period, an error message is displayed when the Transaction Topology flow map tries to open.

When you open a **Currently In-Process** Transaction Topology flow map, the data in the flow map is not automatically updated in the Dashboard. Click the **Refresh** button to update the current data.

Aggregated Topology

The Aggregated Topology displays the path taken by all Business Transaction instances in the given time period for completed transactions. The Aggregated Topology can only be viewed for a 15-minute time period.

Note: The flow map for the in-process Business Transactions does not reflect the selected time range. Only transactions that are currently in-process are shown.

When you set the **View** to **Past 15 minutes**, the current time is used for the end time and the range is always 15 minutes (see Chapter 15, “Aggregated Topology” for details on setting the time). Since the aggregation happens on a 5-minute boundary, the time period is automatically adjusted to align with 5-minute time periods. For example, if the current time is 6:18, the data displayed is within the time range of 6:03 to 6:18.

To include late arriving events, there is a delay in the aggregation process. This delay is configurable on the **Admin > TransactionVision > Business Availability Settings** page using the **Time slice shift in job (in seconds)** option (see “HP Business Availability Settings Page” on page 379 for details). The default value for the delay is 6 minutes.

Due to this delay, the topology flow map does not show data for all the transactions in the past 15 minutes. For example, at 6:18, the Aggregated Transaction Topology flow map only contains data for transactions that occurred between 6:00 and 6:10. The Business Transactions that occurred between 6:10 and 6:15 have not been aggregated yet.

Instance Topology

The Instance Topology flow map overlays the topology for a single instance on top of the Business Transactions aggregated topology. The overlaid topology allows you to compare the path taken by a Business Transaction instance with the path taken by other instances of that Business Transaction in a 15-minute window that surrounds the start time of the selected instance. For example, if the instance started at 5:13, the time range for the aggregate data is 5:05 -5:20.

Note: The flow map for the in-process instances of a Business Transaction does not reflect the selected time range. Only instances that are currently in-process are shown.

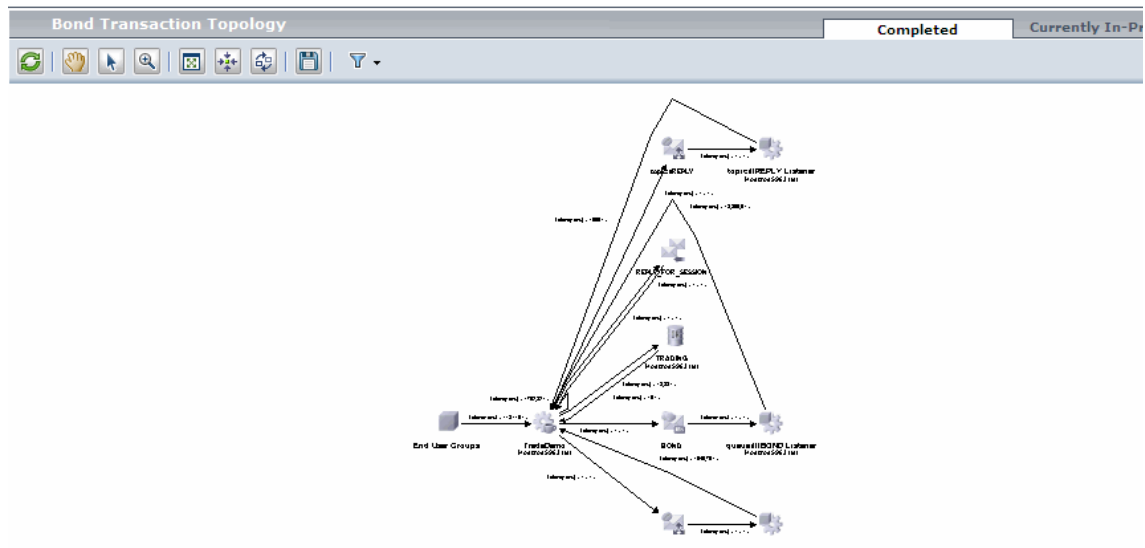
Transaction Topology Flow Map User Interface

This section describes:

- Transaction Topology Flow Map for the Aggregated Topology on page 476
- Transaction Topology Flow Map for the Dashboard on page 484
- Transaction Topology Flow Map for the Instance Topology on page 486

🔍 Transaction Topology Flow Map for the Aggregated Topology

The following is an example of an Aggregated Transaction Topology flow map.



Description	<p>Displays a graphic representation of the aggregated flow for the selected Business Transaction.</p> <p>The Transaction Topology flow map includes metrics for the traffic between each flow component and associated KPI data.</p> <p>To access: Select Applications > TransactionVision > Transaction Topology > Aggregated Topology. The Transaction Topology flow map automatically displays for the selected transaction.</p>
--------------------	---

Important Information	<ul style="list-style-type: none"> ➤ On opening, the Transaction Topology flow map displays information corresponding to completed Business Transactions within the active time frame. For more information on how the active time frame is calculated, see “Response Time Calculation and Breakdown” on page 447. <p>NOTE: The flow map for the in-process Business Transactions does not reflect the selected time range. Only transactions that are currently in-process are shown.</p> <ul style="list-style-type: none"> ➤ The title bar displays the name of the Business Transaction. ➤ The flow for the transaction is indicated by icons representing flow components, and arrows representing the flow links. See “Transaction Topology Flow Map Component Icons” on page 482.
Useful Links	<ul style="list-style-type: none"> ➤ “Completed and Currently In-Process Tabs” on page 467 ➤ “Time Frame for the Transaction Topology Flow Map” on page 473 ➤ “Transaction Topology Flow Map Tooltips” on page 467 ➤ “Transaction Topology Flow Map Components” on page 469 ➤ “Transaction Topology Flow Map Structure” on page 470 ➤ “Transaction Topology Flow Map Metrics” on page 471 ➤ “Transaction Topology Flow Map Component Icons” on page 482




Transaction Topology Flow Map Tabs







The following tabs are included:

GUI Element	Description
Completed	Opens the TransactionVision Transaction Topology flow map for completed instances of the Business Transaction. For details, see “Completed and Currently In-Process Tabs” on page 467.
Currently In-Process	Opens the TransactionVision Transaction Topology flow map for instances of the Business Transaction that are in process. For details, see “Completed and Currently In-Process Tabs” on page 467.

Toolbar Options



The following toolbar options are included:

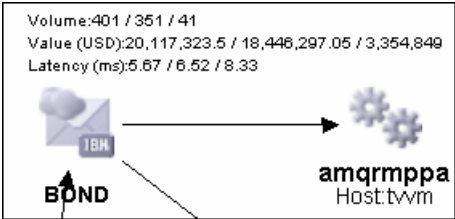
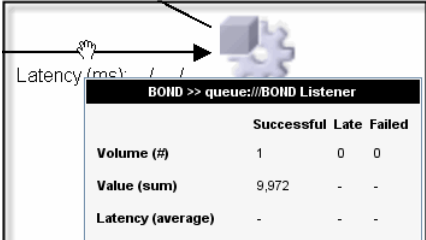
GUI Element	Description
	Click the Refresh button to update the Transaction Topology flow map to show the most current data for the selected time frame. After refreshing, all default settings for the Transaction Topology flow map are restored.
	Click the Pan button, then click and drag in the Transaction Topology flow map to view different areas of the map.
	Click the Select button, then: <ul style="list-style-type: none">➤ click an entity in the map to select it.➤ click and drag an entity in the map to reposition it within the Transaction Topology flow map. For example, you can click on the elbow of a link arrow and drag it to a different location.➤ click and drag over an area to select the entities within that area. Click again outside the selected area to deselect.

GUI Element	Description
	Click the Zoom button, then click and drag in the map to zoom in and out.
	Click the Fit to window button to resize the map to fit the available viewing space.
	Click the Move selected to center button to move the currently selected items to the center of the viewing area.
	Click the Rearrange button to restore the map to its default appearance.
	Click the Save as image button to open a dialog box where you define the target location and required settings for saving the Transaction Topology flow map as an image.
	<p>Click the Display Data button to open a list of data display options. Select the check box for each option that you want displayed in the Transaction Topology flow map.</p> <ul style="list-style-type: none"> ➤ CI Additional Data. Displays the name or IP of the host running the component (when this information is available) under each relevant component icon. ➤ Value. Displays value metrics for each link, alongside each link arrow. ➤ Volume. Displays volume metrics for each link, alongside each link arrow. ➤ Latency. Displays latency metrics for each link, alongside each link arrow. <p>For more information on these options, see “Transaction Topology Flow Map Metrics” on page 471.</p> <p>Note: If an option is not selected, you can still see the metrics information in the tooltip for the link, or the host information in the tooltip for a component.</p>

Transaction Topology Flow Map Components

The following elements are included (unlabeled GUI elements are shown in angle brackets):








GUI Element	Description
<p><Flow component icon></p> <p>Example:</p> 	<p>Each flow component (for example, Application, Servlet, MQ Queue, and so forth) is represented by an icon.</p> <p>For a reference list of the icons used in the Transaction Topology flow map, see “Transaction Topology Flow Map Component Icons” on page 482.</p> <p>If you selected to show CI Additional Data in the Display Data list, then the Transaction Topology flow map displays the name or IP of the host running the component (when this information is available) under each relevant component icon.</p> <p>Tooltip: Hold the pointer over a component icon to display a tooltip containing information for that component, including KPI information. For a J2EE application, the tooltip displays System KPI information.</p> <p>For the Web Clients and End User Group components, the tooltip displays KPI information relating to end user transactions that are mapped to the Business Transaction.</p> <p>For more information on the data in the tooltips, see “Transaction Topology Flow Map Tooltips” on page 467.</p>
<p><Status indicator></p> <p>Example:</p> 	<p>Status indicators are displayed in the bottom-right corner of a J2EE application component icon. The status indicator shows the worst status held by the KPIs attached to the J2EE CI.</p> <p>For more information, see “Transaction Topology Flow Map Tooltips” on page 467.</p>





GUI Element	Description
<p><Link arrow> Example:</p>	<p>The link arrows, each connecting two component icons, indicate the direction of the transaction flow.</p> <p>If the Transaction Topology flow map is set to show volume, value, or latency data (according to the selections made in the Display Data list), then the Transaction Topology flow map displays these metrics, broken down by three transaction state categories, along the edge of each link arrow. The format <successful>/<late>/<failed> is used to display the metrics across each transaction state category, for example: Volume: 100/2/1</p> <div><p>Volume:401 / 351 / 41 Value (USD):20,117,323.5 / 18,446,297.05 / 3,354,849 Latency (ms):5.67 / 6.52 / 8.33</p></div> <p>For an explanation of the metrics and transaction state categories, see “Transaction Topology Flow Map Metrics” on page 471.</p> <p>Tooltip: Hold the pointer over a link arrow to display a tooltip that shows the metrics broken down by transaction state as shown in the following example:</p> <div></div>

Transaction Topology Flow Map Component Icons

The following table lists the component icons that are used for the Business Transaction Topology.

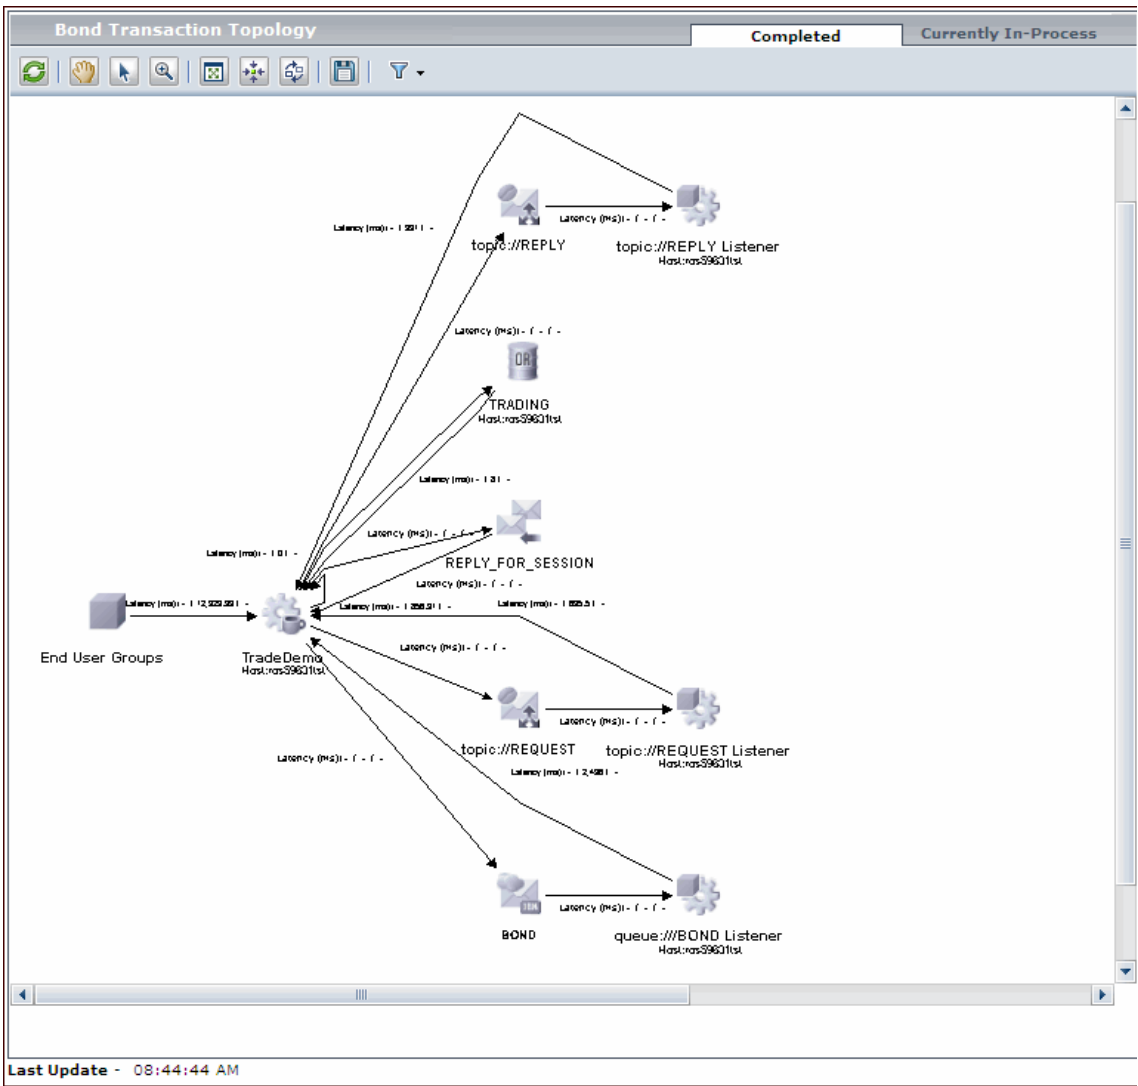
Note: The database component is not included in the table—the icon used for database varies, depending on the database type.

Icon	Component Name
	Program
	J2EE Application
	Web Clients
	IBM WebSphere MQ Queue
	IBM WebSphere MQ Cluster Queue
	JMS Topic
	JMS Queue

Icon	Component Name
	CICS Transaction
	IMS Program
	Batch Job
	End User Groups

Transaction Topology Flow Map for the Dashboard

The following is an example of an aggregated Transaction Topology flow map for the Dashboard.



Description	<p>Displays a graphic representation of the aggregated flow for the selected Business Transaction.</p> <p>The Transaction Topology flow map includes metrics for the traffic between each flow component and associated KPI data.</p> <p>To access: In Applications > Dashboard > Console, display the Business Transactions view in View Explorer, select a Business Transaction to enable the Flow Map tab, and select the Flow Map tab.</p>
Important Information	<ul style="list-style-type: none"> ➤ On opening, the Transaction Topology flow map displays information corresponding to the active time frame. For more information on how the active time frame is calculated, see “Response Time Calculation and Breakdown” on page 447. ➤ The title bar displays the name of the Business Transaction. ➤ The last updated time displays at the bottom of the Transaction Topology flow map in the Dashboard. ➤ The flow for the transaction is indicated by icons representing flow components, and arrows representing the flow links. See “Transaction Topology Flow Map Component Icons” on page 482.
Useful Links	<ul style="list-style-type: none"> ➤ “Completed and Currently In-Process Tabs” on page 467 ➤ “Time Frame for the Transaction Topology Flow Map” on page 473 ➤ “Transaction Topology Flow Map Tooltips” on page 467 ➤ “Transaction Topology Flow Map Components” on page 469 ➤ “Transaction Topology Flow Map Structure” on page 470 ➤ “Transaction Topology Flow Map Metrics” on page 471 ➤ “Transaction Topology Flow Map Component Icons” on page 482

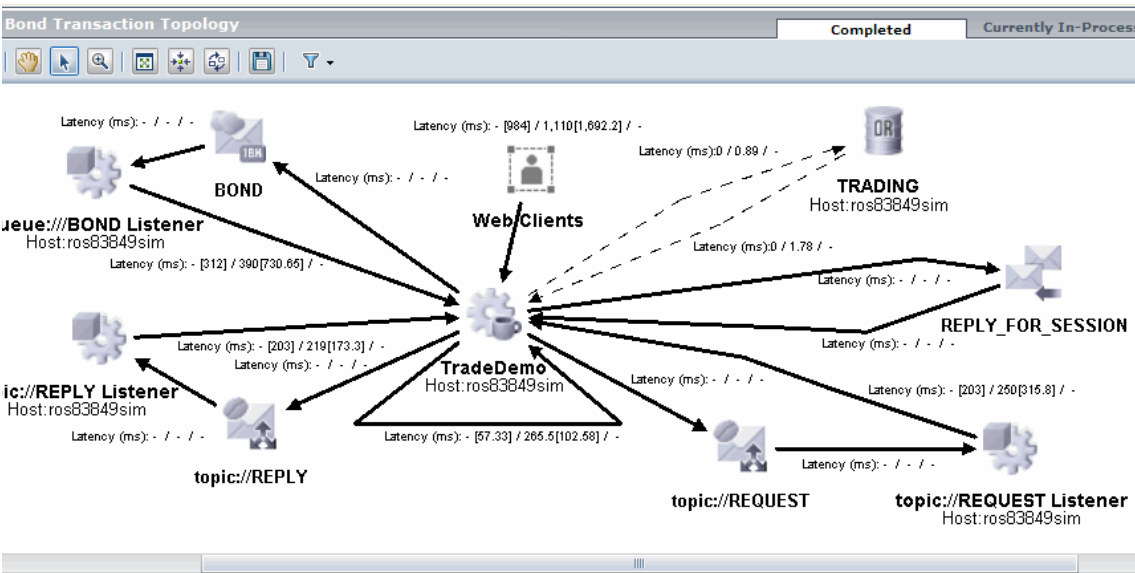
Transaction Topology Flow Map for Dashboard Settings

See “Transaction Topology Flow Map for the Aggregated Topology” on page 476 for a complete description of:

- Transaction Topology Flow Map Tabs
- Toolbar Options
- Transaction Topology Flow Map Components
- Transaction Topology Flow Map Component Icons

Transaction Topology Flow Map for the Instance Topology

The following is an example of an Instance Transaction Topology flow map.



Description	<p>Displays a flow map similar to the Aggregated Transaction Topology flow map, except that this flow map displays the flow of a specific instance of a Business Transaction, overlaid on the Aggregated Topology for the Business Transaction CI.</p> <p>The Transaction Topology flow map includes metrics for the traffic between each flow component and associated KPI data.</p> <p>To access:</p> <ul style="list-style-type: none">➤ Select Applications > TransactionVision > Transaction Topology > Instance Topology. Select a transaction instance and click Generate.➤ On the Applications > TransactionVision > Transaction Topology > Aggregated Topology page, click any of the drill-down icons to the Instance Transaction Topology. See “Aggregated Topology Page” on page 449 for descriptions of the drill-down icons.
--------------------	---

Important Information	<ul style="list-style-type: none">➤ On opening, the Transaction Topology flow map for Instance Topology displays information corresponding to completed instances of the Business Transaction within the active time frame. This is the time frame that matches the data collection period. For details on the active time frame calculation see “Response Time Calculation and Breakdown” on page 447. <p>NOTE: The flow map for the in-process instances of the Business Transactions does not reflect the selected time range. Only instances that are currently in-process are shown.</p> <ul style="list-style-type: none">➤ The title bar displays the name of the Business Transaction instance.➤ The flow for the transaction is indicated by icons representing flow components, and arrows representing the flow links. See “Transaction Topology Flow Map Component Icons” on page 482
Useful Links	<ul style="list-style-type: none">➤ “Instance Topology” on page 453➤ “Completed and Currently In-Process Tabs” on page 467➤ “Time Frame for the Transaction Topology Flow Map” on page 473➤ “Transaction Topology Flow Map Tooltips” on page 467➤ “Transaction Topology Flow Map Components” on page 469➤ “Transaction Topology Flow Map Structure” on page 470➤ “Transaction Topology Flow Map Metrics” on page 471➤ “Transaction Topology Flow Map Component Icons” on page 482

Transaction Topology Flow Map for Instance Topology Settings

See “Transaction Topology Flow Map for the Aggregated Topology” on page 476 for a complete description of:

- Transaction Topology Flow Map Tabs
- Toolbar Options
- Transaction Topology Flow Map Components
- Transaction Topology Flow Map Component Icons

18

Component Topology Analysis

This chapter includes:

Concepts

- Component Topology Analysis Overview on page 492
- Viewing the Component Topology Analysis on page 492
- Edges on page 494
- Set Background Color on page 501
- Modify Layouts and Properties on page 502
- Printing the Component Topology Analysis Graphs on page 508
- Adjusting the Zoom Level on page 512
- Configure Edge and Node Labels on page 514
- Modify Component Groupings on page 519

Tasks

- Viewing Statistics on page 524

Reference

- Component Topology User Interface on page 525

Component Topology Analysis Overview

The Component Topology Analysis provides a detailed graphic representation of the interaction between all system components for which Sensors collect event information. The Analyzer correlates events across host, program, and thread boundaries. These correlated events are used to draw a visual map of the flow of information throughout the system.

Viewing the Component Topology Analysis

You can view the Component Topology Analysis graph in one browser using one of two modes, open multiple views of the graph, and move components to adjust the view.

View Modes

TransactionVision provides two modes for the Component Topology Analysis: dynamic and static.

- **Dynamic mode.** In Dynamic mode, the Component Topology graph is created by comparing each event in the project to the current query and correlating events that match the query. This mode enables you to structure your query to view very specific information, but it can be time-consuming to perform the analysis and display the graph.
- **Static mode.** Rather than evaluating all project events against a query, the static view is created from statistics gathered during event collection. This enables the Analyzer to draw the graph much faster than in dynamic mode. However, it limits the granularity of event data because the static view can only distinguish between the program level, rather than the program instance level used in the dynamic view. The static view also limits queries to statistics occurring within a specified time range. During event collection, statistics for the static mode display are stored in a statistics cache. Periodically, this cache is written to the database, making the results available for viewing the Component Topology Analysis. Entries in the `<TVISION_HOME>/config/services/StatisticsCache.properties` file, control the operation of the cache. For information about the cache properties, see the the *HP TransactionVision Deployment Guide* PDF.

