

HP Business Service Management

for the Windows and Linux operating systems

Software Version: 9.13

RTSM Developer Reference

Document Release Date: May 2012

Software Release Date: May 2012



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.

Copyright Notices

© Copyright 2005 - 2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

AMD and the AMD Arrow symbol are trademarks of Advanced Micro Devices, Inc.

Google™ and Google Maps™ are trademarks of Google Inc.

Intel®, Itanium®, Pentium®, and Intel® Xeon® are trademarks of Intel Corporation in the U.S. and other countries.

iPod is a trademark of Apple Computer, Inc.

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

Microsoft®, Windows®, Windows NT®, Windows® XP, and Windows Vista® 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 Open Group.

Acknowledgements

This product includes software developed by the Apache Software Foundation (**<http://www.apache.org>**).

This product includes software developed by the JDOM Project (**<http://www.jdom.org>**).

This product includes software developed by the MX4J project (**<http://mx4j.sourceforge.net>**).

Documentation Updates

The title page of this document 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

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 online support 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 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 register for an HP Passport ID, go to:

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

To find more information about access levels, go to:

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

Table of Contents

Welcome to This Guide	11
-----------------------------	----

PART I: CREATING DISCOVERY AND INTEGRATION ADAPTERS

Chapter 1: Adapter Development and Writing	19
Adapter Development and Writing Overview	20
Content Creation	21
Adapter Development Cycle	21
Data Flow Management and Integration	24
Associating Business Value with Discovery Development	26
Researching Integration Requirements	28
Developing Integration Content.....	31
Developing Discovery Content	34
How to Implement a Discovery Adapter.....	37
Step 1: Create an Adapter.....	40
Step 2: Assign a Job to the Adapter	49
Step 3: Create Jython Code	51
Chapter 2: Discovery Content Migration Guidelines	53
Discovery Content Migration Guidelines Overview.....	54
Version 9.x New Infrastructure Features	54
Guidelines for Developing Cross-Data Model Scripts.....	59
Implementation Tips.....	59
How to Access BTO Data Model Documentation Online.....	60
Package Migration Utility.....	61
Troubleshooting and Limitations	61

Chapter 3: Developing Jython Adapters.....	63
HP Data Flow Management API Reference	64
How to Create Jython Code	65
How to Support Localization in Jython Adapters.....	79
How to Work with Discovery Analyzer	90
How to Run Discovery Analyzer from Eclipse	99
How to Record DFM Code.....	109
Jython Libraries and Utilities	112
Chapter 4: Developing Generic Database Adapters.....	117
Generic Database Adapter Overview	119
TQL Queries for the Generic Database Adapter	119
Reconciliation.....	120
Hibernate as JPA Provider.....	121
How to Prepare for Adapter Creation	124
How to Prepare the Adapter Package	129
How to Upgrade the Generic DB Adapter from 9.00 or 9.01 to 9.02 and Later	133
How to Configure the Adapter – Minimal Method	134
How to Configure the Adapter – Advanced Method	138
How to Implement a Plugin.....	144
How to Deploy the Adapter	148
How to Edit the Adapter	148
How to Create an Integration Point.....	148
How to Create a View.....	149
How to Calculate the Results	150
How to View the Results	150
How to View Reports.....	150
How to Enable Log Files	150
How to Use Eclipse to Map Between CIT Attributes and Database Tables	151
Adapter Configuration Files	161
Out-of-the-Box Converters.....	186
Plugins	191
Configuration Examples	191
Adapter Log Files	203
External References	205
Troubleshooting and Limitations	205

Chapter 5: Developing Java Adapters	207
Federation Framework Overview	208
Adapter and Mapping Interaction with the Federation Framework	214
Federation Framework Flow for Federated TQL Queries.....	215
Federation Framework Flow for Population.....	231
Adapter Interfaces.....	233
How to Debug Adapter Resources	236
How to Add an Adapter for a New External Data Source	236
How to Implement the Mapping Engine.....	245
How to Create a Sample Adapter	247
XML Configuration Tags and Properties.....	249
Chapter 6: Developing Push Adapters	251
Developing Push Adapters Overview	252
Differential Synchronization.....	252
How to Prepare the Mapping Files	253
How to Write Jython Scripts	255
How to Support Differential Synchronization	259
How to Build an Adapter Package	261
Mapping File Schema	263
Mapping Results Schema.....	278
Chapter 7: Viewing KPIs in External Applications	285
How to Set Up an Adapter to View KPIs in External Applications....	286

PART II: USING APIS

Chapter 8: Introduction to APIs	291
APIs Overview.....	292
Chapter 9: RTSM (HP Universal CMDB) API	293
Conventions	294
Using the HP Universal CMDB API.....	294
General Structure of an Application	296
How to Put the API Jar File in the Classpath	299
How to Create an Integration User	299
HP Universal CMDB API Reference.....	302
Use Cases	302
Examples.....	303