Switching Between View Modes

By default, the Component Topology Analysis view is created in static mode. Click **Switch to Dynamic Mode** to display the Dynamic Component Topology graph. To change from dynamic mode to static mode, click **Switch to Static Mode** on the Dynamic Component Topology graph.

To use a different query in dynamic mode, select the query from the query menu in the upper right corner of the display. Select **New Query** from the query menu to create a new query, or select **Edit Query** to modify the current query. For information about creating or modifying queries, see Chapter 11, “Queries.”

Moving Components

To move any component in the graph, select the component in the view area and drag it to the desired location. TransactionVision automatically adjusts all edges to and from the component.

Opening Multiple Views

You may open new browser windows using multiple Component Topology Analysis windows concurrently, using different queries to display different result sets.

If you open a Component Layout Analysis view in a new browser window, any changes you make in one window are reflected in the other window when it is refreshed. If you start a new browser instance and log into TransactionVision in the new browser, however, a new TransactionVision session is started. If you use multiple sessions, changes you make in one session affect that session only.

Show/Hide the Status Bar

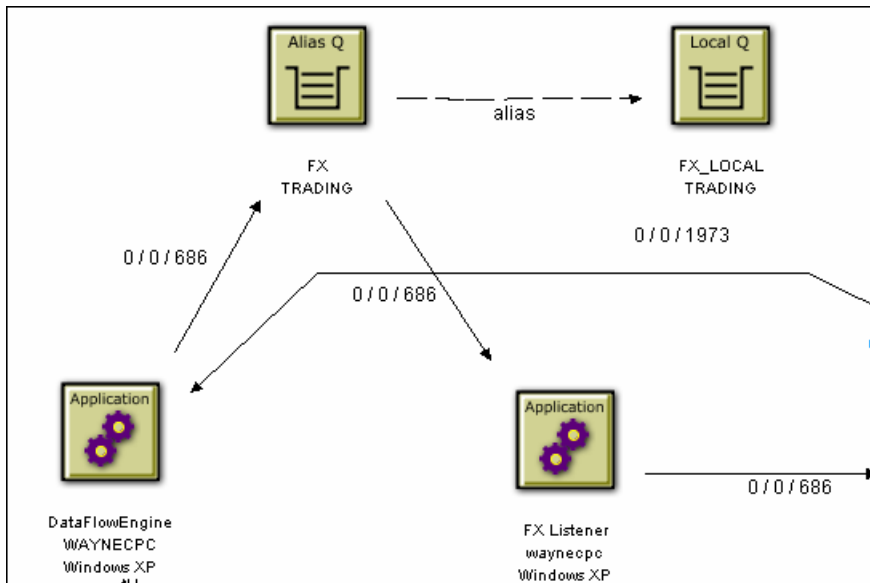
Use the **Show/Hide** arrow buttons to show, hide, or resize the status message pane.

To hide the status bar, click the down arrow button. To show a hidden status bar, click the up arrow button. To resize the status bar, point the cursor to either arrow button so that the cursor changes to a double arrow and drag.

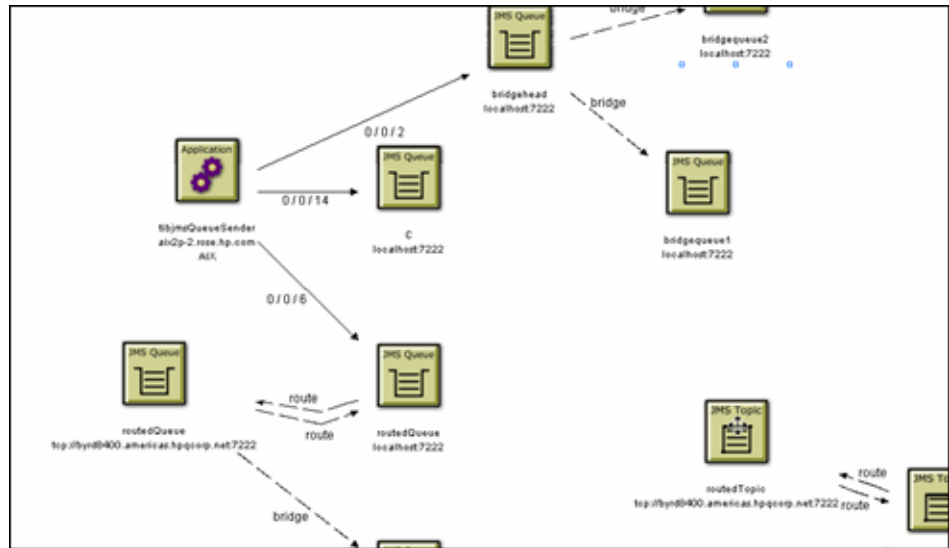
Edges

Resource icons represent system components and the connecting lines, called edges, show the relationship between the resource icons. For WebSphere MQ, JMS, and CICS events, solid edges represent message flows. For servlet and EJB events, solid edges represent control flow.

Dashed edges represent a relationship between icons, such as between a remote, alias, or model queue and the local queue it is associated with. In the following example, the DataFlowEngine and FX Listener programs write to and read from the FX queue, respectively. The FX queue is an alias queue for the local FX_LOCAL queue.

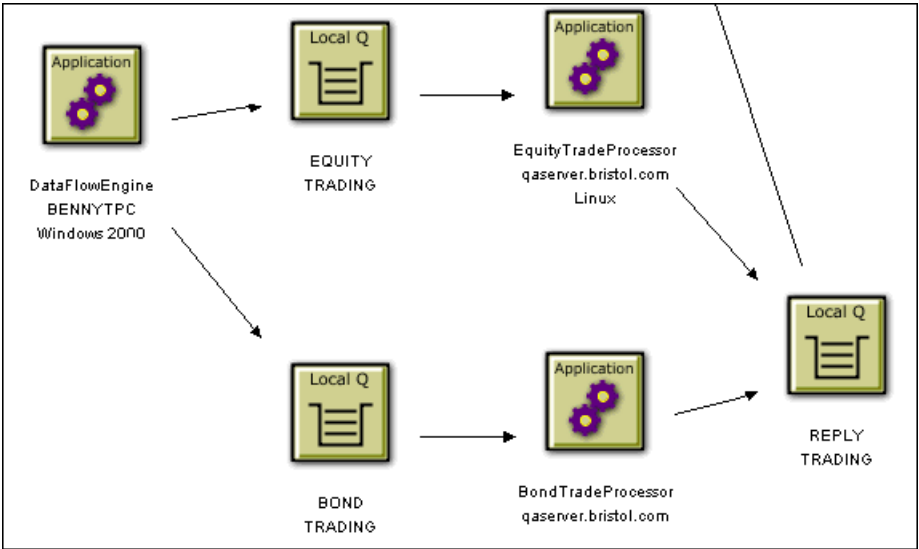


For TIBCO EMS queue objects, edges represent route and bridge relationships between objects, as in the following example:



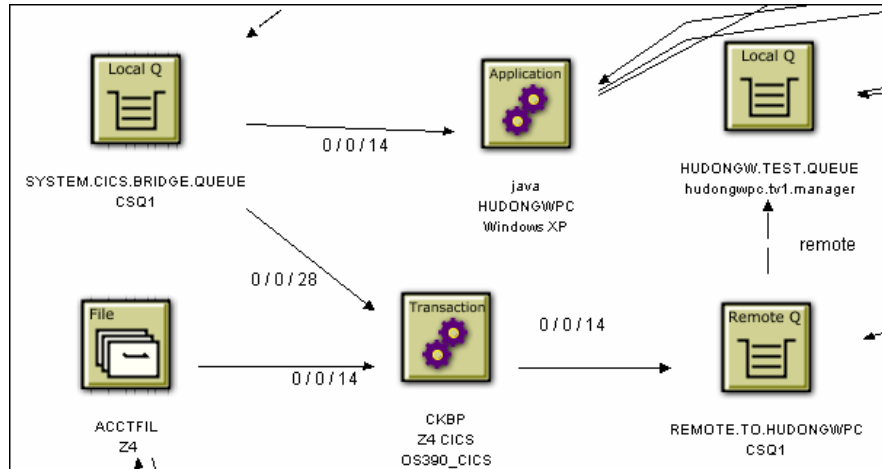
WebSphere MQ and JMS Events

For WebSphere MQ and JMS events, an edge represents message flow. For example, the following diagram shows that the DataFlowEngine application puts messages on the EQUITY and BOND queues of the TRADING queue manager. The EquityTradeProcessor and BondTradeProcessor programs retrieve messages from the EQUITY and BOND queues, respectively, and put a message on REPLY.



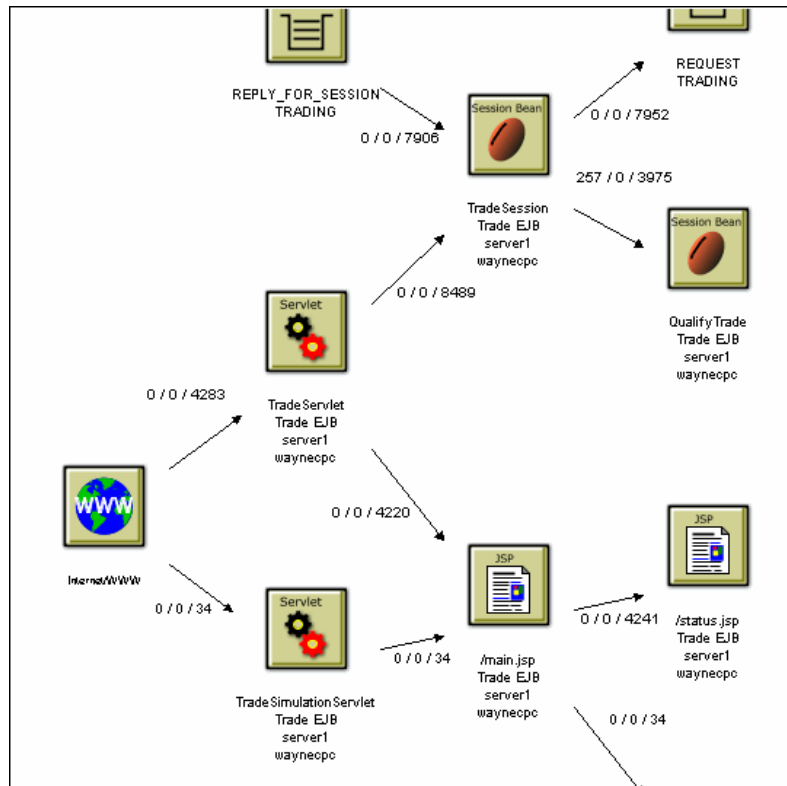
CICS Events

Like WebSphere MQ and JMS events, edges represent message flow for CICS events. For example, in the following diagram, the CKPB CICS transaction reads from the SYSTEM.CICS.BRIDGE.QUEUE queue and the ACCTFIL file. It then writes to the REMOTE.TO.HUDONGWPC queue, which is a remote queue for HUDONG.TEST.QUEUE..



Servlet and EJB Events

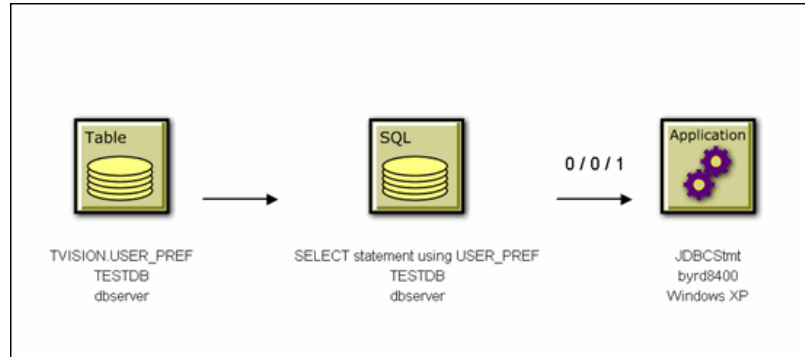
For servlet and EJB events, edges represent control flow. In the following example, the web page calls the TradeServlet and TradeSimulationServlets. TradeServlet calls the TradeSession EJB, while TradeSimulationServlet calls a Java server page.



JDBC Events

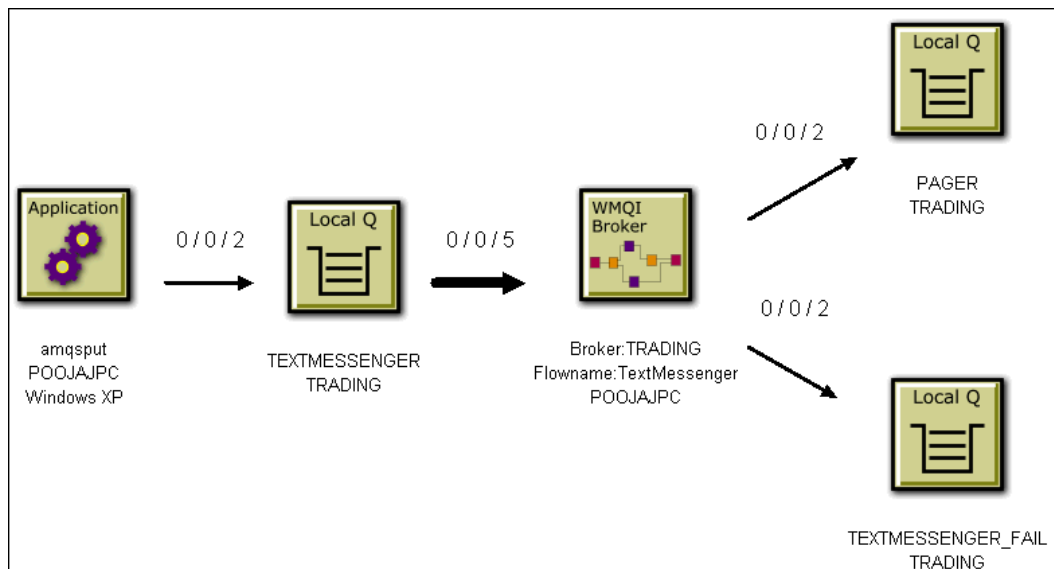
For JDBC events, edges represent access to database resources. An application has a link to an SQL object, which represents the type of SQL statement and which tables it accesses. Various statistics on this operation are available on this edge. Additionally, an SQL object has links to database table objects that were accessed by the SQL statement.

In the following example, you see the JDBCStmt application execute an SQL statement that does a SELECT on the table TVISION.USER_PREF.



WebSphere Business Integration Message Broker (WMIMB) Events

In the following example, the component label of the broker node shows the WBI broker name and the corresponding message flow name instead of listing the name of the broker process.

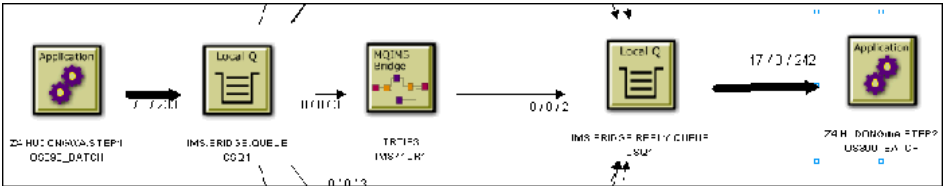


The following table shows the node label meanings for WBI brokers with respect to the various distributed program grouping criteria in the Component Topology Analysis:

Grouping Criteria	Node Label
Individual Thread	Broker and Message Flow name
Individual Process	Broker name
Program Name/Host Combination	Fixed string "WBI Broker"
Program Name	Fixed string "WBI Broker"
Host	N/A

WebSphere MQ-IMS Bridge Events

The following example shows a batch job HUDONGWA with a job step, STEP1, putting messages onto a MQ-IMS bridge input queue. On receiving these messages on the input queue, the WebSphere MQ-IMS bridge invokes an IMS transaction TRTIB3 in IMS job IMS71CR1. That transaction places reply messages on the IMS.BRIDGE.REPLY.QUEUE, which is read by the job step, STEP2 of the job HUDONGWA.



Stand-alone Java Applications

When monitoring JMS events from stand-alone Java applications, the Component Layout view displays the top level class name as the program name. The top level class name is the outmost name in the stack trace. For example, the class name in the following stack trace is taken from the last line in the following example:

```
at java.lang.Thread.dumpStack(Thread.java:993)
at Test.<init>(Test.java:29)
at Test.main(Test.java:62)
```

However, this class does not necessarily contain the public static void main(String[] args). Instead, it might be a class derived from java.lang.Thread, which runs in other threads than the main thread.

If the correct program name is not displayed, set the com.bristol.tvision.programname system property in the command line when running the stand-alone application. For example, to set the program name to MyProgram, start it as follows:

```
java -Dcom.bristol.tvision.programname=MyProgram ...
```

Set Background Color

You can specify the background color of the Component Topology Analysis graph. To specify a background color, choose the **Layout > Background Color** menu item. The Choose Graph Background Color dialog appears.

There are three ways to select a background color:

- On the Swatches page, click the desired color. The preview area shows what your choice looks like.
- On the HSB page, click the desired color in the color area. Use the sliding bar to adjust the hue, saturation, and brightness levels individually.
- On the RGB page, enter specific values for red, green, and blue, or click the desired value on the scale.

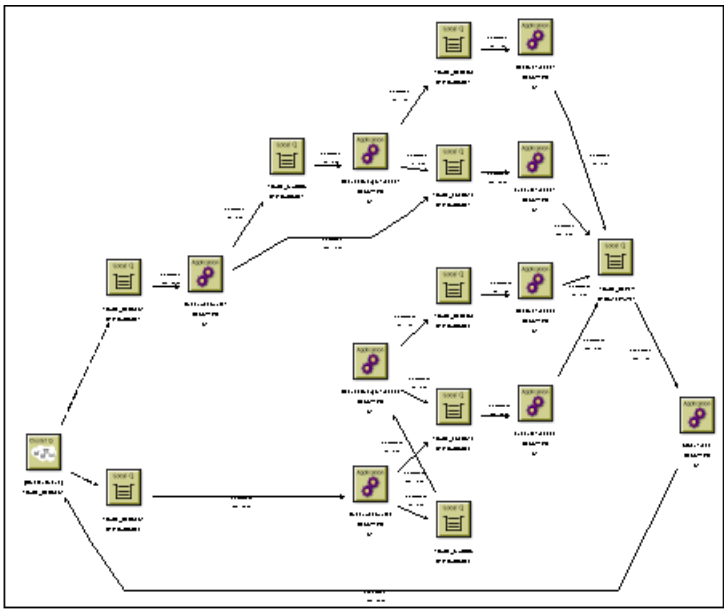
Click **OK** to make your changes active, **Cancel** to close the dialog without activating your changes, or **Reset** to reset the background color to the default.

 **Modify Layouts and Properties**

By default, the Component Topology Analysis is arranged in a hierarchical layout according to the flow of messages through your system. While this layout is typically the best choice for e-business systems and therefore recommended, you may choose to display the Component Topology Analysis in a symmetric layout. Note that using a different layout may make it more difficult to follow the flow of messages throughout your system.

Hierarchical Layout

To select the hierarchical layout (if it is not the current layout), choose the **Layout > Hierarchical Style** menu item. Components are arranged in a hierarchical organization based on the flow of messages through your system, as in the following example.



To modify layout properties, choose the **Layout > Layout Properties** menu item to open the Layout Properties dialog box. Make desired changes on the Hierarchical page, then click **OK**.

Orientation

By default, components are arranged from left to right, but you can change the orientation from top to bottom, bottom to top, or right to left.

- **Left To Right.** Positions components so that the roots of the graph are near the left of the page, and the leaves of the graph are near the right. The levels are vertical.
- **Top To Bottom.** Positions components so that the roots of the graph are near the top of the page, and the leaves of the graph are near the bottom. The levels are horizontal.
- **Right To Left.** Positions components so that the roots of the graph are near the right of the page, and the leaves of the graph are near the left. The levels are vertical.
- **Bottom To Top.** Positions components so that the roots of the graph are near the bottom of the page, and the leaves of the graph are near the top. The levels are horizontal.

Level Alignment

Level alignment refers to the alignment of components on the same level. It is similar to text alignment within a paragraph. If Orientation is horizontal (Top To Bottom or Bottom To Top), you can set the Level Alignment as either Top, Center or Bottom. If Orientation is vertical (Left To Right or Right To Left), you can set the Level Alignment as either Left, Center or Right.

Spacing

The concept of frames is used to control the margins around components. These options allow you to set the spacing between levels of components and between component within the same level in a graph by adjusting the frames that surround them.

- **Variable Level Spacing.** When enabled, variable level spacing considers the density of edges between adjacent levels when adjusting the level spacing. For example, if too many edges cross at the same point between levels, then the level spacing is automatically increased. Consequently, not all levels would have the same spacing. Variable level spacing is usually used in combination with orthogonal edge routing.
- **Proportional Spacing.** Sets the spacing around each node to a fractional value (between 0 and 1) based on the node size.
- **Constant Spacing.** Allows the spacing around each node to be a constant value, independent of the size of the node.
- **Between Levels.** This field allows you to set the spacing between each level in the graph. In left-to-right and right-to-left views, levels run vertically. In top-to-bottom and bottom-to-top views, levels run horizontally.
- **Between Nodes.** This field allows you to set the spacing between adjacent components within a level.

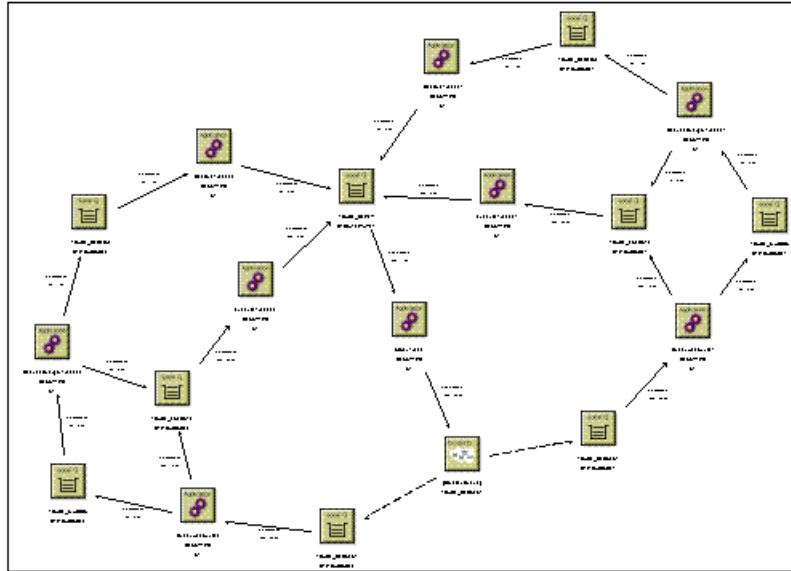
Edge Routing

The following options allow you to tailor edge routing:

- **Orthogonal Routing.** When enabled, each edge is drawn as a sequence of line segments that run parallel to the x- and y-axes. Edges have bend points of 90 degrees only. The orthogonal routing style works best in combination with variable level spacing and ports.
- **Straight Edges.** When enabled, this setting merges the incoming edges of a node into a single fork-like structure. Clearer drawings result provided the graph is not too dense. This feature is only available when orthogonal routing is enabled.

Symmetric Layout

To select the symmetric layout (if it is not the current layout), choose the **Layout > Symmetric Style** from the shortcut menu. The symmetric layout distributes components evenly with very few edge crossings, as in the following example.



To modify layout properties, choose the **Layout > Layout Properties** menu item to open the Layout Properties dialog box. Make desired changes on the Symmetric page, then click **OK**.

Spring Options

These controls allow you to set the options relating to the spring model of the graph.

- **Node Spacing.** This value is used in determining the white space between components. It specifies the length of each edge as a function of the width and height of the components to which the edge connects. A node spacing of zero specifies that edge lengths be just long enough to keep components from touching — there is no white space between components. The larger the node spacing value, the greater the amount of white space between components.

- **Random Seed.** The initial layout of the graph is keyed to a random number sequence. The value set in this field is used to generate the random number sequence to set initial coordinates for each components. Each different non-negative value can produce a different layout. This field is active when Incremental Layout is not checked.



General Properties

The General page of the Properties dialog enables you to specify settings that are independent of the layout style. To modify general layout properties, choose the **Layout > Layout Properties** menu item to open the Layout Properties page. Make desired changes on the General page, then click **OK**.

Spacing Model

The concept of frames is used to control the margins around physical objects. These options control which model is used in determining the amount of white space around the perimeter of objects in the graph. Two spacing models are available, the Proportional Spacing model and the Constant Spacing model. Additional layout-specific spacing options are available, some of which are affected by the spacing model selected here.

- **Proportional Spacing.** When enabled, proportional spacing is used to determine the amount of white space around the perimeter of an object based on the area of the object. For example, if a component's width and height increase, the white space between the component and an adjacent component also increases. The margins of the graph can be controlled through the Graph Margin Spacing fields below. You can set independent proportional spacings for the left, right, bottom, and top sides of the graph. Note that additional spacing controls are available on the other dialog pages.
- **Constant Spacing.** When enabled, a constant space is allocated between like objects in a graph. For example, in hierarchical layout a fixed distance is set between each pair of adjacent components within a level, regardless of the width and height of the nodes. The margins of the graph can be controlled through the Graph Margin Spacing fields below. You can set independent constant spacing values for the left, right, bottom, and top sides of the graph. Note that additional spacing controls are available on the other dialog pages.

Graph Margin Spacing

This group of fields affects the margins for graphs. Separate frame values are stored for the Constant Spacing and Proportional Spacing models, but only one spacing model is in effect at any given time.

Labeling

This option allows you to tailor the automatic positioning of edge labels.

Select **Incremental Layout** to maintain the relative positioning of labels from a graph's previous layout whenever the graph is modified. Incremental Layout results in graphs that maintain the same overall form when changes are made.

Disconnected Properties


The generic term “disconnected objects” is used to refer to both disconnected components and connected subgraphs, unless otherwise noted. A disconnected component is one which has no edges connected to it. A connected subgraph is a graph in which there is an undirected path between each pair of components. In other words, each component can be reached from all other components by following the edges that connect them, regardless of edge direction. If there is no such path between any two components, then either the components belong to different connected components or one of the components is a disconnected node.

To modify disconnected layout properties, choose the **Layout > Layout Properties** menu item to open the Layout Properties dialog box. Make desired changes on the Disconnected page, then click **OK**. These options only have effect if more than one disconnected object exists within the graph.

These settings allow you to edit the proportional and constant spacing between disconnected objects of a graph.

- **Constant Spacing.** Specify the minimum amount of white space between adjacent rows and columns.
- **Proportional Spacing.** Specify a fractional value. The margin around disconnected objects is calculated as the product of this fractional value and half the height of the node.

Printing the Component Topology Analysis Graphs


To print the Component Topology Analysis graph using the current print setup, choose the  > **Print Graph** menu item. TransactionVision invokes the standard printing mechanism for your platform.

Note: Setting the page orientation to landscape on the Windows Print dialog has no effect. Instead, use the TransactionVision Page Setup dialog to set page orientation if you wish to print in landscape mode. For more information about this dialog, see “Page Setup” on page 511.

When you print a graph, a Java warning appears asking permission to print. This message is generated by the Java security feature to prevent unauthorized use of printers. The first time you print in a session, it appears multiple times; it appears once after the first time you print. To disable this message, add the following lines in the java.policy file in your home directory on the client computer:

```
grant {  
    permission java.lang.RuntimePermission "queuePrintJob";  
};
```

Print Preview

To display a print preview, choose the  > **Print Preview** menu item. The Print Preview window opens. It is a secondary window that displays a preview of the printed graph. The Print Preview window allows you to preview graph images before you print them so you can see in advance how changes in print setup options affect your printed images.


Note: Do not change the printing orientation options on the **Print Setup > Page Setup** dialog through the Print Preview window. Instead, use the TransactionVision **Page Setup** dialog to change the orientation options. For more information about this dialog, see “Page Setup” on page 511.

In the Print Preview window, you can zoom in on a preview image by moving the mouse over it and clicking.

You may also use the following buttons to control the appearance of the print preview:

Button	Description
Print	Invokes the standard printing mechanism for your platform.
Print Setup	Opens the Print Setup dialog for setting printing options.
Zoom In	Increases the zoom level for the preview image.
Zoom Out	Decreases the zoom level for the previous image.
Zoom Percentage	Shows the current zoom percentage. You may select a pre-defined zoom percentage from the list or enter a custom zoom percentage.
Fit in Window	Scales the preview image to display the complete image in the preview window.
Close	Closes the Print Preview window.

Print Setup

To customize the look of your printed graph, choose the  > **Print Setup** menu item. The Print Setup dialog opens.

This dialog offers options for tailoring the look of your printed Component Topology Analysis or Transaction Analysis view.

Set the following options, then click **OK** to apply your changes. Click **Cancel** to close the dialog without applying your changes, **Default** to revert to default values, or **Page Setup** to display the Page Setup dialog for selecting paper size, orientation, and margins.

Option	Description
Print	Select the part of the graph you want to print. You may print the entire graph, only the part of the graph currently visible in the view window, or only the selected graph component.
Scale By	Set options related to page size. Select Pages and specify the number of page rows and columns to scale the image to fit the specified page size. The image is not scaled if Actual Size is selected. Choose Zoom Level to scale the image to the zoom level specified in the view.
Print Caption	Select to print the page caption in the Caption field.
Caption	Enter text for the page caption. You must select Print Caption to enable caption printing. To select the font for the caption, click Font to display the Choose Font dialog.
Position	Select the position on the page for the caption to be printed.
Margins	Set the distances in inches between the edge of the paper and the edge of the graph. Set margins for the top, bottom, left, and right edges.
Print Page Numbers	Select to print page numbers on each page when printing multiple pages. Page numbers are printed just outside the crop marks.
Print Crop Marks	Select to print crop marks when printing multiple pages. Crop marks are the lines drawn on the margins of pages that connect to other pages.
Print Border	Select to print a border around the graph. Click Color to open the Choose Color dialog for setting the border color.
Print Background	Select to print the background color. To set the background color, choose the Layout > Background Color menu item on the Component Topology or Transaction Analysis view.
Print Grid	Select to print a background grid.

Page Setup

Click **Page Setup** on the Print Setup dialog to open the Page Setup dialog.

Use this screen to set up page options for printing Component Topology or Transaction Analysis graphs. Set the following options, then click **OK** to activate your changes or **Cancel** to close the dialog without activating your changes.

Option	Description
Paper Size	Select a page style from the list.
Paper Source	Select a source from the list.
Margins	Set the distance between the text and the edge of the printed page (in inches). Set margins for the top, bottom, left, and right edges.
Orientation	Select whether to use Portrait or Landscape for setting up the page. Portrait is default (the height of the page is greater than the width). For Landscape, the width of the page is greater than the height.
Sample	Displays a preview of the selected settings.

Caption Font

Click **Font** on the Print Setup dialog to open the Choose Font dialog.

Use this dialog to select the font to use for the caption when printing a Component Topology or Transaction Analysis graph.

Select a font family, text attributes, and font size from the drop-down lists. The sample window shows how text on the graphic looks.

Click **OK** to make your changes active or **Cancel** to close the dialog without activating your changes.






Border Color

Click **Color** on the Print Setup dialog to open the Choose Color dialog. Use this dialog to set the border color for printed Component Topology Analysis and Transaction Analysis graphs.

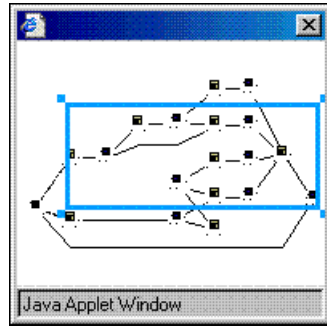
Set the border color in the same way you set the background color for the graph. For more information, see “Set Background Color” on page 501.

Adjusting the Zoom Level

By default, TransactionVision sizes the Component Topology Analysis graph to fit in the view window. There are several ways to adjust the zoom level:

- Choose the **Zoom > Zoom > Zoom In** menu item to increase the zoom level.
- Choose the **Zoom > Zoom > Zoom Out** menu item to decrease the zoom level.
- Click  or choose **Zoom > Selective Zoom** to change to selective zoom mode. The cursor changes to the selective zoom cursor. Click on the view area and then drag the cursor to select an area to zoom. TransactionVision adjusts the zoom level so that the selected area fills the view area. Click  to return to selection mode.
- Click  or choose **Zoom > Interactive Zoom** to change to interactive zoom mode. The cursor changes to the interactive zoom cursor. Click in the view area and drag the cursor to the top or left to decrease the zoom level or drag to the bottom or right to increase the zoom level. Click  to return to selection mode.
- Click  or choose **Zoom > Zoom > Fit Window** to adjust the zoom level so that the entire graph appears in the view window.
- Enter a custom zoom percentage or choose a predefined zoom percentage from the zoom percentage list.

- Choose the **Zoom > Show Overview Window** menu item to open a small secondary window that contains a read-only view of the view window. The entire graph is shown within this window, and a highlighted selection box surrounds the region that is displayed in the active window. The highlighted rectangle allows you to change the view by changing the focal point and zoom level of the active graph window. It is provided as a mechanism to help you see where you are in a graph.



- To resize the selection box, click one of the corner handles and drag your cursor to the desired size. Note that the selection box maintains a constant aspect ratio to match the size of the active view window.
- To move the selection box, click inside the selection box and drag your cursor to the desired location.
- To create a new selection box, click outside of the current selection box and drag your cursor to the desired size.

Note: Any time the Component Topology Analysis view is refreshed, the zoom level reverts to Fit Window.

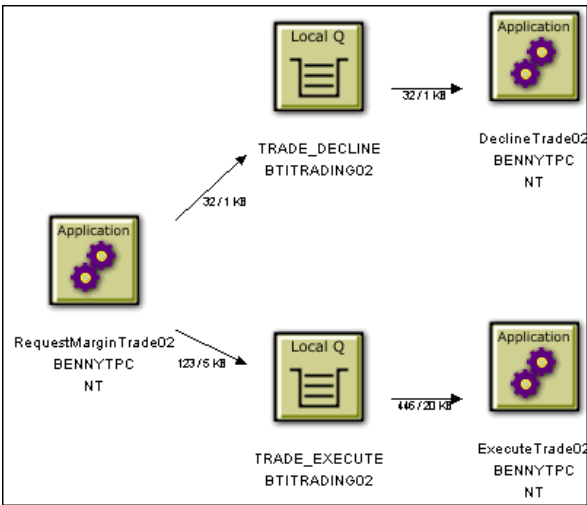
Configure Edge and Node Labels

You may configure edge and node labels to provide more information about the components and events they represent. Select the following options from the Component Topology Analysis **Edge & Node** menu to configure the edge and node labels.

Note: Each menu option shows or hides the label, and the sub-menu options specify the label details.

Message Count/Byte Count

This label indicates the total number of successful message calls between the resources, along with the number of bytes transferred. Note that MQIMS_BRIDGE_ENTRY and MQIMS_BRIDGE_EXIT calls are not included in the message count/byte count.



To specify a unit of measure for the byte count, choose one of the following from the **Edge & Node > Show Byte Counts** menu:

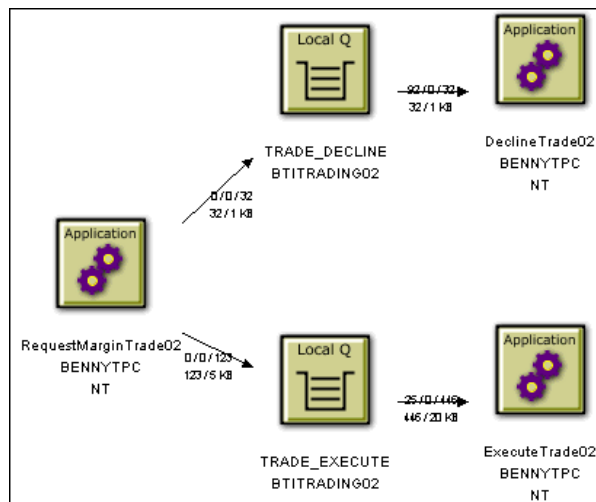
- In MegaBytes
- In KiloBytes
- In Bytes

This menu item is only enabled if the message count/byte count is displayed.

Error, Warning, and Success Count

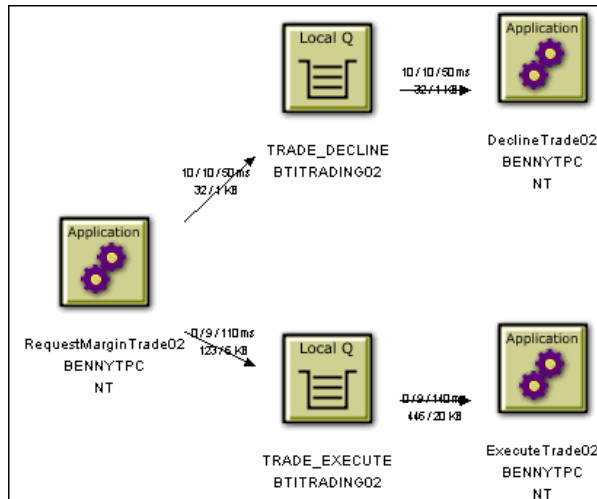
This label indicates the number of error events, warning events, and successful events between components. Events include the following:

- All WebSphere MQ MQGET, MQPUT, MQPUT1, MQINQ, MQSET, MQOPEN, MQIMS_BRIDGE_ENTRY, and MQIMS_BRIDGE_EXIT calls.
- All JMS receive, receiveNoWait, OnMessage, OnException, createReceiver, createSubscriber, createBrowser, publish, send, createSender, createPublisher, and nextElement calls.
- All servlet HTTP_PUT, HTTP_GET, and HTTP_POST requests.



Latency Times

This label indicates the minimum, average, and maximum transaction times for each connection. Use this information to identify potential performance bottlenecks. Transaction times are defined as the time from when a message is put onto the queue and when it was pulled off that queue. For transactions involving a proxy object, TransactionVision shows a latency of N/A.

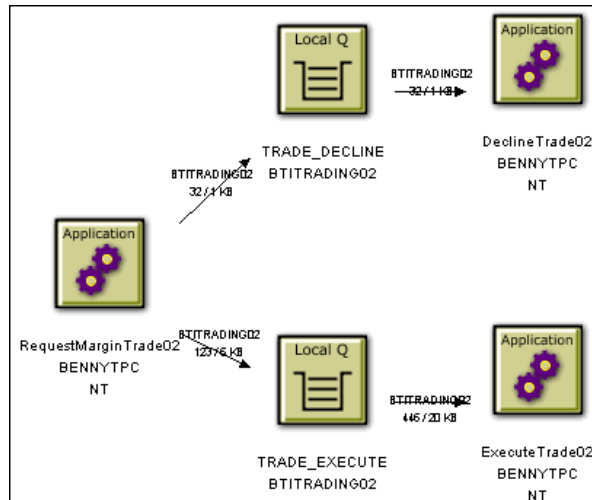


Note: For EJB and servlet events, the time on the edge pointing to a bean represents the time spent in that bean's or servlet's methods; not a latency time.

By default, latency times are calculated using a resolution of 1/100 seconds. To modify the latency resolution, change the value of the `latency_resolution` attribute in the `Analyzer.properties` file. For more information about this file, see the *HP TransactionVision Deployment Guide* PDF.

Edge Label

Edge Labels indicate the connection associated with the edge.



EJB Methods

This label provides details about the EJB Methods. Select the following menu items to display or hide the method information:

- **Show Standard EJB Methods.** Includes standard EJB methods in the edge labels showing the error/warning/success count.
- **Show Access Methods.** Includes access methods in the edge labels showing the error/warning/success count.
- **Show Bean Relationships.** Indicates EJB relationships.

JDBC Methods

This label provides details about the JDBC Methods. Select the following **JDBC Statistics** menu items to display or hide the method information:

- **Cursor Iteration/Get Count.** Indicates the number of result rows that were retrieved using this statement, and the number of calls made to the JDBC resultset to get results from the row.

- **Min/Avg/Max Execute.** Displays statistics showing timing of how long it took to execute this SQL statement.
- **Min/Avg/Max Cursor time.** Displays statistics showing the total time spent iterating through result rows and retrieving data from the JDBC resultset.
- **Avg Time (Cursor, Execute) per statement/Total Statement time.** Displays statistics showing the combined average time and total time spent on operations using this SQL statement.
- **Commit time (Database grouping only).** Displays statistics showing the commit time. Note that this aggregate data is not available on individual statements, so this statistic only displays when viewing your topology with the Database grouping level enabled.
- **Total time (Cursor, Execute, commit)(Database grouping only).** Displays statistics showing the total time spent in JDBC calls, including Cursor time, Execute time and commit time. Note that this aggregate data is not available on individual statements, so this statistic only displays when viewing your topology with the Database grouping level enabled.

Modify Node Labels

If an alias list is associated with the current project, you may choose to display user-defined node name instead of or in addition to system model object names. Select the following **Node Label** menu items to display or hide the node information:

- **Use Alias Name.** Displays the user-defined alias, if defined, for all objects.
- **Use Real Name.** Displays only the system model object name for all objects.
- **Use Both Names.** Displays the user-defined alias, if defined, along with the system model object name for all objects.

See Chapter 10, “Alias Management” for instructions on defining alias names and assigning alias lists to projects.

Modify Component Groupings

Changing component groupings changes the granularity of the information displayed in the view, depending on the level of information you need. For example, viewing each process thread as a separate component provides very fine granularity, but can make it difficult to spot an information flow problem between programs.

Note that proxy objects do not support grouping by thread, process, transaction name, transaction ID, TCB, and PCB. Selecting these grouping criteria have no affect on proxy objects. For more information about proxy objects and the proxy Sensor, see the *HP TransactionVision Deployment Guide* PDF.

Select the following options from the Component Topology Analysis **Grouping** menu to configure the display granularity:

Queues and Queue Managers

Select one of the following **Queue Grouping Criteria** menu items:

- **MQSeries Object/Queue Manager Combination.** Shows each object/queue manager combination as a separate component.
- **Queue Manager.** Shows all objects on a single queue manager as a single component.
- **Show Model Queue Resolution.** Turns model queue resolution on/off. A model queue is a template of a queue definition that is used when creating a dynamic queue. When you call MQOPEN on a model queue, the queue manager creates a temporary or permanent local dynamic queue using the attributes of the model queue. Select this menu item to show model queues in the Component Topology Analysis.

The edges between a model queue node and its model queue definition, remote queue definition, alias queue definition, or cluster queue node representation are displayed as dashed lines because they represent a relationship rather than a message flow.

Distributed Platforms

For UNIX and Windows NT hosts, select one of the following **Distributed Program Grouping Criteria** menu items:

- **Program Name.** Displays each program name on all hosts as a single component.
- **Host.** Displays all programs on a single host as a single component.
- **Program Name/Host Combination.** Shows the default component grouping.

z/OS CICS

For z/OS CICS hosts, select one of the following **z/OS CICS Program Grouping Criteria** menu items:

- **Transaction ID/Region/Host Combination.** Shows the default component grouping.
- **Transaction ID/Host Combination.** Shows the transaction and host combination for all regions as a single component.
- **Region/Host Combination.** Shows all transaction IDs for a single region on a host as a single component.
- **Host.** Shows all transaction IDs for all regions on a single host as a single component.

z/OS Batch

For z/OS Batch hosts, choose one of the following **z/OS Batch Program Grouping Criteria** menu items:

- **Individual Job Step.** Shows each job step as a separate component.
- **Individual Job.** Shows each job as a separate component.
- **Host.** Shows all TCBs on a single host as a single component.
- **Host/Job Name/Step Name Combination.** Shows the job name, job step, and host combination as a single component.
- **Host/Job Name Combination.** Shows all jobs on a host as a single component.

z/OS IMS

For z/OS IMS hosts, select one of the following **z/OS IMS Program Grouping Criteria** menu items:

- **Individual PSB Name.** Shows each PSB as a separate component.
- **IMS Transaction Name.** Shows each transaction as a separate component.
- **IMS Region Identifier.** Shows each region as a separate component.
- **IMS Identifier.** Shows each IMS identifier as a separate component.
- **Job Name.** Shows each job name as a separate component.
- **IMS Type.** Shows each IMS type as a separate component.

z/OS IMS Bridge

For z/OS IMS hosts, select one of the following **z/OS IMS Bridge Grouping Criteria** menu items:

- **Transaction Name.** Shows each transaction name as a separate component.
- **Job Name.** Shows each job name as a separate component.

Servlet/JSP/EJB

For Servlet/JMS/EJB events, select one of the following **Servlet/JSP Grouping Criteria** menu items:

- **Servlet/JSP/EJB.** Shows servlet, JSP, and EJB icons.
- **Web Application.** Shows web application icons.
- **Application Server.** Shows WebSphere Application Server icons.
- **Host.** Shows all programs on a single host as a single component.