Chapter 10: RTSM (HP Universal CMDB) Web Service API	305
Conventions	307
HP Universal CMDB Web Service API Overview	308
HP Universal CMDB Web Service API Reference	310
How to Call the Web Service	311
How to Query the RTSM	311
How to Update the RTSM.....	316
How to Query the BSM Class Model	318
How to Query for Impact Analysis.....	320
UCMDB General Parameters	321
UCMDB Output Parameters	324
UCMDB Query Methods	326
UCMDB Update Methods	340
UCMDB Impact Analysis Methods	343
Use Cases	346
Examples.....	347
Chapter 11: Data Flow Management API	385
Data Flow Management API Overview	386
Conventions	386
Discovery and Dependency Mapping Web Service	387
Call the Web Service	388
Discovery and Dependency Mapping Methods.....	389
Code Sample.....	402
Chapter 12: Internal-Global ID Conversion API	407
Using the Internal-Global ID Conversion API	408
How to Convert RTSM Internal IDs to Global IDs	409
How to Convert Global IDs to RTSM Internal IDs	411
Index	413

Welcome to This Guide

This guide explains how to create and manage adapters that enable you to send and receive data from external data repositories, another CMDB (UCMDB or third-party CMDB such as Atrium) or another RTSM. Note that RTSM is BSM's operational storage and should be populated with CIs that are relevant for operational use cases.

RTSM is an instance of CMDB and all integration adapters and APIs that are available in standalone UCMDB are available in RTSM as well. For other third-party integrations, a customer can develop his own integration adapter using population, federation or data push integrations. The information regarding available third-party integrations can be found in the HP BSM Documentation Library and [HP Live Community](#).

For more information on RTSM integration strategy, see the [RTSM Best Practices PDF](#).

This chapter includes:

- ▶ How This Guide Is Organized on page 12
- ▶ Who Should Read This Guide on page 12
- ▶ How Do I Find the Information That I Need? on page 12
- ▶ Additional Online Resources on page 15
- ▶ Documentation Updates on page 16

How This Guide Is Organized

The guide contains the following chapters:

Part I Creating Discovery and Integration Adapters

Describes how to create adapters.

Part II Using APIs

Describes how to work with the APIs to extract configuration data from the Run-time Service Model (RTSM).

Who Should Read This Guide

This guide is intended for the following users of HP Business Service Management:

- ▶ HP Business Service Management administrators
- ▶ HP Business Service Management platform administrators
- ▶ HP Business Service Management application administrators
- ▶ HP Business Service Management data collector administrators

Readers of this guide should be knowledgeable about enterprise system administration, have familiarity with ITIL concepts, and be knowledgeable about Business Service Management in general and the Operational Database technology specifically.

How Do I Find the Information That I Need?

This guide is part of the HP Business Service Management Documentation Library. This Documentation Library provides a single-point of access for all Business Service Management documentation.

You can access the Documentation Library by doing the following:


- ▶ In Business Service Management, select **Help > Documentation Library**.
- ▶ From a Business Service Management Gateway Server machine, select **Start > Programs > HP Business Service Management > Documentation**.


Topic Types



Within this guide, each subject area is organized into topics. A topic contains a distinct module of information for a subject. The topics are generally classified according to the type of information they contain.

This structure is designed to create easier access to specific information by dividing the documentation into the different types of information you may need at different times.

Three main topic types are in use: **Concepts**, **Tasks**, and **Reference**. The topic types are differentiated visually using icons.

Topic Type	Description	Usage
Concepts 	Background, descriptive, or conceptual information.	Learn general information about what a feature does.

Topic Type	Description	Usage
<p>Tasks</p> 	<p>Instructional Tasks. Step-by-step guidance to help you work with the application and accomplish your goals. Some task steps include examples, using sample data.</p> <p>Task steps can be with or without numbering:</p> <ul style="list-style-type: none"> ▶ Numbered steps. Tasks that are performed by following each step in consecutive order. ▶ Non-numbered steps. A list of self-contained operations that you can perform in any order. 	<ul style="list-style-type: none"> ▶ Learn about the overall workflow of a task. ▶ Follow the steps listed in a numbered task to complete a task. ▶ Perform independent operations by completing steps in a non-numbered task.
	<p>Use-case Scenario Tasks. Examples of how to perform a task for a specific situation.</p>	<p>Learn how a task could be performed in a realistic scenario.</p>

Topic Type	Description	Usage
 Reference	General Reference. Detailed lists and explanations of reference-oriented material.	Look up a specific piece of reference information relevant to a particular context.
	User Interface Reference. Specialized reference topics that describe a particular user interface in detail. Selecting Help on this page from the Help menu in the product generally open the user interface topics.	Look up specific information about what to enter or how to use one or more specific user interface elements, such as a window, dialog box, or wizard.
 Troubleshooting and Limitations	Troubleshooting and Limitations. Specialized reference topics that describe commonly encountered problems and their solutions, and list limitations of a feature or product area.	Increase your awareness of important issues before working with a feature, or if you encounter usability problems in the software.