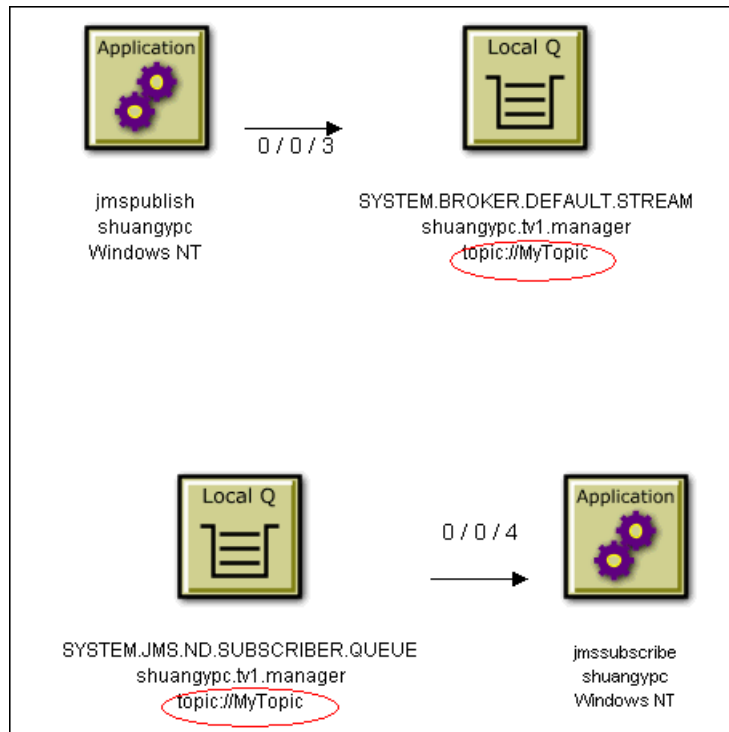
JDBC

For JDBC events, select one of the following **Grouping > Database Grouping Criteria** menu items:

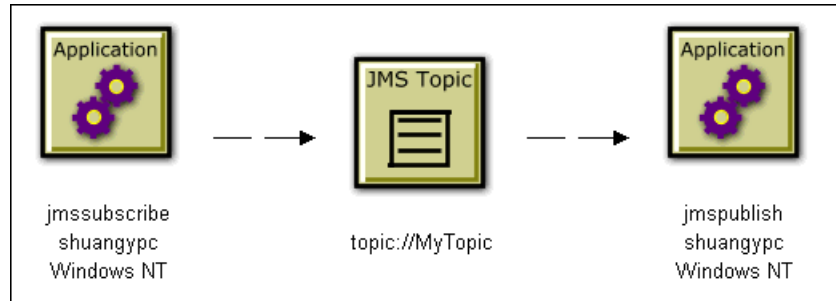
- **Database.** Groups JDBC components by which database they access.
- **Database object.** Groups JDBC components by SQL statements, and shows the tables on which the SQL statements operate.
- **Sql Object.** Groups JDBC components by SQL statements.

Show Publish-Subscribe Topic

By default, TransactionVision is unable to correlate publish-subscribe events unless you monitor the publish-subscribe broker. Therefore, TransactionVision identifies the following events as two separate transactions:



However, if you have enabled publish-subscribe topic correlation, TransactionVision is able to correlate these publish-subscribe topics into a single transaction as follows:



For instructions on enabling TransactionVision to correlate publish-subscribe topics, see the *HP TransactionVision Deployment Guide* PDF. If correlation is enabled, choose **Grouping > Show Publish-Subscribe Topic** to show the correlation.

Group Method Calls On The Same EJB

To group similar method calls on the same EJB as a single edge, choose **Grouping > Group Method Calls On the Same EJB**. Otherwise, an edge is drawn on the graph to represent each method call between EJB beans.

Viewing Statistics

The Component Topology Analysis Static view can display all statistics or statistics from a specific time period.

Viewing All Statistics

To display all events, select the **Show all statistics** option and click **Run Query**.

Viewing Statistics from a Predefined Time Period

To display events from a predefined time period:

- 1 Clear the **Show all statistics** check box, if it is checked.
- 2 Select one of the predefined time periods in the **Set Time to** field.

The date and time fields are automatically populated with the corresponding information. These fields are read-only.

- 3 Click **Run Query**.

Viewing Statistics for a Custom Time Period

To display events from a custom time period:

- 1 Clear the **Show all statistics** check box, if it is checked.
- 2 Select **Custom time** in the **Set Time to** field.
- 3 Enter the desired start and end dates and times in the **Display Static Data from** and **to** fields.
- 4 Click **Run Query**.

The time slice interval set in **StatisticsCache.properties** is the smallest increment of time in which you can view your data in static mode. For example, if your time slice is set to one day, you are not be able to view these statistics on an hourly basis. If you set a specify a custom time period smaller than the time slice, this view shows all events in the time slice associated with the custom time period. For example, if your time slice is set to one day and you specify a custom time period of one hour, the resulting view shows all events for the entire day. If your custom time period spans multiple time slices, the resulting view shows all events from the time slices spanned by the custom time period.

Component Topology User Interface

This section describes:






- Component Topology Analysis Page on page 525

Component Topology Analysis Page

Description	<p>Displays a graphic representation of the interaction between all system components for which Sensors collect event information. The Analyzer correlates events across host, program, and thread boundaries. These correlated events are used to draw a visual map of the flow of information throughout the system.</p> <p>To access: Select Applications > TransactionVision > Transaction Topology > Component Topology Analysis</p>
Important Information	<ul style="list-style-type: none"> ➤ Resource icons represent system components. For information about resource icons, see “Resource Icons” on page 528. ➤ Edges are connecting lines that show the interaction between components. For information about edges, see “Edges” on page 494.

Included in Tasks	“Viewing Statistics” on page 524
Useful Links	<ul style="list-style-type: none">➤ “Viewing the Component Topology Analysis” on page 492➤ “Edges” on page 494➤ “Set Background Color” on page 501➤ “Modify Layouts and Properties” on page 502➤ “Printing the Component Topology Analysis Graphs” on page 508➤ “Adjusting the Zoom Level” on page 512➤ “Configure Edge and Node Labels” on page 514➤ “Modify Component Groupings” on page 519

The following elements are included (unlabeled GUI elements are shown in angle brackets>):






GUI Element	Description
	Click to print the current Component Topology. See “Printing the Component Topology Analysis Graphs” on page 508 for details.
	Click to change to interactive zoom mode. The cursor changes to the interactive zoom cursor. Click in the view area and drag the cursor to the top or left to decrease the zoom level or drag to the bottom or right to increase the zoom level. Click the selection mode icon to return to selection mode.
	Click to change to selective zoom mode. The cursor changes to the selective zoom cursor. Click on the view area and then drag the cursor to select an area to zoom. The zoom level is adjusted so that the selected area fills the view area. Click the selection mode icon to return to selection mode.
	Click to enable section mode.
	Click to adjust the zoom level so that the entire graph appears in the view window.


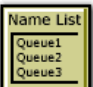


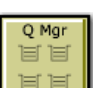
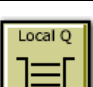
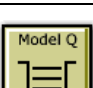
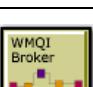
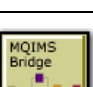
GUI Element	Description
<zoom > drop-down	Select a zoom percentage or enter a custom zoom percentage to adjust the zoom level of the Component Topology Analysis graph.
Display Statistics Data from: <date and time> to: <date and time>	These fields are disabled, but show information that corresponds with the selection in the Set Time to: field.
Edge & Node menu	Menu options to define edge and node labels for the Component Topology Analysis graph. See “Configure Edge and Node Labels” on page 514 for descriptions of each menu item.
Grouping menu	Menu options to define component groupings for the Component Topology Analysis graph. See “Modify Component Groupings” on page 519 for descriptions of each menu item.
Layout menu	Menu options to define the layout of the Component Topology Analysis graph. See “Modify Layouts and Properties” on page 502 for descriptions of each menu item.
Run Query	Click to refresh a view with transactions completed since the view was opened.
Save Settings	Click to save changes and make the current settings the default settings for your user ID.
Set Time to:	<p>Select one of the following predefined time periods:</p> <ul style="list-style-type: none"> ➤ Today ➤ Yesterday ➤ Past Week ➤ This Month ➤ Last 48 hours of data ➤ Custom <p>The date and time fields are automatically populated with the corresponding information.</p>







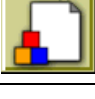



GUI Element	Description
Show all statistics	Check to use all events and click Run Query to view all statistics. Clear if you want to view a pre-defined or custom time period.
Zoom menu	Menu options to adjust the zoom level of the Component Analysis Topology graph. By default, the Component Topology Analysis graph is sized to fit in the view window. See “Adjusting the Zoom Level” on page 512 for descriptions of each menu item.


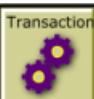
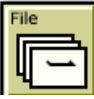





Resource Icons

Resource icons represent system components. Depending on your system, the following resource icons may appear in the Component Topology Analysis:

Icon	Resource Description
	WebSphere MQ cluster node relating multiple queue instances in a clustering environment
	WebSphere MQ distribution list
	WebSphere MQ message queue
	WebSphere MQ remote queue
	WebSphere MQ alias queue

Icon	Resource Description
	WebSphere MQ dead letter queue
	WebSphere MQ namelist
	Program
	Process object
	WebSphere MQ queue manager
	WebSphere MQ local queue
	WebSphere MQ model queue
	WebSphere MQ Integrator broker
	WebSphere MQ IMS Bridge node. If you do not have a custom bean to add queue and queue manager information for IMS Bridge events, an “unknown queue” node is shown for each IMS Bridge node associated with the MQIMS_BRIDGE_ENTRY event. For more information about the TransactionVision IMS Bridge Sensor, see the <i>HP TransactionVision Deployment Guide</i> PDF.

Icon	Resource Description
	Proxy object. Proxy objects involve applications that are not monitored by a TransactionVision Sensor. See the <i>HP TransactionVision Deployment Guide</i> PDF for information about the proxy Sensor and proxy objects.
	Internet
	JMS queue
	JMS topic
	Java server page
	Java servlet
	Web application
	WebSphere web application server
	EJB session bean
	EJB entity bean

Icon	Resource Description
	EJB message-driven bean
	CICS transaction. Note that TransactionVision automatically filters out all CICS transactions that begin with “C” because they are internal CICS transactions. One exception is CKBP, which denotes a WebSphere MQ-CICS bridge transaction. All CICS DPL programs invoked through MQ-CICS bridge have this transaction name.
	CICS file
	CICS Transient Data (TD) queue
	CICS indirect TD queue. Similar to a TD queue, the indirect TD queue is an alias to other TD queues.
	CICS temporary storage (TS) queue. Used for temporarily passing data, users can create and destroy TS queues on the fly.
	A user defined system resource object that participates in one or more user events included by the query. All user defined objects share the same icon with the label “Resource.” For information about implementing user events, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.
	SQL queue. Represents access and execution of an SQL statement.

Troubleshooting and Limitations

For TransactionVision project views that contain an extremely large number of events, the Dynamic Topology view may return no data due to server memory limitations.

19

Transaction Detail Reports

This chapter includes the main concepts, tasks, and reference information for TransactionVision Detail reports.

Concepts

- Transaction Detail Reports Overview on page 534
- Event Analysis Overview on page 535

Tasks

- Run a Report on page 539
- Enable Session Tracking on page 541
- Set Event Analysis Display Options on page 542

Reference

- Event Attributes on page 545
- Transaction Detail Reports User Interface on page 550



Transaction Detail Reports Overview

TransactionVision includes a number of reports that focus on displaying various aspects of transaction instance data. They are located under **Applications > TransactionVision > Transaction Detail**, and include the following:

Report Name	Description	Use This Report To...
Event Analysis	Displays all events in a project that match the current query conditions. For more details, see “Event Analysis Page” on page 551.	see event details, displayed according to any query parameters set.
Service Level Analysis	Displays transaction response times and availability service level analysis. For more details, see “Service Level Analysis Report” on page 559	see transaction response times, and availability service level analysis over time.
Transaction Tracking	Displays the progress of selected transactions, and offers drill-down analysis. For more details, see “Transaction Tracking Report” on page 562.	drill down into individual transactions for analysis.
Web Session	Displays transaction activity for web sessions. For more details, see “Web Session Report” on page 565.	check on transaction details during specific web sessions.

Event Analysis Overview

The Event Analysis report shows all the events in a project that match the current query conditions. By default, this report displays a number of attributes for each event as separate columns. See “Event Attributes” on page 545 for a list of all attributes that can be displayed.

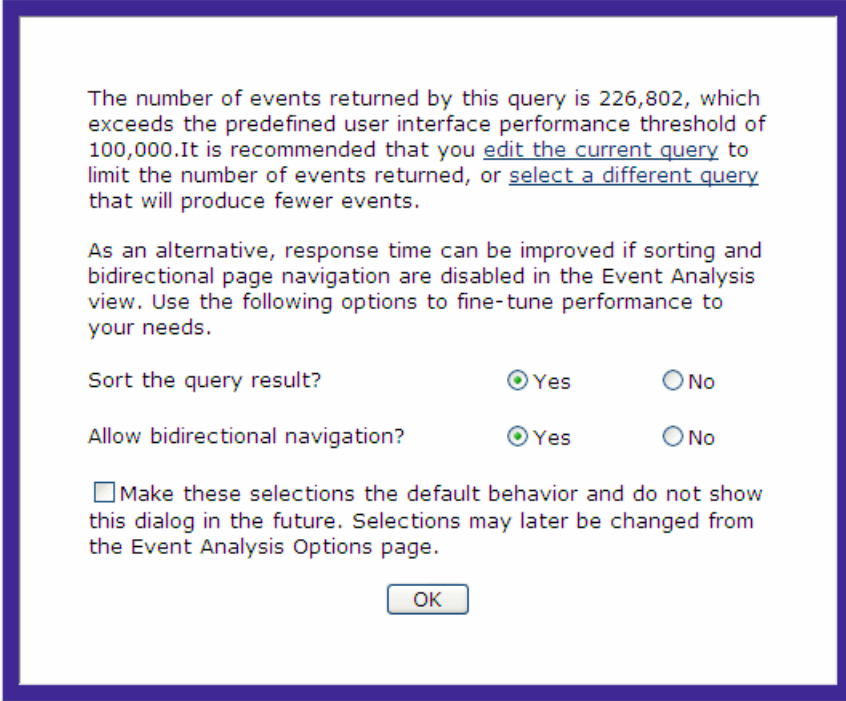
Queries and Event Analysis

Query conditions reduce the number of events displayed by restricting the hosts, queue managers, programs, times, and APIs associated with displayed events. See “Event Analysis Page” on page 551 for descriptions of the query options. For more information about creating or modifying queries, see “Queries” on page 317.

Sorting and Navigating Large Numbers of Events

Sorting large numbers of events in the Event Analysis report can take a long time. Enabling bi-directional navigation (so that you can scroll to previous as well as following pages in the report) for large numbers of events also slows performance. To improve performance in the Event List, you may choose not to sort results and/or enable bi-directional navigation for query results that exceed a certain threshold.

By default, the following dialog is displayed if the number of events that match the current query is greater than 100,000:



The number of events returned by this query is 226,802, which exceeds the predefined user interface performance threshold of 100,000. It is recommended that you [edit the current query](#) to limit the number of events returned, or [select a different query](#) that will produce fewer events.

As an alternative, response time can be improved if sorting and bidirectional page navigation are disabled in the Event Analysis view. Use the following options to fine-tune performance to your needs.

Sort the query result? ☒ Yes ☐ No

Allow bidirectional navigation? ☒ Yes ☐ No

☐ Make these selections the default behavior and do not show this dialog in the future. Selections may later be changed from the Event Analysis Options page.

To change the default number of events for this message to be displayed, change the value on the Event Analysis Options page. See “Set Event Analysis Display Options” on page 542 for instructions.

The recommended action in response to this message is to modify your query to produce fewer events. If you do not wish to modify the query, you may improve response time by disabling sorting of events in this view and/or disabling bi-directional navigation.

Changing Default Settings

Once you specify whether to sort the query result and allow bi-directional scrolling, you may choose to make your selection the new default behavior. If you check the box to make this selection the default behavior, the Sort dialog does not appear in the future.

You may also set sorting and navigation defaults in the Event Analysis Options dialog. See “Set Event Analysis Display Options” on page 542 for instructions.

Event Details

The Event Details page (see “Event Details Page” on page 558) displays the contents of the event data structure fields. If the event is for any of the following calls, the window displays the contents of the user data buffer for the event in a separate pane below the other fields:

- WebSphere MQ MQPUT, MQPUT1, or MQGET calls. Note that for IBM z/OS Batch events, the WebSphere MQ connection handle seen by the user application and the one reported by the Analyzer are always different. The user application sees the SPC pointer, while the Analyzer reports the real connection handle.
- JMS send or receive calls.
- Servlet HTTP_POST or HTTP_GET requests. When a servlet includes other servlets to compose its contents, the user data for each included servlet shows the portion of contents it generated. The user data for the parent servlet shows all the contents generated by all included servlets, plus its own data.

Note: If the data collection filter assigned to the project is configured to collect the API Names Only data range, TransactionVision is unable to display details for events recorded with that filter. See “Data Collection Filters” on page 215 for more information.

See “Event Details Page” on page 558 for a description of this page.

Drill Down to Diagnostics

With Diagnostics integrated into Business Availability Center, new troubleshooting facilities are available. You can drill down from the TransactionVision Event Analysis report to Diagnostics to get the following types of information:

- The drilldown to a Diagnostics Hosts view provides system metrics like CPU utilization and memory utilization that can help you determine if the problem is a system issue.
- The drilldown to a Diagnostics Server Requests view can help you study the performance of these server requests over time and look at detailed instance trees (call profiles, methods, exceptions and faults) associated with the server requests, to identify where the issue lies.

From the Server Request data you can navigate to the corresponding probe entity and look at the application server metrics (for example, JMX metrics) to see if the issue lies with the application server.

To perform the drilldown

- 1 Select any event in the Event Analysis report and click **Details** to display the Event Details.



- 2 Click the **Drill to Diagnostics** button located next to a TransactionVision host.

The Diagnostics UI opens with the Hosts view displayed. The TransactionVision host is selected, data for the host is shown in the graph, and metrics for the host are displayed in the details pane. This is useful for troubleshooting as you see system metrics and can identify if the problem is a system issue.

Some of the metrics available in the details pane for the host system include: Average CPU Utilization, Average Memory Utilization, Average Disk IO in Bytes/Second, Average Network IO in Bytes/Second.

From this Hosts view in Diagnostics you can see the performance of the host system charted over various time periods. Also you can check to see if the system has recently experienced a performance degradation or if the system has consistently under performed.

Run a Report

This task describes how to run a report.

This task includes the following steps:

- “Prerequisites” on page 539
- “Navigate to the Reports” on page 539
- “Select the Report” on page 539
- “Configure Report Settings” on page 540
- “Generate Report” on page 540.

1 Prerequisites

Before you can see transactions in your reports you must define them. See “Transaction Definition Editor” on page 283, for details on how to define and classify transactions.

2 Navigate to the Reports

Go to **Applications > TransactionVision**.

3 Select the Report

The reports are categorized into **Transaction Infrastructure**, **Transaction Summary**, and **Transaction Detail** sections. Check these sections for information about the available reports:

- “Transaction Detail Reports” on page 533
- “Transaction Infrastructure Reports” on page 569
- “Transaction Summary Reports” on page 587.

4 Configure Report Settings

All reports have settings that allow you to select what to display in the report.

Example

The Web Session report offers the following parameters for configuration:

- **Reporting Time Period:** Today, Yesterday, This Week, Last Week, This Month, Last Month, This Year, Last Year, Specified Time Period
- **Results per page.**

The following report shows all transactions within the last week, displaying 20 transactions per page.

Report Parameters	
Reporting Time Period:	Last Week
Results per page:	20
<div>Generate Report Printable View Save Settings</div>	

5 Generate Report

Click **Generate Report**, or **Printable View** if you wish to see the report in a format easy to print, or **Save Settings** if you believe this is a report you may wish to run again. See “Save Report Settings” on page 590 for more details on why and how to save report settings.

Enable Session Tracking

This task describes how to enable analysis and classification of Session ID with a servlet's transaction, to allow classifying and analyzing Session ID data for the Web Session report, see “Web Session Report” on page 565.

To enable session tracking:

- 1 Modify the `<TVISION_HOME>/config/xdm/ transaction.xdm` file to include the field to store the Session ID. The Web Session report uses data from the `/Transaction/Sessionid` field.

Follow this example to modify the file:

```
<Column name="sessionid" type="VARCHAR" size="32"
description="SessionId">
<Path>/Transaction/SessionId</Path>
</Column>
```

Note: After this change, you will need to recreate your project schema on any existing project, and restart the TransactionVision UI/Job Server and analyzer.

- 2 Create an attribute rule in your transaction classification rules that populates the Session ID from your event data. For example:

Attributes	
/Transaction/SessionId	
new value rule	
Value	
XPath	/Event/Technology/Servlet/Session/ID

- 3 Start collecting events.

The transaction class using this rule should now be setting the Session ID. This can be verified by going to the Transaction Tracking report and seeing transactions listed in the Session ID column or by running the Web Session report.

Set Event Analysis Display Options

This task describes how to customize the appearance of the Event Analysis report.

This task includes the following steps:

- “Prerequisites” on page 542
- “Open the Event Analysis Options Dialog” on page 542
- “Set sorting events and navigation defaults” on page 543
- “Set number of user data bytes to display” on page 543
- “Specify date and time” on page 543
- “Specify the time zone” on page 544
- “Specify columns to be shown” on page 544
- “Click Finish” on page 544

1 Prerequisites

- Select a project as described in “Projects” on page 43.
- Set up the transaction events, see “Transaction Definition Editor” on page 283.
- Configure the queries as described in “Queries” on page 317.

2 Open the Event Analysis Options Dialog

Click **Options** on the Event Analysis page. The Event Analysis Options dialog displays. See “Event Analysis Page” on page 551 for details of this page.

3 Set sorting events and navigation defaults

Set the default behavior for sorting events and navigating in the Event Analysis report. These options are useful for improving performance when large numbers of events match the query. You can change the following defaults:

- Number of events that are displayed per page.
- Threshold number at which the warning dialog is displayed. See “Sorting and Navigating Large Numbers of Events” on page 535 for details.
- Sort events chronologically when the number of events exceeds the threshold.
- Allow bi-directional navigation when the number of events exceeds the threshold.
- Show most recent events first in the Event Analysis page.

See “Event Analysis Options Page” on page 555 for descriptions of each option.

4 Set number of user data bytes to display

Enter the number of bytes of user buffer data to display in the Event Details dialog. See “Event Analysis Options Page” on page 555 for details.

5 Specify date and time

Click the **Show the date along with the time** check box to turn on/off the display of the date in addition to the time in the Entry Time and Exit Time columns.

Select whether times in the **Entry Time** and **Exit Time** columns should be displayed in millisecond or microsecond format. Note that only events generated on the z/OS and UNIX platforms are recorded with microsecond precision. See “Event Analysis Options Page” on page 555 for details.

6 Specify the time zone

Select the desired time zone for the **Entry Time** and **Exit Time** columns. The timestamp for all events is recorded in Greenwich Mean Time, but all times are displayed in local time for the selected time zone. See “Event Analysis Options Page” on page 555 for details.

7 Specify columns to be shown

You can add, hide and sort the columns shown in the Event Analysis report. To show all columns associated with one or more technologies, check the check box for the desired technology. The associated columns for each selected technology (shown in the following table), are automatically moved to the **All Selected Columns** list. Clearing a check box automatically moves associated columns to the **All Available Columns** list. For descriptions of the columns, see “Event Attributes” on page 545.

- a** To show a hidden column, select the column name in the **All Available Columns** list and click the right arrow.
- b** To hide a visible column, select the column name in the **All Selected Columns** list and click the left arrow. The **All Selected Columns** list must contain at least one column.
- c** To change the order of the columns in the Event Analysis, select a column name in the **All Selected Columns** list and click the up and down arrows.

8 Click Finish

Click **Finish** to see the resulting Event Analysis display.

Event Attributes

The Event Analysis report shows all the events in the project that match the current query conditions. By default, this report displays a number of attributes for each event as separate columns.

The following table lists all the available attribute columns: by default some columns (for example, those pertaining only to z/OS events) are hidden. See “Set Event Analysis Display Options” on page 542 for instructions on showing and hiding attribute columns.

Column	Description
API Name	The WebSphere MQ or JMS API name, the HTTP request, or EJB methods name.
Application Server	For servlet events, the name of the WebSphere MQ application server associated with the event. For WebSphere Application Server 5.0 Express Edition, the application server name default is server1; other names are an alias for this server name.
CICS API Type	For CICS events, the API type (File Control, Interval Control, Program Control, Program Start, Start Attach, Task End, Task Start, Temporary Storage, and Transient Data).
CICS Appl ID	The name by which CICS is known to the network.
CICS Resource ID	For CICS events, the resources used by CICS transactions. Currently, these resources are file name, transient data (TD) queue name, and temporary storage (TS) queue name.
CICS Response Code	For CICS events, the CICS response code associated with the event.
CICS Start Code	A two character code indicating how the CICS transaction associated with the event was started. Refer to the CICS Application Programming Reference (ASSIGN command) for the range of Startcode values and meanings.
CICS SYSID	The CICS SYSID associated with the event.
CICS Task	The CICS task associated with the event.

Column	Description
CICS Terminal ID	The CICS Terminal ID associated with the event.
CICS Transaction	The CICS transaction ID associated with the event.
Completion Code	The completion code in the WebSphere MQ API call return from the WebSphere MQ library.
Connection Name	For all technologies, the connection name represents the connection the event is using. For WebSphere MQ, the connection name and queue manager are typically the same, except in the case where you open a queue on one queue manager using a connection (specified with the hConn parameter) to a different queue manager over a transmission queue. For example, suppose your program opens queue QM1.Q using a connection to queue manager QM1, then later opens QM1.Q using a connection to queue manager QM2. This event shows QM1.Q as the object name, QM1 as the queue manager, and QM2 as the connection name.
Data Size	For WebSphere MQ events, The size in bytes of the user data buffer for MQGET, MQPUT, and MQPUT1 calls. For JMS events, the size of JMS user data in send and receive calls. For Servlet events, the size of the returning HTML page.
Database	For JDBC events, the database name that these events are connected to.
EJB Name	For EJB events, the EJB name associated with the event. If the Program Name column is displayed, the EJB name also appears as the program name.
Entry Time	The time stamp of when the event initiated.
Event Time	The time used to order the event chronologically. It may be either the value of the event entry time or the event exit time, depending on the API. For some event types such as those that receive messages, the important time is not when the event was first called, but when it returned (when the message was received). For those events, the event time is the event exit time; for all other events, the event time is the event entry time.
Exit Time	The time stamp of when the event finished.

Column	Description
Host Name	The host running the application that called the WebSphere MQ or JMS API or the host the WebSphere application server is running on.
Host IP	The host IP address of the machine the events occurred on.
i5/OS Job Name	The i5/OS job name associated with the event.
JDBC Class	For JDBC events, the JDBC class associated with the event.
JMS Class	For JMS events, the JMS class associated with the event.
JMS Exception Code	For JMS events, the exception code associated with the event.
JMS Queue	For JMS events, the name of the JMS queue associated with the event.
Object Name	The WebSphere MQ object (queue, distribution list, namelist, etc.) associated with the event.
Program Name	The name of the program making the event.
Queue Manager	The WebSphere MQ queue manager associated with the WebSphere MQ event. The queue may not be always be the same as the connection name. For example, it is different when you are putting to a queue on a different queue manager from your connection over a transmission queue.
Reason Code	The reason code in the WebSphere MQ API call return from the WebSphere MQ library.
Reply To Queue	The WebSphere MQ or JMS queue in the Reply To field.
Reply To Queue Manager	The WebSphere MQ or JMS queue manager in the Reply To field.
Servlet	For servlet events, the name of the servlet associated with the event.
SQLCode	For JDBC events, SQL code used by JDBC. The code is based on the underlying database vendor.
SQL Objects	For JDBC events, tables, stored procedures, views, this event accesses.

Column	Description
SQLState	For JDBC events, SQL State values used by JDBC. The values are based on the underlying database vendor.
SQL Type	For JDBC events, type of SQL statement — batched SQL statements such as DELETE, INSERT, MERGE, SELECT, Stored Procedure, UPDATE.
Status Code	For servlet events, the status code associated with the event. Note: The TransactionVision servlet Sensor is unable to get the response status code and any headers in the response object for servlet events generated from an HTTP_HEAD request. Therefore, the status code value for HTTP_HEAD events is always N/A.
Technology	The technology used by the event. Currently, WebSphere MQ, JMS, Servlet, and EJB technologies are supported.
Topic	For JMS events, the name of the topic associated with the event.
Txn ID	For transaction events, the query on the transaction identifier.
Txn Class	For transaction events, the query on the selected transaction class.
Txn Start Time	For transaction events, the query on the transaction start time.
Txn End Time	For transaction events, the query on the transaction end time.
Txn Response Time	For transaction events, the query on the response time, in milliseconds, of the transaction.
Txn State	For transaction events, the events from transactions that are in a Completed/Processing/Unknown state.
Txn Result	For transaction events, the events from transactions that have a Failed/Success/Unknown result.
Txn Label	For transaction events, the events from transactions based on matching the transaction label.

Column	Description
Txn SLA State	For transaction events, the query on the SLA state of Aged out/None/Violated.
User Event Class	For user events, the class associated with the event. For information about implementing user events in your application, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.
User Event Completion Code	For user events, the completion code associated with the event. For information about implementing user events in your application, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.
User Event Status	For user events, the status associated with the event. For information about implementing user events in your application, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.
User Event Object	For user events, the object associated with the event. For information about implementing user events in your application, see the <i>HP TransactionVision Advanced Customization Guide</i> PDF.
Web Application	For servlet events, the name of the web application associated with the event.
z/OS IMS ID	The z/OS IMS identifier associated with the event.
z/OS IMS PSB	The z/OS IMS PSB name associated with the event.
z/OS IMS Region ID	The z/OS IMS region identifier associated with the event.
z/OS IMS Region Type	The z/OS IMS region type associated with the event.
z/OS IMS Transaction	The z/OS IMS transaction name associated with the event.
z/OS Job Name	The z/OS Batch job name associated with the event.
z/OS Job Step	The z/OS Batch step name associated with the event.

Transaction Detail Reports User Interface

This section describes:

- Event Analysis Page on page 551
- Event Analysis Options Page on page 555
- Event Details Page on page 558
- Service Level Analysis Report on page 559
- Transaction Tracking Report on page 562
- Web Session Report on page 565

 **Event Analysis Page**

Description	<p>Displays all the events in the project that meet query conditions. By default, this page displays a attributes for each event in column format.</p> <p>To access: Select Applications > TransactionVision > Transaction Detail > Event Analysis</p>
Important Information	<ul style="list-style-type: none"> ➤ Color-Coded Completion Codes: Events with a warning or error completion code are displayed in a different color from other events, making them easy to spot. ➤ Sort Order: If the events are sorted, an arrow to the left of the first attribute column heading indicates the sort order. A down arrow indicates that events are sorted in chronological order. An up arrow indicates reverse chronological order. If the events are not sorted, no arrow is displayed. For information about sorting, see “Sorting and Navigating Large Numbers of Events” on page 535. ➤ Navigation: Events are displayed in pages according to the number of events per page specified in the Event List Options dialog. See “Set Event Analysis Display Options” on page 542 for instructions on using this dialog. See “Event Analysis Options Page” on page 555 for descriptions of this dialog.
Included in Tasks	“Set Event Analysis Display Options” on page 542
Useful Links	<ul style="list-style-type: none"> ➤ “Event Attributes” on page 545 ➤ “Save Report Settings” on page 590 ➤ “Sorting and Navigating Large Numbers of Events” on page 535

Event Analysis Contents

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Clear Checks	Click to deselect all selected events (useful if selected events are not near each other in the Event Analysis page).
Compare	<p>Select two events and click the Compare button to see two events compared side by side.</p> <p>A dialog displays the contents of the event data structure fields for two events side-by-side. If the events are both for the same call type (for example, both for MQGET events or both JMS send events), any fields with different values are highlighted in yellow. Any two EJB events are comparable also. The difference is highlit and the event detail tree nodes are sorted under the same level.</p> <p>Scroll lock is automatically set for the left and right panes. Scrolling in one window automatically scrolls the other window, so you can compare fields easily.</p>
Detail	Click the Detail button to show event details.
Options	Click the Options button to customize the appearance of the Event Analysis page. See “Event Analysis Options Page” on page 555 for details.
Run Query	Click to refresh a view with events collected since the view was opened.

GUI Element	Description
<navigation links>	<p>If your Event Analysis report contains more than one page of events, use the following navigation links to navigate between pages:</p> <ul style="list-style-type: none"> ➤ Page Number: Enter the page number of the page you wish to display. The numbers to the left of this field indicate the current page number and the total number of pages. For example, [130/731] indicates that the current page is page 130 of 731 total pages ➤ First Page: Displays the first page of events. This navigation link only appears after page 10 or greater has been displayed. ➤ Prev 10: Goes back 10 pages. This navigation link only appears after page 10 or greater has been displayed. ➤ Prev: Displays the previous page of events. This navigation link does not appear on page 1. ➤ Next: Displays the next page of events. ➤ Next 10: Goes forward 10 pages. ➤ Last Page: Displays the last page of events.
<check box>	<p>Click to select an event to perform operations such as viewing only events in the same business or local transaction or displaying event details.</p> <p>You may select two events. If you select a third event, the first selected event is automatically deselected so that only the two most recent events clicked are selected.</p> <p>Click a selected check box to deselect it.</p>
<number of events>	Displays the total number of events collected and which events are being shown.
Open these events in	Enables you to show selected events in the Component Topology Analysis page as well as the Event Analysis page.

GUI Element	Description
Query	Select one of the following from the drop-down list: <ul style="list-style-type: none">➤ All Events➤ Last 24 Hours➤ New Query➤ Edit Query See “Queries” on page 317 for details.
View events as:	By default, the Event Analysis report displays all events in the project that match the current query conditions. Select an event, then one of the following: <ul style="list-style-type: none">➤ Business Transaction — to display only events in the same business transaction as the selected event.➤ Local Transaction — to display only events in the same local transaction as the selected event.➤ No Transaction — to display all events after viewing only events in a local or business transaction.

Event Analysis Options Page

Description	Enables you to set the default behavior for sorting events and navigating in the Event Analysis report. These options are useful for improving performance when large numbers of events match the query. To access: Applications > TransactionVision > Transaction Detail > Event Analysis and click Options .
Included in Tasks	“Set Event Analysis Display Options” on page 542
Useful Links	“Sorting and Navigating Large Numbers of Events” on page 535.

Event Analysis Options

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Number of events to display per page	Set the number of events to display per page
Minimum number of events for a warning message box to be displayed	Enter a new value to change the threshold number at which the dialog is displayed.
Display a warning message box when number of events exceeds the minimum count	Check to display the dialog shown in “Sorting and Navigating Large Numbers of Events” on page 535 each time the number of events in this view exceeds the threshold.
Sort events by event time when the number of events exceeds the minimum count	Check to sort events chronologically when the number of events exceeds the threshold.

GUI Element	Description
Allow bi-directional navigation of the event list when the number of events exceeds the minimum count	Check to allow bi-directional navigation when the number of events exceeds the threshold.
Show event list in reverse chronological order	Check to show most recent events first in the Event Analysis report.

Detail and Time Options

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Detail Options	Enter the number of bytes of user buffer data to display in the Event Details dialog.
Show the date along with the time	In the Time Column Format section, check to turn on/off the display of the date in addition to the time in the Entry Time and Exit Time columns.
Show times out to the	In the Time Column Format section, select whether times in the Entry Time and Exit Time columns should be displayed in millisecond or microsecond format. Note that only events generated on the z/OS and UNIX platforms are recorded with microsecond precision.
Time Zone Options	Select the desired time zone for the Entry Time and Exit Time columns. The timestamp for all events is recorded in Greenwich Mean Time, but all times are displayed in local time for the selected time zone.
Detail Options	Enter the number of bytes of user buffer data to display in the Event Details dialog.

Event Analysis Columns

Technology	All Selected Columns
Common	Event Time, Entry Time, Exit Time, API Name, Host Name, Host IP, Program Name, Technology, Data Size
EJB	EJB Name
JDBC	SQL Type, JDBC Class, SQL Objects, Database, SQLCode, SQLState
JMS	JMS Class, JMS Queue, Topic, JMS Exception Code
CICS	CICS Task, CICS SYSID, CICS Transaction, CICS API Type, CICS Terminal ID, CICS Appl ID, CICS Start Code, CICS Resource ID, CICS Response Code
Servlet	Status Code, Servlet, Web Application, Application Server
User Event	User Event Class, User Event Completion Code, User Event Status, User Event Object
WebSphere MQ	Completion Code, Reason Code, Connection Name, Queue Manager, Object Name, Reply To Queue Manager, Reply To Queue
WebSphere MQ z/OS Batch	z/OS Job Name, z/OS Job Step
WebSphere MQ z/OS IMS	z/OS IMS Region ID, z/OS IMS Region Type, z/OS IMS Id, z/OS IMS Transaction, z/OS IMS PSB
WebSphere MQ i5/OS	i5/OS Job Name
Transactions	Txn ID, Txn Class, Txn Start Time, Txn End time, Txn Response time, Txn State, Txn Result, Txn Label, Txn SLA State



Event Details Page

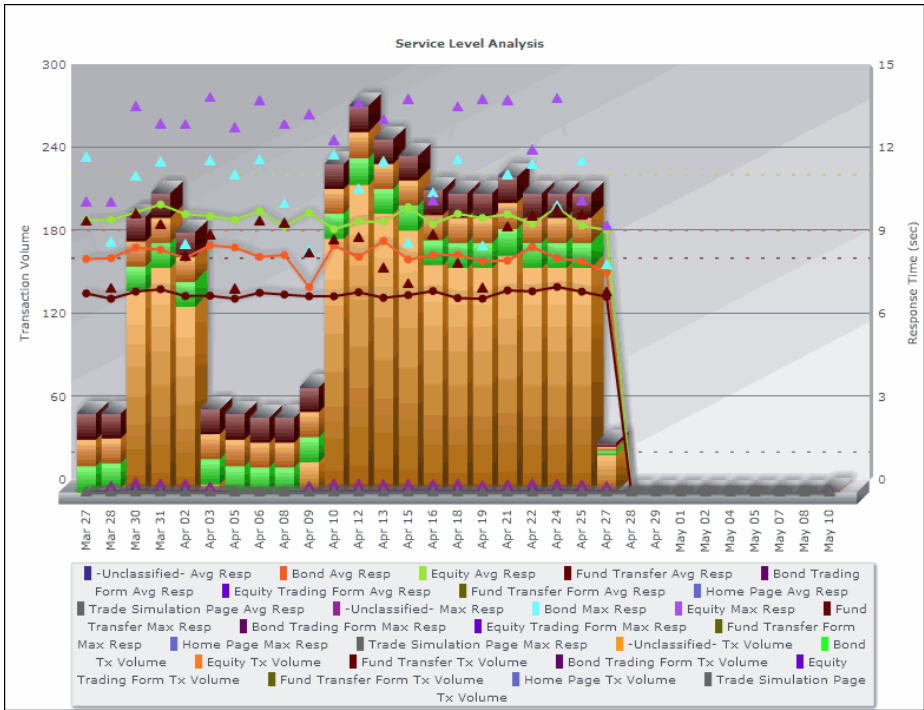
Description	<p>Displays additional information about any event in the Event Analysis report. The Event Details dialog interprets and displays all of the information in the MQMD header information, the dead letter queue header, and the user data buffer associated with the event.</p> <p>The event details are displayed in a hierarchy, enabling you to view different levels of detail.</p> <p>To access: Applications > TransactionVision > Transaction Detail > Event Analysis and click Details.</p>
Useful Links	"Event Details" on page 537

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
-	Click to collapse the details below a single node. A + indicates a collapsed field. Click to expand.
+	Click to expand the details below a single node.
Collapse All	Click to collapse all nodes.
Expand All	Click to expand all nodes.
Save as XML	Click to save a copy of the current event as an XML document. This option is available for Internet Explorer and Mozilla Firefox.

Service Level Analysis Report

The following is an example of the Service Level Analysis report.



Description	Displays transaction response times and availability service level analysis. Assists monitoring of Service Level Agreement response requirements. To access: Select Applications > TransactionVision > Transaction Details > Service Level Analysis Report .
Important Information	Linking to the Transaction Tracking report: <ul style="list-style-type: none">➤ From the Service Level Analysis report, click a section of a vertical bar to run the Transaction Tracking report on all the transactions the bar represents.➤ Click a triangle in the Service Level Analysis report to show the Transaction Details view of the Transaction Tracking report, for the specific transaction represented.
Useful Links	➤ For more detail on individual transactions, see “Transaction Tracking Report” on page 562.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Reporting Time Period	Specify the time period for a project, or All Time Periods .
Include Transaction Classes	Select from the list of transaction classes, or click Select All , or Unselect All .
Series Selection	Select Avg Response Time , Max Response Time and/or Volume .
Combine Transaction Class Totals	Check if you wish to combine the transaction class totals.
Output Results as	Select whether you wish your report displayed as a Graph or a Table .

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Tool-tip>	Hover the cursor over a data item and a tool-tip displays its legend and value.
<Transaction Classes>	Each selected transaction class appears as a color in the report, and is notated in the legend at the base of the report.
Average Response Time / Maximum Response Time	For each transaction class, triangles represent the transaction with the maximum response time, and circles represent the average response time.
Transaction Volume	For each transaction class, the vertical bars represent transaction volume.

 **Transaction Tracking Report**

The following is an example of the Transaction Tracking report.

Transaction Tracking Report (6249 transactions, showing 1-20)						
Tx Class	Start Time	Response Time	Completion State	Result State	SLA State	
-Unclassified-	06/07/2005 10:47:46.851	0.02	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:46.893	0.02	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:46.923	0.02	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:47.019	0.04	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:47.156	0.04	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:47.283	0.04	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:47.397	0.01	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:47.427	0.03	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:48.756	0.04	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:48.804	0.01	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:48.913	0.04	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:50.832	0.05	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:50.963	0.14	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:51.164	0.11	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:51.358	0.08	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:51.504	0.06	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:51.633	0.05	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:51.765	0.05	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:51.889	0.05	Unknown	Unknown	None	
-Unclassified-	06/07/2005 10:47:52.011	0.05	Unknown	Unknown	None	
First Page Prev 10 Prev Go to page: <div>1</div> Next Next 10 Last Page						

Description	Displays transactions and allows drill-down into individual transactions to analyze behavior. To access: Select Applications > TransactionVision > Transaction Detail > Transaction Tracking Report .
Important Information	To show details on a particular transaction from this report, click the transaction class (Tx Class) link. A Transaction Details page displays, with options to display the transaction in the Event or Component Topology Analysis views.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Include Transaction Classes	Select the Transaction Classes to be included in the report.
Only show transactions exceeding service level	Select if you wish to filter for transactions that have exceeded the desired service level.
Reporting Time Period	Select a time period for the report from the drop-down list, or All Time Periods.
Results per page	Specify the number of transactions to show per page of results. If the number of transactions that match the setting exceeds this number, multiple pages of results are generated.
Search for	Enter text for a specific search.
Search on Attribute	Select an attribute to search on, if desired.
Show results in	Select to sort the results in Ascending or Descending order.
Sort by Attribute	<p>Select a custom attribute to sort results by, if desired. Custom attribute columns are available if you have added rules to a transaction definition. Use the Transaction Definition editor to add rules (see “Transaction Definition Editor” on page 283), or manually edit the transaction definition file (see the <i>HP TransactionVision Advanced Customization Guide</i> PDF for details).</p> <p>Note: If the Transacting Tracking report is set to sort on the Class ID attribute, results are sorted in the order in which the Transaction Classifications were defined in the Transaction Definition editor, and not in alphabetical order by Transaction Class Name.</p>

GUI Element	Description
Transaction Result Code	Filter on one or more of the following: Unspecified, Success, Failed.
Transaction State	Filter on Unspecified, In-Process and/or Completed transaction states.

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Completion State	Displays the completion state of the transaction. Only transactions where the completion state matches that selected in the report settings display in the report.
Label	By default transaction labels are null. See the <i>HP TransactionVision Advanced Customization Guide</i> PDF for information on modifying transaction attributes.
Response Time	Displays the response time of the transaction. If the report settings were configured to show only those transactions that exceed service level agreements then only those display.
Result State	Displays the result code of the transaction. Only transactions where the result state matches that selected in the report settings display in the report.
Start Time	Displays the start time of the transaction.
Tx Class	The transaction class of the transaction. By default transactions are unclassified, see “Transaction Definition Editor” on page 283 for how to set up transaction classification classes.
End User Name	The user name of the transaction initiator.
End User Start IP	The IP address used for the initiation of the transaction.
End User End IP	The IP address used for the completion of the transaction.

GUI Element	Description
Country, State, City	These 3 columns indicate the location where the transaction took place.
SLA State	None, or Violated.

Web Session Report

The following is an example of the Web Session report.