Additional Online Resources

Troubleshooting & Knowledge Base accesses the Troubleshooting page on the HP Software Support Web site where you can search the Self-solve knowledge base. Choose **Help > Troubleshooting & Knowledge Base**. The URL for this Web site is <http://h20230.www2.hp.com/troubleshooting.jsp>.

HP Software Support accesses the HP Software Support Web site. This site enables you to browse the Self-solve knowledge base. You can also post to and search user discussion forums, submit support requests, download patches and updated documentation, and more. Choose **Help > HP Software Support**. The URL for this Web site is www.hp.com/go/hpsoftwaresupport.

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 user ID, go to:

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

HP Software Web site accesses the HP Software Web site. This site provides you with the most up-to-date information on HP Software products. This includes new software releases, seminars and trade shows, customer support, and more. Choose **Help > HP Software Web site**. The URL for this Web site is www.hp.com/go/software.

Documentation Updates

HP Software is continually updating its product documentation with new information.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to the HP Software Product Manuals Web site (<http://h20230.www2.hp.com/selfsolve/manuals>).

Part I

Creating Discovery and Integration Adapters

1

Adapter Development and Writing

This chapter includes:

Concepts

- ▶ Adapter Development and Writing Overview on page 20
- ▶ Content Creation on page 21
- ▶ Adapter Development Cycle on page 21
- ▶ Data Flow Management and Integration on page 24
- ▶ Associating Business Value with Discovery Development on page 26
- ▶ Researching Integration Requirements on page 28
- ▶ Developing Integration Content on page 31
- ▶ Developing Discovery Content on page 34

Tasks

- ▶ How to Implement a Discovery Adapter on page 37
- ▶ Step 1: Create an Adapter on page 40
- ▶ Step 2: Assign a Job to the Adapter on page 49
- ▶ Step 3: Create Jython Code on page 51

Concepts

Adapter Development and Writing Overview

Prior to beginning actual planning for development of new adapters, it is important for you to understand the processes and interactions commonly associated with this development.

The following sections can help you understand what you need to know and do, to successfully manage and execute a discovery development project.

This chapter:

- ▶ Assumes a working knowledge of Run-time Service Model and some basic familiarity with the elements of the system. It is meant to assist you in the learning process and does not provide a complete guide.
- ▶ Covers the stages of planning, research, and implementation of new discovery content for Run-time Service Model, together with guidelines and considerations that need to be taken into account.
- ▶ Provides information on the key APIs of the Data Flow Management Framework. For full documentation on the available APIs, see the *HP UCMDB API Reference*. (Other non-formal APIs exist but even though they are used on out-of-the-box adapters, they may be subject to change.)

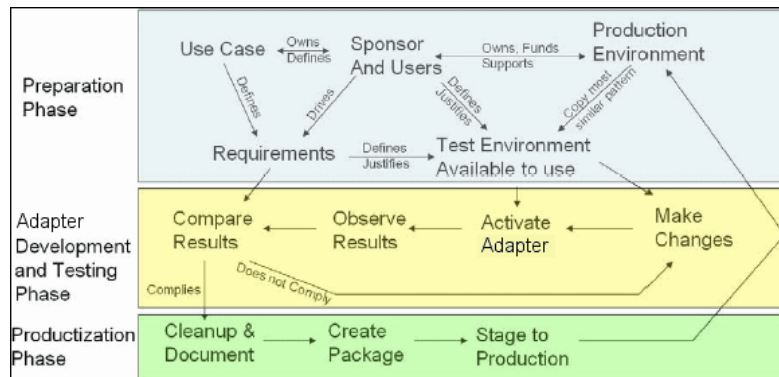
Content Creation

This section includes:

- "Adapter Development Cycle" on page 21
- "Data Flow Management and Integration" on page 24
- "Associating Business Value with Discovery Development" on page 26
- "Researching Integration Requirements" on page 28

Adapter Development Cycle

The following illustration shows a flowchart for adapter writing. Most of the time is spent in the middle section, which is the iterative loop of development and testing.




Each phase of adapter development builds on the last one.

Once you are satisfied with the way the adapter looks and works, you are ready to package it. Using either the BSM Package Manager or manual exporting of the components, create a package *.zip file. As a best practice, you should deploy and test this package on another BSM system before releasing it to production, to ensure that all the components are accounted for and successfully packaged. For details on packaging, see "Package Manager" in the *RTSM Administration Guide*.

The following sections expand on each of the phases showing the most critical steps and best practices:

- Research and Preparation Phase
- Adapter Development and Testing
- Adapter Packaging and Productization


Research and Preparation Phase



The Research and Preparation phase encompasses the driving business needs and use cases, and also accounts for securing the necessary facilities to develop and test the adapter.