Web Session (40 sessions, sho	
Session Start	Session End
11/21/2008 12:16:13.962	11/21/2008 12:20:04.464
11/21/2008 12:31:14.193	11/21/2008 12:32:18.155
11/21/2008 12:46:14.727	11/21/2008 12:50:04.476
11/21/2008 13:01:14.155	11/21/2008 13:05:04.083
11/21/2008 13:16:13.983	11/21/2008 13:20:04.675
11/21/2008 13:31:14.614	11/21/2008 13:35:05.157
11/21/2008 13:46:14.743	11/21/2008 13:50:07.657
11/21/2008 14:01:14.871	11/21/2008 14:05:06.502
11/21/2008 14:16:14.698	11/21/2008 14:20:05.838
11/21/2008 14:22:15.996	11/21/2008 14:26:07.255
11/21/2008 14:28:14.981	11/21/2008 14:32:05.827
11/21/2008 14:34:15.073	11/21/2008 14:38:03.654
11/21/2008 14:40:15.165	11/21/2008 14:44:04.117
11/21/2008 14:46:15.256	11/21/2008 14:50:04.018
11/21/2008 14:52:15.247	11/21/2008 14:56:04.293
11/21/2008 14:58:15.339	11/21/2008 15:02:03.381
11/21/2008 15:04:15.129	11/21/2008 15:05:17.985
11/21/2008 15:05:58.530	11/21/2008 15:08:05.338
11/21/2008 15:10:20.451	11/21/2008 15:14:08.582
11/21/2008 15:16:20.543	11/21/2008 15:20:09.092

Description	Displays transaction activity for a web session. To access: Select Applications > TransactionVision > Transaction Detail > Web Session Report
Important Information	In order to generate this report your Analyzer must be configured to extract and store Session ID information, as this is not stored by default, see “Enable Session Tracking” on page 541.
Useful Links	“Enable Session Tracking” on page 541

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Reporting Time Period	Offers predefined time periods, or click the calendar icon and select a date, choose Date Range and enter, or select All Time Periods and include all statistics for the project.
Results per page	Enter the number of results to list on each page.

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Session Start	Displays the session start time. Click on this link to display the Transaction Tracking report on the transactions within this session. See “Transaction Tracking Report” on page 562 for details.
Session End	Displays the session end time.
Transaction Count	Displays the total number of transactions within this session.
Total (Dollars)	Displays the dollar value of all transaction within this session.

20

Transaction Infrastructure Reports

This chapter includes the main concepts, tasks, and reference information for TransactionVision Infrastructure reports.

Concepts

- Transaction Infrastructure Reports Overview on page 570

Reference

- Transaction Infrastructure Reports User Interface on page 571

Transaction Infrastructure Reports Overview

TransactionVision includes a number of reports that show information about the different infrastructure elements that transactions execute across. They are located under **Applications > TransactionVision > Transaction Infrastructure**, and include the following:

Report Name	Description	Use This Report To...
Application Server Statistics	Displays the performance of application server components. For more details, see “Application Server Statistics Report” on page 572.	view average latency and call count statistics on the activity of any EJB and Servlet components.
Message Latency Analysis	Shows how message latency is performing on a given queue. For more details, see “Message Latency Analysis Report” on page 575.	see a chart of average and maximum latency on any message queue.
Performance Dashboard	Displays IT infrastructure performance, updated regularly according to the parameters set. For more details, see “Performance Dashboard Report” on page 577.	see a real-time view of summary information and performance statistics for Transactions, Application Servers, WMQ, and JMS activity.
SQL Statistics	Displays overall statistics on all executed SQL statements. For more details, see “SQL Statistics Report” on page 580.	see details of JDBC activity.
WebSphere MQ and JMS Statistics	Displays statistics on the activity of WebSphere MQ and JMS components. For more details, see “WebSphere MQ and JMS Statistics Report” on page 583.	see Queue and Queue Manager activity including MQGET/MQPUT and JMS event status.

Transaction Infrastructure Reports User Interface

This section describes:

- Application Server Statistics Report on page 572
- Message Latency Analysis Report on page 575
- Performance Dashboard Report on page 577
- SQL Statistics Report on page 580
- WebSphere MQ and JMS Statistics Report on page 583

 **Application Server Statistics Report**

The following is an example of the Application Server Statistics report.

EJB Statistics		
EJB	EJB Method	Call count [sucess/error] Late
BondTradeProcessor		
	onMessage	1796/0
EquityTradeProcessor		
	onMessage	1778/0
FundTransferProcessor		
	onMessage	919/0
Order		
	ejbCreate	2704/0
	ejbLoad	2704/0
	ejbPostCreate	2704/0
	ejbStore	5408/0
	getData	5408/0
	getId	10815/0
	getType	5408/0
	process	2704/0
	setData	2704/0
	setId	2704/0
	setType	2704/0
QualifyTrade		
		42/0
	EquityTradeProcessor	21/0
	Order	42/0
	getData	42/0
	qualify	2788/186
	queue://TRADING/REPLY_FOR_SESSION	42/0
	queue://TRADING/REQUEST	21/0
	setData	21/0
	setType	21/0
	waynecpc	21/0
TradeSession		
	ejbCreate	2890/0
	processTrade	2889/0
Servlet Statistics		
Servlet	Call count [sucess/error]	Late
TradeSession	5355/0	

Description	<p>Displays statistics on the activities of your EJB and Servlet components. The report breaks down activity by EJB and EJB method, and by Servlet, showing the average latency and call count for each.</p> <p>To access: Select Applications > TransactionVision > Transaction Infrastructure > Application Server Statistics Report</p>
Important Information	<p>EJB and Servlet filter: By specifying a ';' delimited list of EJB's and/or servlets you can limit the resulting output. This filter also accepts names that include wildcards such as a prefix or suffix to the object name - for example '*EJB' or 'EJB*'. </p>

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
EJB Filter	Restrict the report by setting the filter. By default all EJBs in this project are selected.
Project	Select the project you wish to report on.
Reporting Time Period	Select the reporting time period from the drop-down list of options. To specify a custom time period, choose Date Range and specify the start and end dates and times. Click the calendar icon to select a date from a calendar control. To include all statistics for the entire project, choose All Time Periods .
Saved Settings	If this report was run before and the settings saved, it appears in this drop-down. If you are the user who saved these report settings, a trash can icon appears, offering the possibility to delete it.
Servlet Filter	Restrict the report by setting the filter. By default all servlets in this project are selected.

EJB Statistics

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Average Latency (sec)	Average latency statistics, in seconds, for the EJB.
Call Count (success/error)	Click the call count link to view a chart and table showing additional information about an EJB method (displayed in a separate window).
EJB	EJB name.
EJB Method	EJB method.

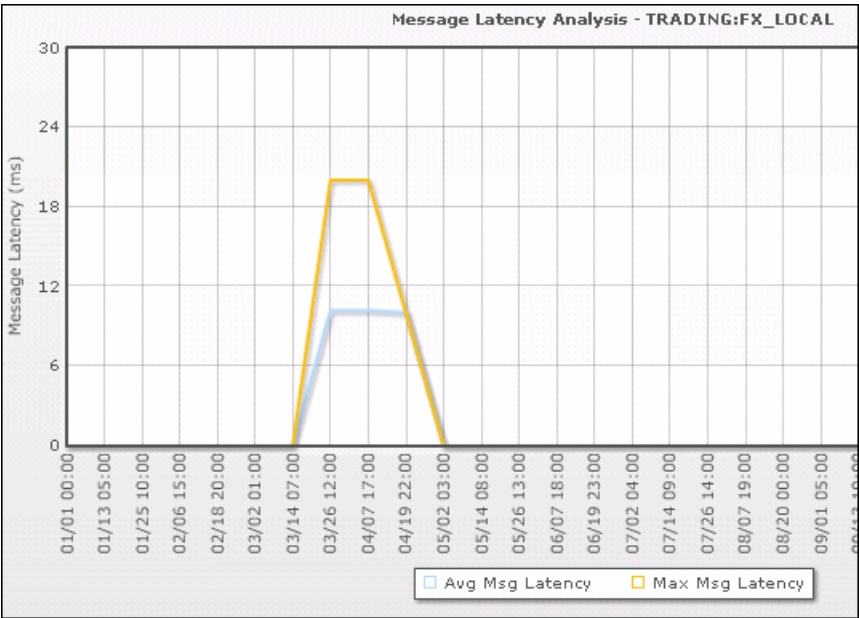
Servlet Statistics

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Average Latency (sec)	Average latency statistics, in seconds, for the servlet.
Call Count (success/error)	Click the call count link to view a chart and table showing additional information about a servlet (displayed in a separate window).
Servlet	Servlet name.

 **Message Latency Analysis Report**

The following is an example of the Message Latency Analysis report.



Description	<p>Displays a chart showing the average and maximum latency on a selected message queue.</p> <p>To access: Select Applications > TransactionVision > Transaction Infrastructure > Message Latency Analysis</p>
Important Information	<p>To view further details on a data item in a chart, move the cursor over the item and a tool-tip displays the values.</p>

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

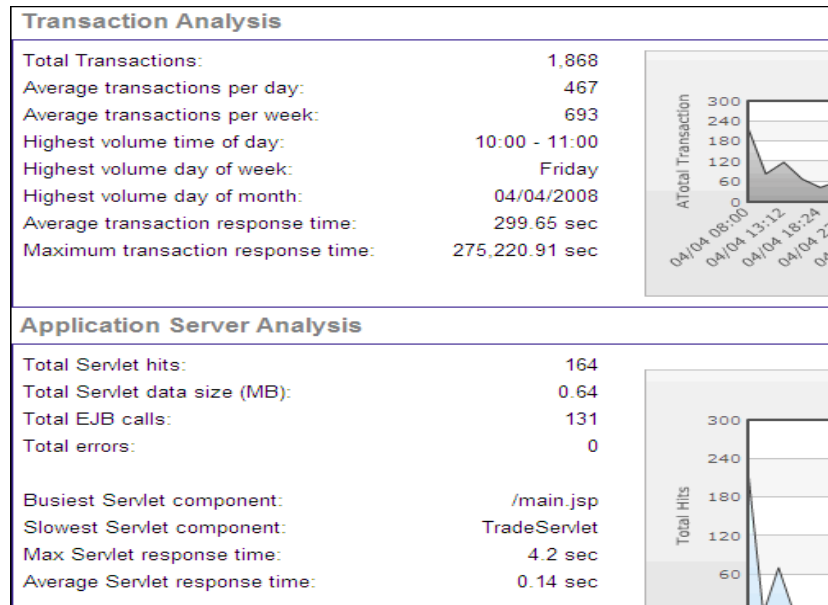
GUI Element	Description
Latency Options	Specify whether to report the Average Latency , Max Latency , or both.
Output results as	Specify whether to show results in a graph, a table, or both.
Project	Select the project you wish to report on.
Queue	Select WebSphere MQ Queue, or JMS Queue, and then the specific projects from the drop-downs.
Reporting Time Period	Select the reporting time period from the drop-down list of options. To specify a custom time period, choose Date Range and specify the start and end dates and times. Click the calendar icon to select a date from a calendar control. To include all statistics for the entire project, choose All Time Periods .
Saved Settings	If this report was run before and the settings saved, it appears in this drop-down. If you are the user who saved these report settings, a trash can icon appears, offering the possibility to delete it.

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
<Parameters>	The Average and/or Max Latency, depending on the parameters set.
Message Latency (ms)	The measurement of message latency.
Message Queue	Full name of the Message queue.
Reporting Time Period	The time being reported on.

Performance Dashboard Report



Description	Displays summary information and performance statistics for Transactions, Application Servers, WMQ, JMS, and JDBC activity. To access: Select Applications > TransactionVision > Transaction Infrastructure > Performance Dashboard
Important Information	For each category, the report provides summary information and applicable performance statistics, including a graph that plots the activity/volume over the specified time interval.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
Project	Select the project you wish to report on.
Reporting Time Period	Select the reporting time period from the drop-down list of options. To specify a custom time period, choose Date Range and specify the start and end dates and times. Click the calendar icon to select a date from a calendar control. To include all statistics for the entire project, choose All Time Periods .
Saved Settings	If this report was run before and the settings saved, it appears in this drop-down. If you are the user who saved these report settings, a trash can icon appears, offering the possibility to delete it.

Report Content

Data is displayed in the following categories:

GUI Element	Description
Transaction Analysis	<ul style="list-style-type: none"> ➤ Total Transactions ➤ Average transactions per day ➤ Average transactions per week ➤ Highest volume time of day ➤ Highest volume day of week ➤ Highest volume day of month ➤ Average transaction response time ➤ Maximum transaction response time
Application Server Analysis	<ul style="list-style-type: none"> ➤ Total Servlet hits ➤ Total Servlet data size (MB) ➤ Total EJB calls ➤ Total errors ➤ Busiest Servlet component ➤ Slowest Servlet component ➤ Max Servlet response time ➤ Average Servlet response time ➤ Busiest EJB Method ➤ Slowest EJB Method ➤ Max EJB Response time: ➤ Average EJB Response time
WebSphere MQ Analysis	<ul style="list-style-type: none"> ➤ Total messages sent ➤ Total messages received ➤ Total megabytes sent ➤ Total megabytes received ➤ Total errors ➤ Highest traffic queue ➤ Highest traffic queue manager ➤ Highest traffic application

GUI Element	Description
JMS Analysis	<ul style="list-style-type: none">➤ Total messages sent➤ Total messages received➤ Total megabytes sent➤ Total megabytes received➤ Total errors➤ Highest Traffic queue/topic➤ Highest traffic application
JDBC Analysis	<ul style="list-style-type: none">➤ Total SQL executed➤ Total Errors➤ Max Insert Execution time➤ Max Select Execution time➤ Max Cursor time➤ Max Batch Execution time➤ Max Other Execution time

 **SQL Statistics Report**

The following is an example of the SQL Statistics report.

JDBC Statistics (18 statements, showing 1-18)										
Program	SQL statement	Database	Total Time	Execute Time (ms)			Cursor Time (ms)			Calls [errors]
			(ms)	Min	Avg	Max	Min	Avg	Max	
JDBCTimingTest	SELECT statement using TESTTIMING	TVISION	311	0	3	30	15	28.1	47	1010 17000 10 [0]
JDBCStmt	SELECT statement using USER_PREF	TESTDB	118	40	40	40	78	78	78	2 3 1 [0]
JDBCStatementTest	INSERT statement using TESTTABLE	TVISION	80	0	6.15	70	-	-	-	- - 5 [8]
JDBCStatementTest	SELECT statement using TESTTABLE	TVISION	40	40	40	40	0	0	0	0 0 1 [0]
JDBCStmt	SELECT statement using USER_PREFS	TESTDB	10	10	10	10	-	-	-	- - 0 [1]
JDBCResultTest	SELECT statement using TESTSG	TVISION	10	10	10	10	-	-	-	- - 1 [0]
	INSERT									

Description	Displays overall statistics on all executed SQL statements. To access: Select Applications > TransactionVision > Transaction Infrastructure > SQL Statistics Report
Important Information	By default, all programs and tables used in the project are included in this report. To restrict the report, check the program filter and/or table manager filter check box and enter what to include. Separate multiple programs or tables with commas. You may include the * wildcard character as the prefix or suffix to an object name.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Include Statement Types	Select the statement types to be included in the report: Batch, Call, Delete, Insert, Merge, Select, Update.
Program Filter	By default all programs are included, to restrict the report use this filter, see Important Information.
Reporting Time Period	Offers predefined time periods, or click the calendar icon and select a date, choose Date Range and enter, or select All Time Periods and include all statistics for the project.
Results per page	Enter the number of results to list on each page.
Show results in	Select an order to display results.
Sort by	Select the column on which to sort results.
Table Filter	By default all tables are included, to restrict the report use the filter, see Important Information.

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Calls (errors)	Number of times the statement was called. The number of calls with errors is in brackets.
Cursor Iterations	The number of result rows that were retrieved using this SQL statement.
Cursor Time	How much time was spent interacting with the JDBC result set (iterating over the cursor, retrieving database values).
Database	The database used.
Execute Time	Min/Avg/Max timing information for how long it took to execute the SQL statement.
Get count	The number of data elements retrieved from the result rows of this statement.
Program	The program for which the data is displayed.
SQL Statement	Clicking on the SQL Statement in the report expands out the full SQL Statement.
Total Time	Total time spent executing the SQL, and interacting with the cursor (if applicable).

WebSphere MQ and JMS Statistics Report

The following is an example of the WebSphere MQ and JMS Statistics report.

WebSphere MQ Statistics				
Queue Manager	Queue	Applications Interacting	Average Latency (sec)	Puts [success/warning/error]
unknown		2	-	0 [0/0/0]
	TRADING	2	-	0 [0/0/0]
TRADING		2	23.23	354 [354/0/0]
	FX	1	-	0 [0/0/0]
	SYSTEM.BROKER.DEFAULT.STREAM	2	0.12	0 [0/0/0]
	SYSTEM.BROKER.ADMIN.STREAM	1	-	0 [0/0/0]
	FX_LOCAL	1	0.15	0 [0/0/0]
	EQUITY	1	0.01	0 [0/0/0]
	REPLY_FOR_SESSION	1	-	0 [0/0/0]
	BOND	1	0.09	0 [0/0/0]
	SYSTEM.JMS.REPORT.QUEUE	2	0	<u>116</u> [116/0/0]
	SYSTEM.JMS.PS.STATUS.QUEUE	1	26.85	<u>222</u> [222/0/0]
	SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE	2	0.03	<u>16</u> [16/0/0]

JMS Statistics				
Queue Manager	Queue	Applications Interacting	Average Latency (sec)	Puts [success/warning/error]
TRADING		8	0.06	32 [32/0]
	FX	1	0.12	<u>3</u> [3/0]
	SYSTEM.BROKER.DEFAULT.STREAM	4	0.08	<u>16</u> [16/0]
	FX_LOCAL	1	0.09	0 [0/0]
	EQUITY	2	0.01	<u>2</u> [2/0]
	REPLY_FOR_SESSION	2	0.02	<u>8</u> [8/0]
	BOND	2	0.04	<u>3</u> [3/0]
	SYSTEM.JMS.D.CC.SUBSCRIBER.QUEUE	4	0.04	0 [0/0]

Description	Displays statistics on WebSphere MQ and JMS infrastructure activity, assisting successful queue management. To access: Select Applications > TransactionVision > Transaction Infrastructure > WebSphere MQ and JMS Statistics Report
Important Information	Click an item in the Total link in the table or the Total Count bar in the chart, to view transaction details for all the transactions for a given day, in a new window.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Reporting Time Period	Offers predefined time periods, or click the calendar icon and select a date, choose Date Range and enter, or select All Time Periods and include all statistics for the project.
Queue Filter	By default all queues in the project are included in the report. To restrict the report, check the queue filter and enter the desired queues (or the * wildcard character), separating queues with commas.
Queue Manager Filter	By default all queue managers in the project are included in the report. To restrict the report, check the queue manager filter and enter the desired queue managers (or the * wildcard character), separating them with commas.

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Applications Interacting	The number of applications interacting within this queue or queue manager.
Average Latency (sec)	The average latency for this queue or queue manager.
Gets (success, warning, error)	Click the event count link for a queue to view a chart and table showing additional information about the Get events for that queue. The results display in a new window.
Puts (success, warning, error)	Click the event count link for a queue to view a chart and table showing additional information about the Put events for that queue. The results display in a new window.
Total/Total Count	Move the cursor over the Total Count bar in the chart to display a tooltip showing its value.
Queue	The queue for which the data is displayed.
Queue Manager	The queue manager for which the data is displayed.

21

Transaction Summary Reports

This chapter includes the main concepts, tasks, and reference information for TransactionVision Summary reports.

Concepts

- Transaction Summary Reports Overview on page 588
- Save Report Settings on page 590

Reference

- Transaction Summary Reports User Interface on page 591

Transaction Summary Reports Overview

TransactionVision includes a number of reports that show summary, overview or over time transaction class information. They are located under **Applications > TransactionVision > Transaction Summary**, and include the following:

Report Name	Description	Use This Report To...
Capacity Planning	Displays transactional throughput data, and projects future capacity levels. For more details, see “Capacity Planning Report” on page 592.	see projections of future transaction response times based on current trends.
Transaction Dashboard	Shows how current transactions are performing. For more details, see “Transaction Dashboard Report” on page 595.	see transaction response times in real time.
Transaction Over Time	Displays the count or monetary values of various measurement correlations and transaction states, over a period of time. For more details, see “Transaction Over Time Report” on page 598.	see at a glance how successfully your transactions are performing.
Transaction Performance Summary	Displays a performance summary of selected transactions. For more details, see “Transaction Performance Summary Report” on page 603.	check on transaction response times.

Report Name	Description	Use This Report To...
Transaction Scorecard	Displays transaction performance success/failure rates. For more details, see “Transaction Scorecard Report” on page 607.	see average transaction response times within a specific time period.
Transaction Volume Analysis	Displays the volume of successful, failed, or late transactions. For more details, see “Transaction Volume Analysis Report” on page 610.	view the volume of successful, failed, or late transactions.

 **Save Report Settings**

You can save report settings after running a report. If a report could be useful again at another time, save the settings and give the report a meaningful name.

Example

A manager needs to see all Fund Transfers for the past week. Select the **Transaction Scorecard Report**, set the **Reporting Time Period** parameter to **This Week**, then select **Fund Transfer** from the **Include Transaction Classes** list. If desired, click **Save Settings** and give the report a meaningful name for future use.

Report Parameters

Reporting Time Period:

This Week

Include Transaction Classes:

☐ -Unclassified-

☐ Bond

☐ Equity

☒ Fund Transfer

Select AllUnselect All

Generate Report

Printable View

Save Settings

Save the current settings as:

Report Name:

Weekly Fund Transfers

Description:

All Fund Transfers from present week

OK

Cancel

Saved reports are displayed in the **Saved Settings** drop-down on the original report page:



Because this user created this report type, the trash icon displays on the right, giving the option to delete it.

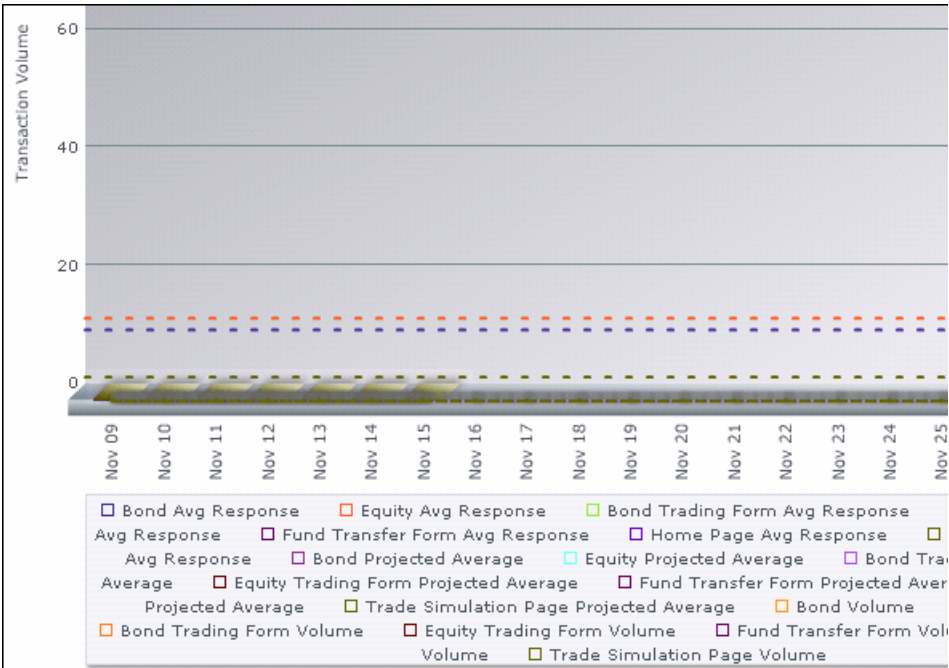
Transaction Summary Reports User Interface

This section describes:

- Capacity Planning Report on page 592
- Transaction Dashboard Report on page 595
- Transaction Over Time Report on page 598
- Transaction Performance Summary Report on page 603
- Transaction Scorecard Report on page 607
- Transaction Volume Analysis Report on page 610
- Business Transactions Filter Selection Dialog Box on page 613

 **Capacity Planning Report**

The following is an example of the Capacity Planning report.