- 1** When planning to modify an existing adapter, the first technical step is to make a backup of that adapter and ensure you can return it to its pristine state. If you plan to create a new adapter, copy the most similar adapter and save it under an appropriate name. For details, see "Resources Pane" in the *RTSM Data Flow Management Guide*.
- 2** Research how the adapter should collect data:
 - Use External tools/protocols to obtain the data
 - Develop how the adapter should create CIs based on the data
 - You now know what a similar adapter should look like
- 3** Determine most similar adapter based on:
 - Same CIs created
 - Same Protocols used (SNMP)
 - Same kind of targets (by OS type, versions, and so on)
- 4** Copy entire package.

- 5 Unzip into work space and rename the adapter (XML) and Jython (.py) files.



Adapter Development and Testing

The **Adapter Development and Testing** phase is a highly iterative process. As the adapter begins to take shape, you begin testing against the final use cases, make changes, test again, and repeat this process until the adapter complies with the requirements.

Startup and Preparation of Copy

- Modify XML parts of the adapter: Name (id) in line 1, Created CI Types, and Called Jython script name.
- Get the copy running with identical results to the original adapter.
- Comment out most of the code, especially the critical result-producing code.

Development and Testing

- Use other sample code to develop changes
- Test adapter by running it
- Use a dedicated view to validate complex results, search to validate simple results

Adapter Packaging and Productization

The **Adapter Packaging and Productization phase** accounts for the last phase of development. As a best practice, a final pass should be made to clean up debugging remnants, documents, and comments, to look at security considerations, and so on, before moving on to packaging. You should always have at least a readme document to explain the inner workings of the adapter. Someone (maybe even you) may need to look at this adapter in the future and will be aided greatly by even the most limited documentation.

Cleanup and Document

- ▶ Remove debugging
- ▶ Comment all functions and add some opening comments in the main section
- ▶ Create sample TQL and view for the user to test

Create Package

- ▶ Export adapters, TQL, and so on with the Package Manager. For details, see "Package Manager" in the *RTSM Administration Guide*.
- ▶ Check any dependencies your package has on other packages, for example, if the CIs created by those packages are input CIs to your adapter.
- ▶ Use Package Manager to create a package zip. For details, see "Package Manager" in the *RTSM Administration Guide*.
- ▶ Test deployment by removing parts of the new content and redeploying, or deploying on another test system.

Data Flow Management and Integration

DFM adapters are capable of integration with other products. Consider the following definitions:

- ▶ DFM collects specific content from many targets.
- ▶ Integration collects multiple types of content from one system.

Note that these definitions do not distinguish between the methods of collection. Neither does DFM. The process of developing a new adapter is the same process for developing new integration. You do the same research, make the same choices for new vs. existing adapters, write the adapters the same way, and so on. Only a few things change:

- ▶ The final adapter's scheduling. Integration adapters may run more frequently than discovery, but it depends on the use cases.
- ▶ Input CIs:
 - ▶ Integration: non-CI trigger to run with no input: a file name or source is passed through the adapter parameter.
 - ▶ Discovery: uses regular, RTSM CIs for input.

For integration projects, you should almost always reuse an existing adapter. The direction of the integration (from Run-time Service Model to another product, or from another product to Run-time Service Model) may affect your approach to development. There are field packages available for you to copy for your own uses, using proven techniques.

From Run-time Service Model to another project:

- ▶ Create a TQL that produces the CIs and relations to be exported.
- ▶ Use a generic wrapper adapter to execute the TQL and write the results to an XML file for the external product to read.

Note: For examples of field packages, contact HP Software Support.

To integrate another product to Run-time Service Model, depending on how the other product exposes its data, the integration adapter acts differently:

Integration Type	Reference Example to Be Reused
Access the product's database directly	HP ED
Read in a csv or xml file produced by an export	HP ServiceCenter
Access a product's API	BMC Atrium/Remedy

Associating Business Value with Discovery Development

The use case for developing new discovery content should be driven by a business case and plan to produce business value. That is, the goal of mapping system components to CIs and adding them to the RTSM is to provide business value.

The content may not always be used for application mapping, although this is a common intermediate step for many use cases. Regardless of the end usage of the content, your plan should answer these questions of this approach:

- ▶ Who is the consumer? How should the consumer act on the information provided by the CIs (and the relationships between them)? What is the business context in which the CIs and relationships are to be viewed? Is the consumer of these CIs a person or a product or both?
- ▶ Once the perfect combination of CIs and relationships exists in the RTSM, how do I plan on using them to produce business value?
- ▶ What should the perfect mapping look like?
 - ▶ What term would most meaningfully describe the relationships between each CI?
 - ▶ What types of CIs would be most important to include?
 - ▶ What is the end usage and end user of the map?
- ▶ What would be the perfect report layout?

Once the business justification is established, the next step is to embody the business value in a document. This means picturing the perfect map using a drawing tool and understanding the impact and dependencies between CIs, reports, how changes are tracked, what change is important, monitoring, compliance, and additional business value as required by the use cases.

This drawing (or model) is referred as the **blueprint**.

For example, if it is critical for the application to know when a certain configuration file has changed, the file should be mapped and linked to the appropriate CI (to which it relates) in the drawn map.

Work with an SME (Subject Matter Expert) of the area, who is the end user of the developed content. This expert should point out the critical entities (CIs with attributes and relationships) that must exist in the RTSM to provide business value.

One method could be to provide a questionnaire to the application owner (also the SME in this case). The owner should be able to specify the above goals and blueprint. The owner must at least provide a current architecture of the application.

You should map critical data only and no unnecessary data: you can always enhance the adapter later. The goal should be to set up a limited discovery that works and provides value. Mapping large quantities of data gives more impressive maps but can be confusing and time consuming to develop.

Once the model and business value is clear, continue to the next stage. This stage can be revisited as more concrete information is provided from the next stages.

Researching Integration Requirements

The prerequisite of this stage is a **blueprint** of the CIs and relationships needed to be discovered by DFM, which should include the attributes that are to be discovered. For details, see "Adapter Development and Writing Overview" on page 20.

This section includes the following topics:

- "Modifying an Existing Adapter" on page 28
- "Writing a New Adapter" on page 29
- "Model Research" on page 29
- "Technology Research" on page 29
- "Guidelines for Choosing Ways to Access Data" on page 30
- "Summary" on page 31

Modifying an Existing Adapter

You modify an existing adapter when an out-of-the-box or field adapter exists, but:

- it does not discover specific attributes that are needed
- a specific type of target (OS) is not being discovered or is being incorrectly discovered
- a specific relationship is not being discovered or created

If an existing adapter does some, but not all, of the job, your first approach should be to evaluate the existing adapters and verify if one of them almost does what is needed; if it does, you can modify the existing adapter.

You should also evaluate if an existing field adapter is available. Field adapters are discovery adapters that are available but are not out-of-the-box. Contact HP Software Support to receive the current list of field adapters.

Writing a New Adapter

A new adapter needs to be developed:

- When it is faster to write an adapter than to insert the information manually into the RTSM (generally, from about 50 to 100 CIs and relationships) or it is not a one-time effort.
- When the need justifies the effort.
- If out-of-the-box or field adapters are not available.
- If the results can be reused.
- When the target environment or its data is available (you cannot discover what you cannot see).

Model Research

- Browse the BSM class model (CI Type Manager) and match the entities and relations from your **blueprint** to existing CITs. It is highly recommended to adhere to the current model to avoid possible complications during version upgrade. If you need to extend the model, you should create new CITs since an upgrade may overwrite out-of-the-box CITs.
- If some entities, relations, or attributes are lacking from the current model, you should create them. It is preferable to create a package with these CITs (which will also later hold all the discovery, views, and other artifacts relating to this package) since you need to be able to deploy these CITs on each installation of Run-time Service Model.

Technology Research

Once you have verified that the RTSM holds the relevant CIs, the next stage is to decide how to retrieve this data from the relevant systems.

Retrieving data usually involves using a protocol to access a management part of the application, actual data of the application, or configuration files or databases that are related to the application. Any data source that can provide information on a system is valuable. Technology research requires both extensive knowledge of the system in question and sometimes creativity.

For home-grown applications, it may be helpful to provide a questionnaire form to the application owner. In this form the owner should list all the areas in the application that can provide information needed for the blueprint and business values. This information should include (but does not have to be limited to) management databases, configuration files, log files, management interfaces, administration programs, Web services, messages or events sent, and so on.

For off-the-shelf products, you should focus on documentation, forums, or support of the product. Look for administration guides, plug-ins and integrations guides, management guides, and so on. If data is still missing from the management interfaces, read about the configuration files of the application, registry entries, log files, NT event logs, and any artifacts of the application that control its correct operation.

Guidelines for Choosing Ways to Access Data

Relevance: Select sources or a combination of sources that provide the most data. If a single source supplies most information whereas the rest of the information is scattered or hard to access, try to assess the value of the remaining information by comparison with the effort or risk of getting it. Sometimes you may decide to reduce the blueprint if the value or cost does not warrant the invested effort.

Reuse: If Run-time Service Model already includes a specific connection protocol support it is a good reason to use it. It means the DFM Framework is able to supply a ready made client and configuration for the connection. Otherwise, you may need to invest in infrastructure development. You can view the currently supported Run-time Service Model connection protocols: **Data Flow Management > Data Flow Probe Setup > Domains and Probes pane**. For details, see "Domains and Probes Pane" in the *RTSM Data Flow Management Guide*.

You can add new protocols by adding new CIs to the model. For details, contact HP Software Support.

Note: To access Windows Registry data, you can use either WMI or NTCmd.

Security: Access to information usually requires credentials (user name, password), which are entered in the RTSM and are kept secure throughout the product. If possible, and if adding security does not conflict with other principles you have set, choose the least sensitive credential or protocol that still answers access needs. For example, if information is available both through JMX (standard administration interface, limited) and Telnet, it is preferable to use JMX since it inherently provides limited access and (usually) no access to the underlying platform.