Description	Displays transactional throughput data, and projects future capacity levels. To access: Select Applications > TransactionVision > Transaction Summary > Capacity Planning Report
--------------------	--

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Generate Report	Click to generate the report.
Printable View	Click for a report view that is easy to print.
Save Settings	Click to save the settings of this report. See “Save Report Settings” on page 590.
Project	Select the project you wish to report on.
Saved Settings	If this report was run before and the settings saved, it appears in this drop-down. If you are the user who saved these report settings, a trash can icon appears, offering the possibility to delete it.
Combine Transaction Class Totals	Check this if you wish to show combined totals for all selected transaction classes, instead of totals for each selected class.
Forward Projections	Specify the number of periods into the future to project: One Period , Two Periods , or Three Periods , of the same type as the reporting time period.
Include Transaction Classes	Select the transaction classes to include in the report.
Output Results as	Select Graph and/or Table .
Polynomial Degree	Specify the polynomial degree of analysis to perform when making projections: Linear , Quadratic , Cubic , Fourth Order , Fifth Order , Sixth Order .
Reporting Time Period	Select a time period, click the calendar icon and select a date, choose Date Range and enter a date, or select All Time Periods and include all statistics for the project.

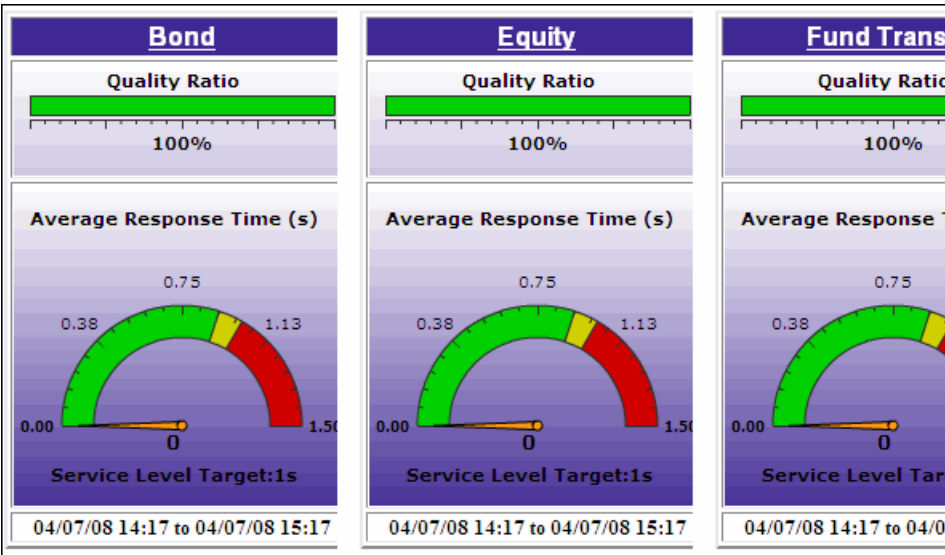
Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Time period>	The x-axis on the graph, or the fields in the table, specify the times, depending on the time period selected.
<Transaction Classes>	The selected transaction classes are plotted on the graph, or listed in the table.
Response Time (sec)	The y-axis on the right side of the graph, or the fields in the table, specifies the average response time for the transaction class.
Transaction Volume	The y-axis on the left side of the graph, or the fields in the table, specify the transaction volume measurements.

 **Transaction Dashboard Report**

The following is an example of the Transaction Dashboard report.



Description	<p>Displays gauges updated in real time of the average transaction response time for each selected transaction class. Additionally, it provides a Quality Ratio measurement which displays the percentage of transactions that have not exceeded the SLA value.</p> <p>To access: Select Applications > TransactionVision > Transaction Summary > Transaction Dashboard</p>
Important Information	<ul style="list-style-type: none">➤ Click the transaction class name on a gauge in this report to run the Transaction Tracking report for that transaction class.➤ Click the quality ratio gauge on the report to run the Transaction Tracking report on transactions that have exceeded required service levels.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Generate Report	Click to generate the report.
Printable View	Click for a report view that is easy to print.
Save Settings	Click to save the settings of this report. See “Save Report Settings” on page 590.
Project	Select the project you wish to report on.
Saved Settings	If this report was run before and the settings saved, it appears in this drop-down. If you are the user who saved these report settings, a trash can icon appears, offering the possibility to delete it.
Dashboard will show data from the past.... minutes	Enter the desired time frame for the report, in minutes. The default value is 60 minutes.
Include Transaction Classes	Select the transaction classes to include in the report.
Refresh Dashboard interval (in minutes)	Enter the desired time frame for refreshing the report, in minutes. TransactionVision continues to update the report results at this interval as long as this report page is open.

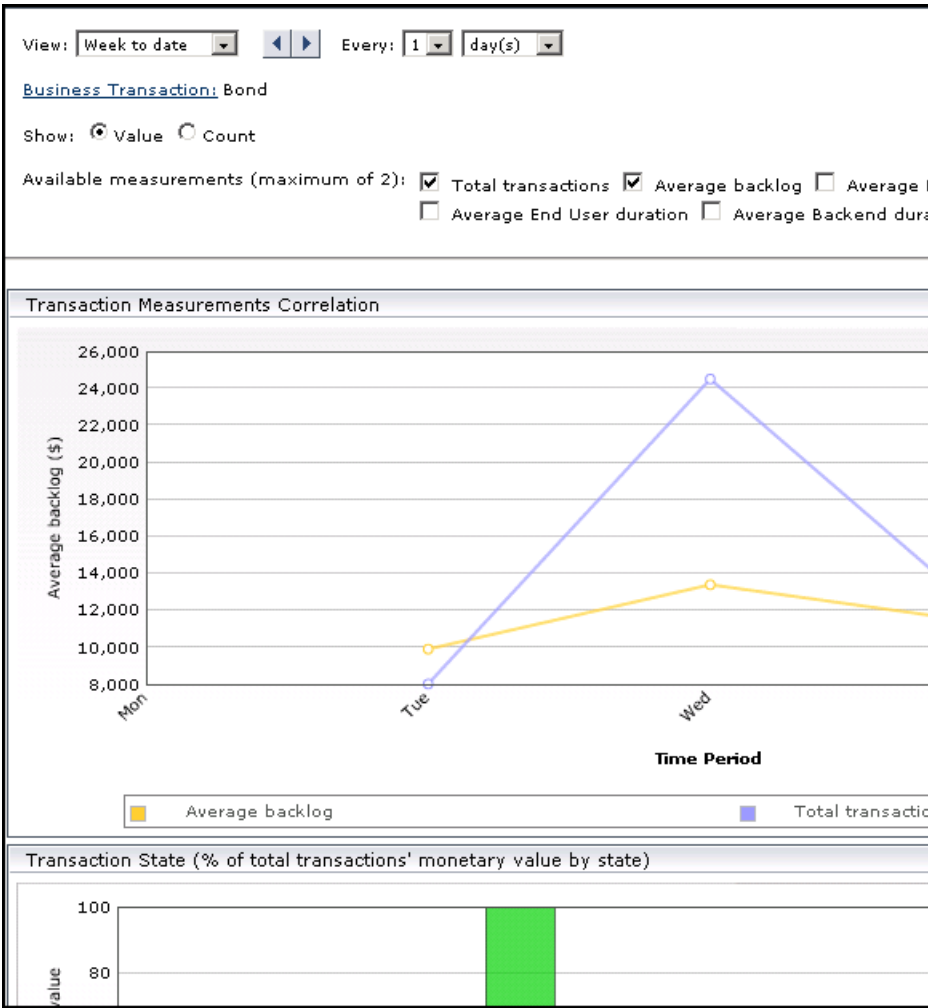
Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<time period>	The time range displays according to the report parameters set.
<transaction class>	A gauge displays for each transaction class.
Average Response Time(s)	Displays the average transaction response time for each selected transaction class
Quality Ratio/ Service Level Target	This displays the percentage of transactions that have not exceeded the SLA value.

 **Transaction Over Time Report**

The following is an example of the Transaction Dashboard report.



Description	<p>Displays the count or monetary values for various measurement correlations and transaction states of selected business transactions, over time.</p> <p>To access: Select Applications > TransactionVision > Transaction Summary > Transaction Over Time Report</p>
Important Information	<p>The metrics (calculated by either volume or by monetary value, depending on the selected mode) are shown in two graphs:</p> <ul style="list-style-type: none">➤ Transaction Measurements Correlation. A line graph showing the correlation between various measurements. Select from total transactions, backlog, average End-to-End duration, average End User duration, and average Backend duration.➤ Transaction State. A bar graph showing the percentage results for transaction states (exceptions, delays, failures) out of total transactions.

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Common report settings>	For details, see “Common Report Elements” in <i>Reports</i> .
Available measurements	<p>Select the check boxes for either one or two measurements, to display the measurements in the Transaction Measurements Correlation graph:</p> <ul style="list-style-type: none">➤ Total transactions. Displays total volume or total value for completed transaction instances.➤ Average backlog. Displays average backlog (transaction instances in-process) count or value for the time period.➤ Average End-to-End duration. Displays the average duration of total completed transaction instances (including the end user perspective), for the time period.➤ Average End User duration. Displays the average duration of completed transaction instances, from the end user perspective alone, for the time period.➤ Average Backend duration. Displays the average duration of completed transaction instance response time, from the TransactionVision perspective alone, for the time period.

GUI Element	Description
Business Transaction	<p>Displays the selected business transaction type, for inclusion in the report.</p> <p>Click Business Transactions to launch the filter selection dialog box, where you can select items to report on, see “Business Transactions Filter Selection Dialog Box” on page 613.</p>
Show	<p>Select the required radio button:</p> <ul style="list-style-type: none">➤ Value. The graphs display results based on the monetary value of each instance for the business transaction.➤ Count. The graphs display results based on the volume of instances for the business transaction.









Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Drill down>	Click a point or a bar in the graphs to drill down in the Transaction Over Time report to a finer granularity for the transaction. For example, in a report showing results per week, you can drill down from a particular week to show results per day in the week. You can then drill down again from a particular day to show results per hour in that day.
Transaction Measurements Correlation graph	Displays the time period on the x-axis, and the units for the selected measurements on the y-axes. Each point on the graph shows the total or average for the preceding time period. Tooltip: Hold the cursor over a point in the graph to see the value for the measurement.
Transaction State	Displays the time period on the x-axis, and percentage on the y-axis. A separate bar is shown for each transaction state during each time period included in the graph. Tooltip: Hold the cursor over a bar in the graph to see the values for the measurement and the percentage calculation.

Transaction Performance Summary Report



The following is an example of the Transaction Performance Summary report.

View: Week to date ◀▶					
Business Transactions : Filtered(9)					
Show: <input type="radio"/> Value <input type="radio"/> Count <input checked="" type="radio"/> Both					
Value					
   					
Business Transaction ▲	Total Value	Avg. End-to-End Duration (sec.)	Avg. End User Duration (sec.)	Avg. Backend Duration (sec.)	
Bond (\$)	200,396,711.79	2.14	0.00	0.01	3
Bond Trading Form (\$)	-	0.32	0.00	0.00	
Equity (\$)	187,509,825.00	2.03	0.00	0.02	2
Equity Trading Form (\$)	-	0.30	0.00	0.00	
Fund Transfer (\$)	189,258,500.00	1.97	0.00	0.01	
Fund Transfer Form (\$)	-	0.34	0.00	0.00	
Home Page (\$)	-	0.05	0.00	0.05	
Trade Simulation Page (\$)	-	0.06	0.00	0.00	
◀					
Count					
   					
Business Transaction ▲	Total count	Avg. End-to-End Duration (sec.)	Avg. End User Duration (sec.)	Avg. Backend Duration (sec.)	Avg. Backlog
Bond	3,347	2.14	0.00	0.01	2.07
Bond Trading Form	235	0.32	0.00	0.00	0.00
Equity	3,457	2.03	0.00	0.02	29.76

Description	<p>Displays metrics for selected transactions. The metrics (calculated by either volume or by monetary value, depending on the selected mode) are shown broken down by different transaction state categories, such as failures or delays.</p> <p>You can use this report to see summarized results for transactions during a set time period.</p> <p>To access: Select Applications > TransactionVision > Transaction Summary > Transaction Performance Summary Report</p>
--------------------	--

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
	<p>Click this button to drill down to the Transaction Over Time report for the transaction currently selected in the table. The report opens using the settings selected for the Transaction Performance Summary report.</p> <p>The Transaction Over Time report shows monetary values if you drill down from the Value table, and volume if you drill down from the Count table.</p>
	<p>Click this button to drill down to the Instance Topology report for the transaction currently selected in the table. The report opens using the settings selected for the Transaction Performance Summary report.</p> <p>The Transaction Over Time report shows monetary values if you drill down from the Value table, and volume if you drill down from the Count table.</p>
<Common report settings>	For details, see “Common Report Elements” in <i>Reports</i> .

GUI Element	Description
Business Transactions	Click Business Transaction to launch the filter selection dialog box, where you can select items to report on, see “Business Transactions Filter Selection Dialog Box” on page 613. Default: All business transactions are selected.
Show	Select the required radio button: <ul style="list-style-type: none"> ► Value. Displays the Value table, showing metrics based on the monetary value for each business transaction instance. ► Count. Displays the Count table, showing metrics based on the volume of business transaction instances in each category. ► Both. Displays both the Value and the Count tables.

Report Content

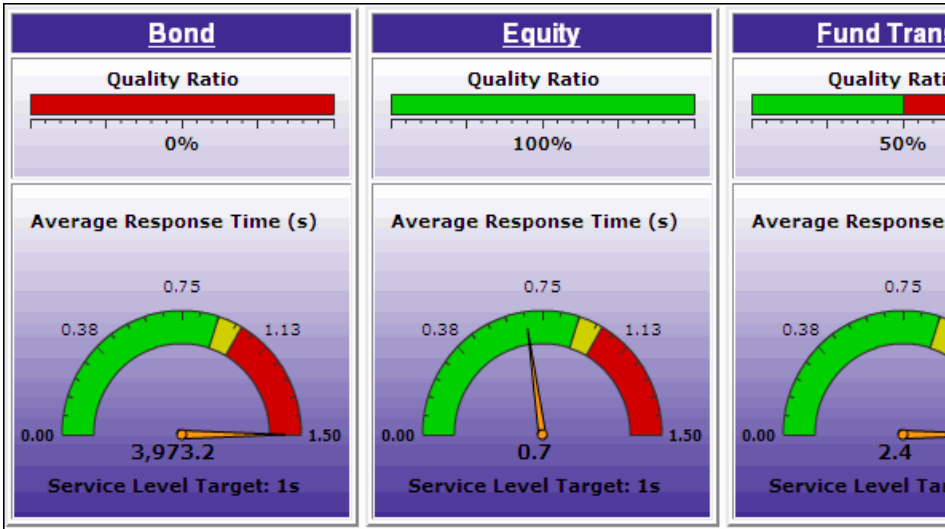
The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Avg. Backlog	(Count table) Displays the average number of backlogged (in-process) instances of the business transaction during the time range.
Avg. Backlog Value	(Value table) Displays the average monetary value of the backlogged (in-process) instances of the business transaction during the time range.
Avg. End-to-End Duration	Displays the average duration (in seconds) for all successfully completed instances of the business transaction, including end user perspective, triggered during the time range.
Avg. End User Duration	Displays the average duration of completed transaction instances, from the end user perspective alone, for the time range.

GUI Element	Description
Avg. Backend Duration	Displays the average duration of completed transaction instances response time, from the TransactionVision perspective alone, for the time range.
Business Transaction	Lists the Business Transactions CIs selected for inclusion in the report.
Delays	Displays the percentage of business transaction instances classified as late, or the percentage cost for the business transaction instances classified as late, during the time range (for both in-process and completed instances). Tooltip: The tooltip for each cell shows the values used to calculate the percentage.
Exceptions	Displays the percentage of business transaction instances classified as exceptions, or the percentage cost for the business transaction instances classified as exceptions, during the time range (for both in-process and completed instances). Tooltip: The tooltip for each cell shows the values used to calculate the percentage.
Failures	Displays the percentage of business transaction instances classified as failures, or the percentage cost for the business transaction instances classified as failures, during the time range (for both in-process and completed instances). Tooltip: The tooltip for each cell shows the values used to calculate the percentage.
Total Count	(Count table) Displays the total number of completed instances of the business transaction during the time range.
Total Value	(Value table) Displays the total monetary value for all completed instances of the business transaction during the time range.

 **Transaction Scorecard Report**

The following is an example of the Transaction Scorecard report.



Description	<p>Displays a gauge of the average transaction response time in a specified time period for each selected transaction class. Additionally, it provides a Quality Ratio measurement which displays the percentage of transactions that have not exceeded the service level value.</p> <p>To access: Select Applications > TransactionVision > Transaction Summary > Transaction Scorecard</p>
--------------------	---

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element	Description
Generate Report	Click to generate the report.
Printable View	Click for a report view that is easy to print.
Save Settings	Click to save the settings of this report. See “Save Report Settings” on page 590.
Project	Select the project you wish to run the report for.
Saved Settings	If this report was run before and the settings saved, it appears in this drop-down. If you are the user who saved these report settings, a trash can icon appears, offering the possibility to delete it.
Include Transaction Classes	Select the transaction classes to include in the report.
Reporting Time Period	Select the reporting time period from the drop-down list of options. To specify a custom time period, choose Date Range and specify the start and end dates and times. Click the calendar icon to select a date from a calendar control. To include all statistics for the entire project, choose All Time Periods .

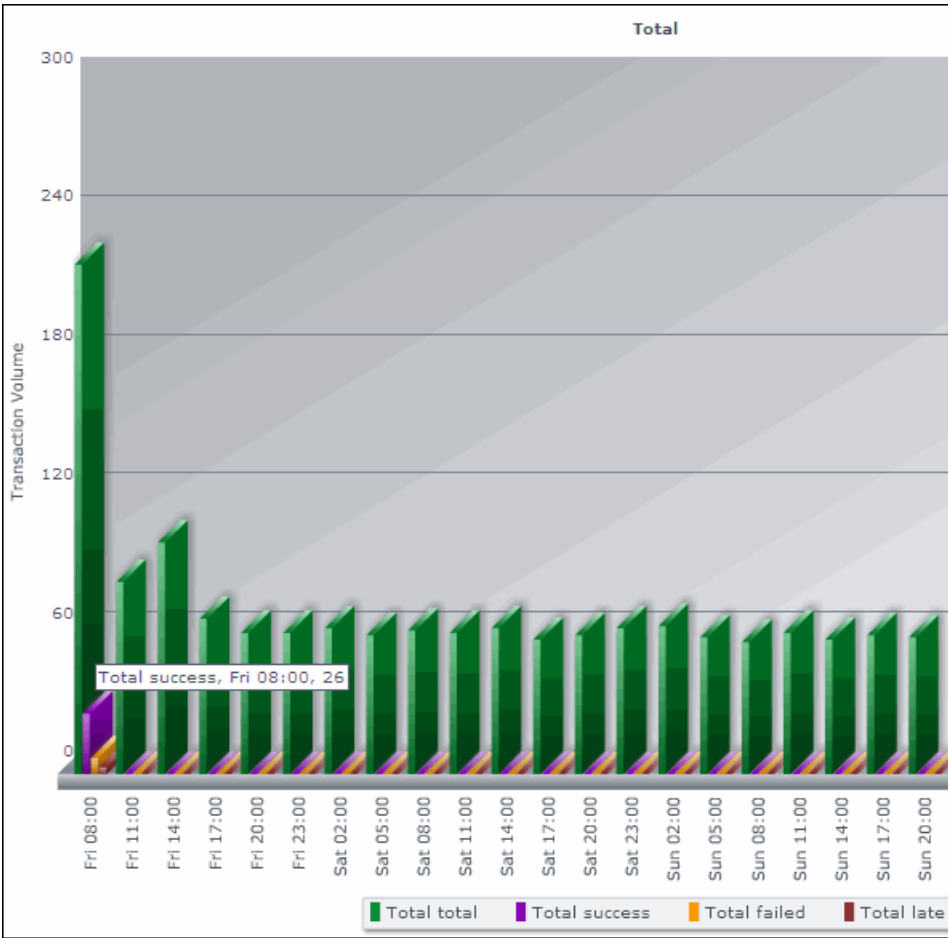
Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<time period>	The time range displays, as set in the report parameters.
<transaction class>	A gauge displays for each transaction class. Click the transaction class name on a gauge to run the Transaction Tracking report for that transaction class.
Average Response Time(s)	Displays the average transaction response time for each selected transaction class
Quality Ratio/ Service Level Target	This displays the percentage of transactions that have not exceeded the SLA value. Click the quality ratio gauge to run the Transaction Tracking report on transactions that have exceeded required service levels.

 **Transaction Volume Analysis Report**

The following is an example of the Transaction Volume Analysis report.



Description	<p>Displays a chart showing the volume of successful, failed, or late transactions for a given transaction class.</p> <p>To access: Select Applications > TransactionVision > Transaction Summary > Transaction Volume Analysis</p>
-------------	--

Report Settings

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<Common report elements>	For details on the user interface, see “Common Report Elements” in <i>Reports</i> .
Combine Transaction Class Totals	Check this if you wish to show combined totals for all selected transaction classes, instead of totals for each selected class.
Include Transaction Classes	Select the transaction classes to include in the report.
Output results as	Output the report as a graph and/or table.
Reporting Time Period	Select the reporting time period from the drop-down list of options. To specify a custom time period, choose Date Range and specify the start and end dates and times. Click the calendar icon to select a date from a calendar control. To include all statistics for the entire project, choose All Time Periods .
Series Selection	Select one or more from: <ul style="list-style-type: none"> ➤ Show successful transactions ➤ Show late transactions ➤ Show failed transactions ➤ Show total transactions

Report Content

The following elements are included (unlabeled GUI elements are shown in angle brackets):

GUI Element	Description
<data item value>	To view more information about a data item in a chart, move the cursor over the data item. TransactionVision displays a tool-tip that shows the value of the data item.
<series>	To run the Transaction Tracking report on a series for a given date, click the bar representing that series and date. TransactionVision displays the results in a separate window.
<time periods>	The times as set in the report parameters.
Transaction Volume	The volume of successful, failed, or late transactions for a given transaction class

Business Transactions Filter Selection Dialog Box

Description	Enables you to select Business Transactions to include in the report. To access: From the report page, click on the Business Transactions/Business Transaction link.
Included in Tasks	Setup for “Transaction Over Time Report” on page 598, and “Transaction Performance Summary Report” on page 603.