Comfort: Some management interfaces may include more advanced features. For example, it might be easier to issue queries (SQL, WMI) than to navigate information trees or build regular expressions for parsing.

Developer Audience: The people who will eventually develop adapters may have an inclination towards a certain technology. This can also be considered if two technologies provide almost the same information at an equal cost in other factors.

Summary

The outcome of this stage is a document describing the access methods and the relevant information that can be extracted from each method. The document should also contain a mapping from each source to each relevant blueprint data.

Each access method should be marked according to the above instructions. Finally you should now have a plan of which sources to discover and what information to extract from each source into the blueprint model (which should by now have been mapped to the corresponding BSM model).

Developing Integration Content

Before creating a new integration, you must understand what the integration's requirements are:

- Should the integration copy data into the RTSM? Should the data be tracked by history? Is the source unreliable?

Population is needed.

- ▶ Should the integration federate data on the fly for views and TQL queries? Is the accuracy of changes to data critical? Is the amount of data too large to copy to the RTSM, but the requested amount of data is usually small?

Federation is needed.

- ▶ Should the integration push data in to remote data sources?

Data Push is needed.

Note: Federation and Population flows may be configured for the same integration, for maximum flexibility.

For details about the different types of integrations, see "Integration Studio" in the *RTSM Data Flow Management Guide*.

Four different options are available for creating integration adapters:

▶ **Jython Adapter**

- ▶ The classic discovery pattern
- ▶ Written in Jython
- ▶ Used for population

For details, see "Developing Jython Adapters" on page 63.

▶ **Java Adapter**

- ▶ An adapter that implements one of the adapter interfaces in the Federation SDK Framework.
- ▶ May be used for one or more of Federation, Population, or Data Push (depending on the required implementation).
- ▶ Written from scratch in Java, which allows writing code that will connect to any possible source or target.
- ▶ Suitable for jobs that each connect a single data source or target.

For details, see "Developing Java Adapters" on page 207.

► Generic DB Adapter

- An abstract adapter based on the Java Adapter and uses the Federation SDK Framework).
- Allows creation of adapters that connect to external data repositories.
- Supports both Federation and Population (with a Java plugin implemented for changes support).
- Relatively easy to define, as it is based mainly on XML and property configuration files.
- Main configuration is based on an **orm.xml** file that maps between BSM classes and database columns.
- Suitable for jobs that each connect a single data source.

For details, see "Developing Generic Database Adapters" on page 117.

► Generic Push Adapter

- An abstract adapter based on the Java Adapter (the Federation SDK Framework) and the Jython Adapter.
- Allows creation of adapters that push data to remote targets.
- Relatively easy to define, as you need only to define the mapping between BSM classes and XML, and a Jython script that pushes the data to the target.
- Suitable for jobs that each connect a single data target.
- Used for Data Push.

For details, see "Developing Push Adapters" on page 251.

The following table displays the capabilities of each adapter:

Flow/Adapter	Jython Adapter	Java Adapter	GDB Adapter	Push Adapter
Population	X	X	X	
Federation		X	X	
Data Push		X		X

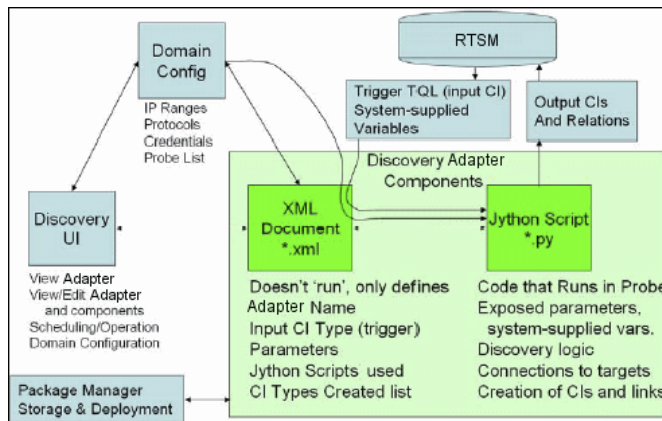
Developing Discovery Content

This section includes:

- "Discovery Adapters and Related Components" on page 34
- "Separating Adapters" on page 35

Discovery Adapters and Related Components

The following diagram shows an adapter's components and the components they interact with to execute discovery. The components in green are the actual adapters, and the components in blue are components that interact with adapters.



Note that the minimum notion of an adapter is two files: an XML document and a Jython script. The Discovery Framework, including input CIs, credentials, and user-supplied libraries, is exposed to the adapter at run time. Both discovery adapter components are administered through Data Flow Management. They are stored operationally in the RTSM itself; although the external package remains, it is not referred to for operation. The Package Manager enables preservation of the new discovery and integration content capability.

Input CIs to the adapter are provided by a TQL, and are exposed to the adapter script in system-supplied variables. Adapter parameters are also supplied as destination data, so you can configure the adapter's operation according to an adapter's specific function.

The DFM application is used to create and test new adapters. You use the Discovery Control Panel, Adapter Management, and Data Flow Probe Setup pages during adapter writing.

Adapters are stored and transported as packages. The Package Manager application and the JMX console are used to create packages from newly created adapters, and to deploy adapters on new systems.

Separating Adapters

Technically, an entire discovery could be defined in a single adapter. But good design demands that a complex system be separated into simpler, more manageable components.

The following are guidelines and best practices for dividing the adapter process:

- ▶ Discovery should be done in stages. Each stage should be represented by an adapter that should map an area or tier of the system. Adapters should rely on the previous stage or tier to be discovered, to continue discovery of the system. For example, Adapter A is triggered by an application server TQL result and maps the application server tier. As part of this mapping, a JDBC connection component is mapped. Adapter B registers a JDBC connection component as a trigger TQL and uses the results of adapter A to access the database tier (for example, through the JDBC URL attribute) and maps the database tier.
- ▶ **The two-phase connect paradigm:** Most systems require credentials to access their data. This means that a user/password combination needs to be tried against these systems. The DFM administrator supplies credentials information in a secure way to the system and can give several, prioritized login credentials. This is referred to as the **Protocol Dictionary**. If the system is not accessible (for whatever reason) there is no point in performing further discovery. If the connection is successful, there needs to be a way to indicate which credential set was successfully used, for future discovery access.

These two phases lead to a separation of the two adapters in the following cases:

- ▶ **Connection Adapter:** This is an adapter that accepts an initial trigger and looks for the existence of a remote agent on that trigger. It does so by trying all entries in the Protocol Dictionary which match this agent's type. If successful, this adapter provides as its result a remote agent CI (SNMP, WMI, and so on), which also points to the correct entry in the Protocol Dictionary for future connections. This agent CI is then part of a trigger for the content adapter.
- ▶ **Content Adapter:** This adapter's precondition is the successful connection of the previous adapter (preconditions specified by the TQLs). These types of adapters no longer need to look through all of the Protocol Dictionary since they have a way to obtain the correct credentials from the remote agent CI and use them to log in to the discovered system.
- ▶ Different scheduling considerations can also influence discovery division. For example, a system may only be queried during off hours, so even though it would make sense to join the adapter to the same adapter discovering another system, the different schedules mean that you need to create two adapters.
- ▶ Discovery of different management interfaces or technologies to discover the same system should be placed in separate adapters. This is so that you can activate the access method appropriate for each system or organization. For example, some organizations have WMI access to machines but do not have SNMP agents installed on them.

Tasks

How to Implement a Discovery Adapter

A DFM task has the aim of accessing remote (or local) systems, modeling extracted data as CIs, and saving the CIs to the RTSM. The task consists of the following steps:

1 Create a DFM adapter.

You configure an adapter file that holds the context, parameters, and result types by selecting the scripts that are to be part of the adapter. For details, see the following section.

2 Create a Discovery job.

You configure a job with scheduling information and a trigger TQL. For details, see "Step 2: Assign a Job to the Adapter" on page 49.

3 Edit Discovery code.

You can edit the Jython or Java code that is contained in the adapter files and that refers to the DFM Framework. For details, see "Step 3: Create Jython Code" on page 51.


To write new adapters, you create each of the above components, each one of which is automatically bound to the component in the previous step. For example, once you create a job and select the relevant adapter, the adapter file binds to the job.

Adapter Code


The actual implementation of connecting to the remote system, querying its data, and mapping it as RTSM data is performed by the Jython code. For example, the code contains the logic for connecting to a database and extracting data from it. In this case, the code expects to receive a JDBC URL, a user name, a password, a port, and so on. These parameters are specific for each instance of the database that answers the TQL query. You define these variables in the adapter (in the Trigger CI data) and when the job runs, these specific details are passed to the code for execution.

The adapter can refer to this code by a Java class name or a Jython script name. In this section we discuss writing DFM code as Jython scripts.


An adapter can contain a list of scripts to be used when running discovery. When creating a new adapter, you usually create a new script and assign it to the adapter. A new script includes basic templates, but you can use one of the other scripts as a template by right-clicking it and selecting **Save as**:



For details on writing new Jython scripts, see "Step 3: Create Jython Code" on page 51. You add scripts through the Resources pane:



The list of scripts are run one after the other, in the order in which they are defined in the adapter:



Note: A script must be specified even though it is being used solely as a library by another script. In this case, the library script must be defined before the script using it. In this example, the `processdbutils.py` script is a library used by the last `host_processes.py` script. Libraries are distinguished from regular runnable scripts by the lack of the `DiscoveryMain()` function.

Step 1: Create an Adapter

An adapter can be considered as the definition of a function. This function defines an input definition, runs logic on the input, defines the output, and provides a result.