The following elements are included (unlabeled GUI elements are shown in angle brackets>):

GUI Element (A-Z)	Description
Cancel	Click to cancel the selection.
OK	Click to confirm the selection and return to the report settings page.
Business Transactions	A list of all available transactions of that type.
View	Any transaction categories for selection display here.

Index

A

- Add Filter page 262
- Aggregated Topology
 - calculating response time 447
- alias management 312
 - defined 311
- analyzer
 - COLLECTING EVENTS status 100
 - configuring to store and use raw events 96
 - create or edit host 72
 - exception message queue 94
 - FAIL status 100
 - NOT COLLECTING EVENTS status 100
 - QUIESCING status 100
 - recovery mode 97
 - RECOVERY status 100
 - ServicesManager script 409, 412
 - status values 100
 - storing and using raw events 95
 - time skew settings 92
 - trace logging 98
 - using ManageQueue to store and use raw events 95
- Analyzer security
 - TIBCO EMS permissions 117
 - WebSphere MQ permissions 113
- Analyzer.properties 94, 98
- analyzer.properties 74
- API name query
 - BTTrace 326
 - CICS 327
 - WebSphere MQ API 342
 - WebSphere MQ IMS Bridge 351
- API type query
 - CICS 327

- application query
 - EJB 331
- application server query
 - EJB 330
 - JDBC 332
- Application Server Statistics report 572
- applications, troubleshooting 427
- Attribute 285

B

- BAC Aggregate TV Sample Data 364
- BAC integration
 - drill downs from RUM 366
- BPI Action 305
- BTTrace
 - API name query 326
 - severity query 327
- BTTRACE API 257
- BTTRACE severity 257
- business transactions 613

C

- Capacity Planning report 592
- CICS
 - API name query 327
 - API type query 327
 - Event Analysis columns 557
 - file name query 327
 - response code query 328
 - SYSID query 328
 - TD queue name query 328
 - Terminal ID query 329
 - Transaction ID query 330
 - TS queue ID query 329
- CICS sensor

- API name filter option 225
 - file name filter option 226
 - response code filter option 226
 - SYSID filter option 227
 - TD queue name filter option 227
 - terminal id filter option 228
 - transaction id filter option 225
 - TS queue id filter option 228
 - user id filter option 226
 - class query
 - JDBC 333
 - JMS 336
 - Classify 285
 - classify events in TransactionVision 372
 - classify section 285
 - client applications 424
 - Committing 284
 - common technology
 - Event Analysis columns 557
 - communication links
 - configuration queue 123, 124, 125, 131, 133, 141, 143, 185, 186, 191, 196
 - configuration queues 106, 107
 - configurations 109
 - connection type 129, 193
 - create new 121, 130
 - create new JMS Provider 142
 - create new SonicMQ 135
 - create new WebLogic JMS 140
 - defined for TransactionVision 17
 - event connection and delivery retry 126
 - event queue 107, 123, 125, 131, 141, 143, 168, 186, 191
 - event queues 107
 - exception message queues 108
 - global communication link templates 117
 - link types 122
 - overview 106
 - SonicMQ configurations 110
 - test 144
 - TIBCO EMS configurations 109
 - troubleshooting 212
 - troubleshooting problems 212
 - understanding 106
 - user event processing 128, 194
 - view global communication link templates 117
 - WebLogic JMS configurations 110
 - WebSphere MQ configurations 109
 - completion code query
 - WebSphere MQ 345
 - completion codes, color-coded in Event Analysis 551
 - component topology analysis
 - configure edge labels 514
 - edges 494
 - JDBC edges 498
 - modify component groupings 519
 - modifying JDBC edge labels 517
 - modifying layouts and properties 502
 - modifying node labels 518
 - printing graphs 508
 - resource icons 528
 - conditional operator 286
 - configuration message expiry, for Analyzers 72
 - configuration queue 107, 123, 124, 125, 131, 133, 141, 143, 185, 186, 191, 196
 - configuration queues, communication links 107
 - connection name query
 - JMS 337
 - WebSphere MQ 342
 - WebSphere MQ IMS Bridge 352
 - connection type 129
 - correlation ID query
 - WebSphere MQ 345
 - correlation ID query, WebSphere MQ IMS Bridge 353
 - CreateSqlScript 385, 427
 - current projects tab
 - TransactionVision 40
- D**
- data collection filters 425
 - Add Filter page 262
 - BTTRACE API 257

- BTTRACE severity 257
- CICS api 225
- CICS SYSID 227
- CICS TD queue name 227
- CICS terminal id 228
- CICS TS queue id 228
- completion codes 249
- connection name 236
- create 258
- data range 222
- defined for TransactionVision 19
- edit 260
- Edit page 263
- EJB application name 230
- EJB application server 229
- EJB exception setting 231
- EJB method 231
- EJB name 230
- event packaging 223
- file name 226
- filter conditions 219, 260
- host 220
- i5/OS job names 257
- JDBC class 231
- JDBC database 232
- JDBC database object 232
- JDBC methods 233
- JDBC result code 233
- JDBC SQL code 233
- JDBC SQL state 234
- JDBC SQL statement 234
- JDBC web applications 234
- JMS class 235
- JMS exception setting 237
- JMS methods 235
- JMS queue 236
- JMS topic 236
- JMS web applications 235
- MQIMS Bridge API 251
- overview 215
- planning 216
- reason codes 249
- response code 226
- servlet application server 237
- servlet client host/IP 238
- servlet exception setting 239
- servlet methods 238
- servlet status code 239
- servlet URI 239
- servlet web application 237
- transaction id 225
- user id 226
- WBI message brokers 251
- WBI message flow 251
- WebSphere MQ API name 248
- WebSphere MQ completion codes 249
- WebSphere MQ objects 249
- WebSphere MQ reason code 249
- WebSphere MQ ReplyTo object 250
- z/OS CICS SYSID 253
- z/OS CICS task 252
- z/OS CICS transactions 253
- z/OS IMS identifier 255
- z/OS IMS PSB 256
- z/OS IMS region identifier 254
- z/OS IMS region type 254
- z/OS IMS transaction 255
- z/OS jobs and steps 252
- data size query
 - JMS 339
 - Servlet 341
 - WebSphere MQ 345
- database object query
 - JDBC 333
- database query
 - JDBC 333
- database schema
 - CreateSqlScript script 385
 - defined for TransactionVision 20
 - DeleteEvents script 392
 - manually deleting 427
- database schemas 65
 - Manager UI description 66
 - overview 65
 - user interface 66
- DataUtil 389
- DB2_RR_TO_RS 128
- DeleteEvents 392

E

- edge labels
 - modifying JDBC 517
- Edit Data Collection Filter page 263
- edit schema associations 314
- EJB
 - application query 331
 - application server query 330
 - Event Analysis columns 557
 - method query 331
 - name query 331
 - status query 332
- EJB sensor
 - application name filter option 230
 - application server filter option 229
 - classes instrumented 418
 - EJB method filter option 231
 - EJB name filter option 230
 - methods instrumented 418
 - troubleshooting 426
- enable session tracking 541
- evaluate section 297
- Event Analysis
 - Event Details description 537
- Event Analysis
 - changing default settings 537
 - CICS columns 557
 - color-coded completion codes 551
 - common columns 557
 - EJB columns 557
 - event attributes 545
 - Event Details description 537
 - event sort order 551
 - JDBC columns 557
 - JMS columns 557
 - navigation links 551
 - options dialog 555
 - overview 535
 - page 551
 - servlet columns 557
 - sorting many events 535
 - transaction columns 557
 - User Event columns 557
 - WebSphere MQ columns 557
 - WebSphere MQ i5/OS columns 557
 - WebSphere MQ z/OS Batch columns 557
 - WebSphere MQ z/OS IMS columns 557
- event analysis 535
 - set display options 542
- event attributes 545
- event collection
 - defined for TransactionVision 17
- event collection threads 128, 188, 194, 199, 203, 207
- event connection retry 126
- event delivery retry 126
- event details page 558
- Event Details view
 - overview 558
- event queue 107, 123, 125, 131, 141, 143, 168, 186, 191
 - maximum message length 188, 193, 199, 203, 207
- event queue message persistency, WebSphere MQ requirement 112
- event queues
 - message persistency requirements for TIBCO EMS 116
 - queue manager requirements for WebSphere MQ 112
 - SonicMQ requirements 114
 - storage requirements for SonicMQ 115
 - storage requirements for TIBCO EMS 116
 - storage requirements for WebSphere MQ 112
 - TIBCO EMS requirements 116
- event queues, communication links 107
- events
 - failure to put on event queue 126
 - storing and using raw events 95
 - troubleshooting missing events 425
- exception class name query
 - JMS 340
- exception code query
 - JMS 339
- exception message queue 94, 127, 188, 193, 199, 203, 207

exception message queues 108
export BPI events 378

F

Failure Mode
 for Analyzers 74
Failure mode
 enabling 75
 how it works 76
file name query
 CICS 327
files
 Trade Demo Events
 TransactionVision 25
filters 257
 i5/OS job names 257
filters_userevent_criteria 239
FlushStatusUtil 98, 399, 400

G

global communication link templates 117

I

i5/OS
 WebSphere MQ job name query 350
icons
 component topology analysis 528
Information 385
infrastructure reports UI 571
Instance Topology
 calculating response time 455
 Transaction Topology flow map 486

J

Java programs
 monitoring 423
JDBC
 application server query 332
 class query 333
 database object query 333
 database query 333
 edges in component topology analysis
 view 498

Event Analysis columns 557
modifying edge labels 517
result code query 334
SQL code query 334
SQL State query 335
SQL statement query 335
SQL statement type query 335
web application query 335

JDBC sensor

class filter option 231
database filter option 232
database object filter option 232
method filter option 233
methods tracked 419
result code filter option 233
SQL code filter option 233
SQL state filter option 234
SQL statement filter option 234
web application filter option 234

JMS

class query 336
connection name query 337
data size query 339
Event Analysis columns 557
exception class name query 340
exception code query 339
method query 337
queue query 338
topic query 337
web application query 336, 341
web application server query 340
WebSphere MQ queue manager query
 338
WebSphere MQ queue query 339

JMS sensor

class filter option 235
connection name filter option 236
exception setting filter option 237
JMS queue filter option 236
method filter option 235
methods tracked 416
topic filter option 236
troubleshooting 427
web application filter option 235

job name query

i5/OS for WebSphere MQ events 350

L

- labels
 - modifying nodes 518
- LD_LIBRARY_PATH environment variable 213, 423
- LIBPATH environment variable 213, 424

M

- manage alias lists 312
- ManageQueue 94, 128
 - using to store and use raw events for the analyzer 95
- Match 286
- match conditions 286
- maximum event message length 125
- message ID query
 - WebSphere MQ 345
- message ID query, WebSphere MQ IMS Bridge 353
- Message Latency Analysis report 575
- message length 125
- message length, WebSphere MQ
 - configuration queue requirement 111
- method query
 - EJB 331
 - JMS 337
 - Servlet 341
- modifying
 - node labels 518

N

- name query
 - EJB 331
- node labels, modifying 518

O

- object query
 - WebSphere MQ 343
 - WebSphere MQ IMS Bridge 353
- optimizing performance 424

P

- PATH environment variable 213, 423
- Performance Dashboard report 577
- performance optimization 424
- Perl programs
 - monitoring 423
- permissions
 - SonicMQ Sensor security 115
 - TIBCO EMS Analyzer security 117
 - TIBCO EMS Sensor security 116
 - WebSphere MQ Analyzer security 113
 - WebSphere MQ Sensor security 113
- printing
 - component topology analysis graphs 508
- project status 51
- projects
 - defined for TransactionVision 19
 - DeleteEvents script 392
 - edit projects page 44
 - project wizard 53
 - UI page 48
- publishing transaction class to CMDB 363

Q

- queries
 - BTTrace API name 326
 - BTTrace severity 327
 - CICS API name 327
 - CICS API type 327
 - CICS file name 327
 - CICS response code 328
 - CICS SYSID 328
 - CICS TD queue name 328
 - CICS Terminal ID 329
 - CICS Transaction ID 330
 - CICS TS queue ID 329
 - EJB application 331
 - EJB application server 330
 - EJB method 331
 - EJB name 331
 - EJB status 332
 - JDBC application server 332
 - JDBC class 333
 - JDBC database 333

- JDBC database object 333
- JDBC result code 334
- JDBC SQL code 334
- JDBC SQL state 335
- JDBC SQL statement 335
- JDBC SQL statement type 335
- JDBC web application 335
- JMS class 336
- JMS connection name 337
- JMS data size 339
- JMS exception class name 340
- JMS exception code 339
- JMS method 337
- JMS queue 338
- JMS topic 337
- JMS web application 336
- JMS WebSphere MQ queue 339
- JMS WebSphere MQ queue manager 338
- Servlet data size 341
- Servlet method 341
- Servlet status code 341
- Servlet web application 341
- Servlet web application server 340
- WebSphere MQ API name 342
- WebSphere MQ completion code 345
- WebSphere MQ connection name 342
- WebSphere MQ correlation ID 345
- WebSphere MQ data size 345
- WebSphere MQ i5/OS job name 350
- WebSphere MQ IMS Bridge API name 351
- WebSphere MQ IMS Bridge
 - connection name 352
- WebSphere MQ IMS Bridge
 - correlation ID 353
- WebSphere MQ IMS Bridge IMS
 - transaction 351
- WebSphere MQ IMS Bridge message ID 353
- WebSphere MQ IMS Bridge object 353
- WebSphere MQ IMS Bridge queue manager 352
- WebSphere MQ message ID 345
- WebSphere MQ object 343
- WebSphere MQ queue manager 343
- WebSphere MQ reason code 345
- WebSphere MQ Reply To object 344
- WebSphere MQ reply to queue manager 344
- WebSphere MQ z/OS CICS SYSID 347
- WebSphere MQ z/OS CICS task 347
- WebSphere MQ z/OS CICS transaction 347
- WebSphere MQ z/OS IMS identifier 349
- WebSphere MQ z/OS IMS PSB 350
- WebSphere MQ z/OS IMS region identifier 348
- WebSphere MQ z/OS IMS region type 348
- WebSphere MQ z/OS IMS transaction 349
- WebSphere MQ z/OS job name 346
- WebSphere MQ z/OS job step 346
- queries analysis 535
- queue depth, WebSphere MQ queue requirement 111
- queue manager
 - troubleshooting channel limitations 426
- queue manager channel limitations 426
- queue manager query
 - WebSphere MQ 343
 - WebSphere MQ for JMS 338
 - WebSphere MQ IMS Bridge 352
- queue parameters
 - SonicMQ queue requirements 115
 - TIBCO EMS queue requirements 116
 - WebSphere MQ queue requirements 113
- queue query
 - WebSphere MQ for JMS 339
- queue query, JMS 338
- queue requirements
 - event queue manager 112
 - event queue message persistency 112
 - event queue storage 112
 - queue depth 111
 - queue parameters 113
- quiescing mode 100

R

- reason code query
 - WebSphere MQ 345
- rebind_sensor 428
- recovery mode 97
- remote queue connection
 - same SonicMQ server 137
- reply to queue manager query
 - WebSphere MQ 344
- reports
 - procedure to run 539
- requirements
 - TransactionVision queue 110
- response code query
 - CICS 328
- result code query
 - JDBC 334
- RUM drill downs for BAC integration 366
- runmqtrm 423

S

- save reports 590
- schemas
 - See database schema 20
- Sensor security
 - SonicMQ permissions 115
 - TIBCO EMS permissions 116
 - WebSphere MQ permissions 113
- sensors
 - trace logging 128, 188, 194, 199, 203, 207
- SensorSetup 426
- Service Level Analysis report 559
- ServicesManager 409, 412
- Servlet
 - data size query 341
 - Event Analysis columns 557
 - method query 341
 - status code query 341
- Servlet sensor
 - instrumented classes 414
 - methods tracked 415
 - troubleshooting 426
- servlet sensor
 - application server filter option 237

- client host/IP filter option 238
- exception setting 239
- servlet method filter option 238
- status code filter option 239
- URI filter option 239
- web application filter option 237
- severity query
 - BTTrace 327
- SHLIB_PATH environment variable 213, 422, 423
- SLA violations 77
- SonicMQ
 - communication link configurations 110
 - create the Sensor connection 135
 - event queue requirements 114
 - event queue storage requirement 115
 - queue parameter requirements 115
 - same server for a remote queue connection 137
 - Sensor security permissions 115
- SQL code query
 - JDBC 334
- SQL state query, JDBC 335
- SQL Statistics report 580
- statement query, JDBC 335
- statement type query, JDBC 335
- StatisticsCache.properties 97
- status code query
 - Servlet 341
- status query
 - EJB 332
- SYSID query
 - CICS 328
 - z/OS CICS for WebSphere MQ events 347

T

- TD queue name query
 - CICS 328
- Terminal ID query, CICS 329
- test, communication link 144

- TIBCO EMS
 - Analyzer security permissions 117
 - communication link configurations 109
 - event queue message persistency requirement 116
 - event queue requirements 116
 - event queue storage requirements 116
 - queue parameter requirements 116
 - Sensor security permissions 116
 - time server 134
- time server
 - TIBCO EMS 134
- time skew
 - settings 92
- topic query, JMS 337
- Trade Demo Sample Events
 - installed files TransactionVision 25
 - using in TransactionVision 24
- transaction class attribute 304
- Transaction Dashboard report 595
- transaction detail reports 533, 534
- transaction detail reports UI 550
- Transaction ID query, CICS 330
- transaction infrastructure reports 569, 570
- transaction monitor CI 363
- Transaction over Time report 598
- Transaction Performance report 603
- transaction query
 - IMS for WebSphere MQ IMS Bridge events 351
 - z/OS CICS for WebSphere MQ events 347
 - z/OS IMS for WebSphere MQ events 349
- Transaction Scorecard report 607
- transaction summary reports
 - defined 588
 - UI 591
- Transaction Topology flow map
 - for Aggregated Topology 476
 - for Dashboard 484
 - for Instance Topology 486
 - overview 465
 - time frame 473
 - user interface 475
 - viewing 466
- Transaction Tracking report 562
- Transaction Volume Analysis report 610
- transactions
 - Event Analysis columns 557
- TransactionVision
 - basic concepts 16
 - components 16
 - data concepts 20
 - drill down to Diagnostics 538
 - overview 16
 - terms and concepts 17
- TransactionVision BAC settings 362
- TransactionVision business settings UI 376
- TransactionVision configuration for Business
 - Process Insight 376
- TransactionVision data in Business
 - Availability Center 362
- TransactionVision Data Sample Job
 - Configuration 379
- TransactionVision drill downs from End User
 - Management reports 366
- TransactionVision integration with Business
 - Process Insight 370
- TransactionVision integrations
 - troubleshooting 381
- TransactionVision projects
 - defined 43
 - UI 44
- TransactionVision queue requirements
 - WebSphere MQ 111
- TransactionVision queues
 - requirements 110
- TransactionVision settings for Business
 - Availability Center 379
- TransactionVision settings for Business
 - Process Insight 369
- transmission queue 186, 191
- triggered programs, monitoring 423
- troubleshooting
 - applications 427
 - communication links 212
 - EJB sensor 426
 - JMS Sensor 427
 - missing events 425
 - monitoring client applications 424

- monitoring Java programs 423
- monitoring Perl programs 423
- monitoring triggered programs 423
- queue manager channel limitations 426
- triggered programs 423
- web user interface 427
- TS queue ID query, CICS 329
- TVISION_BANNER environment variable 212
- TVISION_REPORT_ARGS environment variable 423
- TVISION_SYSLOG environment variable 213

U

- Unicode
 - troubleshooting 435
- User Event
 - Event Analysis columns 557
- user events 128, 194

V

- value 286

W

- web application query
 - JDBC 335
 - JMS 336
 - Servlet 341
- web application server query
 - Servlet 340
- Web Session report 565
- web user interface, troubleshooting 427
- WebLogic JMS
 - communication link configurations 110
- WebSphere MQ
 - API name query 342
 - communication link configurations 109
 - completion code query 345
 - connection name query 342
 - correlation ID query 345
 - data size query 345

- Event Analysis columns 557
- event queue manager requirements 112
- event queue message persistency requirement 112
- event queue storage requirements 112
- i5/OS job name query 350
- message ID query 345
- message length configuration queue requirements 111
- object query 343
- queue depth requirement 111
- queue manager query 343
- queue parameter requirements 113
- reason code query 345
- rebinding sensor after support pac installation 428
- Reply To object query 344
- reply to queue manager query 344
- TransactionVision queue requirements 111
- z/OS CICS SYSID query 347
- z/OS CICS task query 347
- z/OS CICS transaction query 347
- z/OS IMS identifier query 349
- z/OS IMS PSB query 350
- z/OS IMS region identifier query 348
- z/OS IMS region type query 348
- z/OS IMS transaction query 349
- z/OS job name query 346
- z/OS job step query 346
- WebSphere MQ and JMS Statistics report 583
- WebSphere MQ i5/OS, Event Analysis columns 557
- WebSphere MQ IMS Bridge
 - API name query 351
 - connection name query 352
 - correlation ID query 353
 - IMS transaction query 351
 - message ID query 353
 - object query 353
 - queue manager query 352
- WebSphere MQ sensor
 - API name filter option 248
 - BTTRACE API filter option 257
 - BTTRACE severity filter option 257

- completion code filter option 249
- i5/OS job name filter option 257
- MQIMS Bridge API filter option 251
- object filter option 249
- reason code filter option 249
- reply to object filter option 250
- WBI message broker filter option 251
- WBI message flow filter option 251
- z/OS CICS SYSID filter option 253
- z/OS CICS task filter option 252
- z/OS CICS transaction filter option 253
- z/OS IMS identifier filter option 255
- z/OS IMS PSB filter option 256
- z/OS IMS region identifier filter option 254
- z/OS IMS region type filter option 254
- z/OS IMS transaction filter option 255
- z/OS jobs and steps filter option 252
- WebSphere MQ z/OS IMS, Event Analysis columns 557
- WebSphere MQ z/OSBatch, Event Analysis columns 557

X

- XPath 286

Z

z/OS

- batch 18
- CICS API crossing exit mechanism 18
- IMS 18
- WebSphere MQ CICS SYSID query 347
- WebSphere MQ CICS task query 347
- WebSphere MQ CICS transaction query 347
- WebSphere MQ IMS identifier query 349
- WebSphere MQ IMS PSB query 350
- WebSphere MQ IMS region identifier query 348
- WebSphere MQ IMS region type query 348

- WebSphere MQ IMS transaction query option 349
- WebSphere MQ job name query 346
- WebSphere MQ job step query 346