Each adapter specifies input and output: Both input and output are Trigger CIs that are specifically defined in the adapter. The adapter extracts data from the input Trigger CI and passes this data as parameters to the code. (Data from related CIs is sometimes passed to the code too. For details, see "Related CIs Window" in the *RTSM Data Flow Management Guide*.) An adapter's code is generic, apart from these specific input Trigger CI parameters that are passed to the code.

For details on input components, see "Trigger CIs and Trigger Queries" in the *RTSM Data Flow Management Guide*.

This section includes the following topics:

- ▶ "Define Adapter Input (Trigger CIT and Input Query)" on page 41
- ▶ "Define Adapter Output" on page 46
- ▶ "Override Adapter Parameters" on page 48

Define Adapter Input (Trigger CIT and Input Query)

You use the Trigger CIT and Input Query components to define specific CIs as adapter input:

- ▶ The Trigger CIT defines which CIT is used as the input for the adapter. For example, for an adapter that is going to discover IPs, the input CIT is Network.
- ▶ The Input query is a regular, editable query that defines the query against the RTSM. The Input Query defines additional constraints on the CIT (for example, if the task requires a `hostID` or `application_ip` attribute), and can define more CI data, if needed by the adapter.

If the adapter requires additional information from the CIs that are related to the Trigger CI, you can add additional nodes to the input TQL. For details, see "Example of Input Query Definition" on page 43 and "How to Add Query Nodes and Relationships to a TQL Query" in the *RTSM Modeling Guide*.

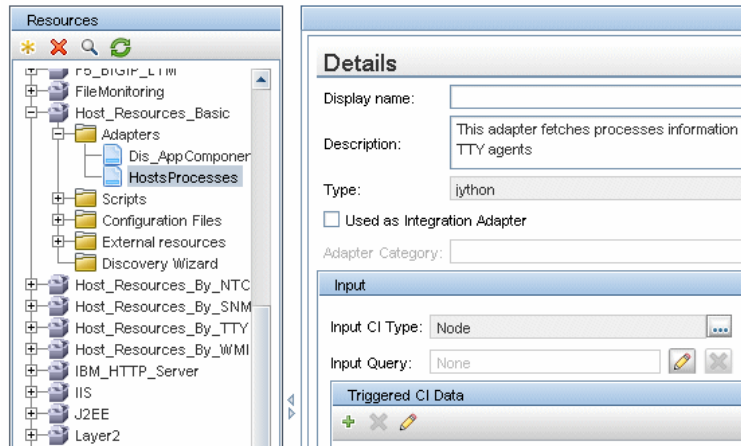
- ▶ The Trigger CI data contains all the required information on the Trigger CI as well as information from the other nodes in the Input TQL, if they are defined. DFM uses variables to retrieve data from the CIs. When the task is downloaded to the Probe, the Trigger CI data variables are replaced with actual values that exist on the attributes for real CI instances.

Example of Trigger CIT Definition:

In this example, a Trigger CIT defines that IP CIs are permitted in the adapter.

- 1** Access **Admin > RTSM Administration > Data Flow Management > Adapter Management**. Select the HostProcesses adapter (**Packages > Host_Resources_Basic > Adapters > HostProcesses**).
- 2** Locate the Input CI Type box. For details, see "Triggered CI Data" in the *RTSM Data Flow Management Guide*.
- 3** Click the button to open the Choose Discovered Class dialog box. For details, see "Choose Discovered Class Dialog Box" in the *RTSM Data Flow Management Guide*.
- 4** Select the CIT.


In this example, the IP CI (Host) is permitted in the adapter:



Example of Input Query Definition

In this example, the Input TQL query defines that the IpAddress CI (configured in the previous example as the Trigger CIT) must be connected to a Node CI.

- 1** Access **Admin > RTSM Administration > Data Flow Management > Adapter Management**. Locate the Input TQL box. Click the **Edit** button to open the Input TQL Editor. For details, see "Input Query Editor Window" in the *RTSM Data Flow Management Guide*.
- 2** In the Input TQL Editor, name the Trigger CI node **SOURCE**: right-click the node and choose **Query Node Properties**. In the **Element Name** box, change the name to **SOURCE**.
- 3** Add a Node CI and a Containment relationship to the IpAddress CI. For details on working with the Input TQL Editor, see "Input Query Editor Window" in the *RTSM Data Flow Management Guide*.




The **IpAddress** CI is connected to a **Node** CI. The input TQL consists of two nodes, **Node** and **IpAddress**, with a link between them. The **IpAddress** CI is named **SOURCE**.

Example of Adding Variables to the Input TQL Query:

In this example, you add `DIRECTORY` and `CONFIGURATION_FILE` variables to the Input TQL query created in the previous example. These variables help to define what must be discovered, in this case, to find the configuration files residing on the hosts that are linked to the IPs you need to discover.

- 1 Display the Input TQL created in the previous example.
Access **Admin > RTSM Administration > Data Flow Management > Adapter Management**. Locate the Triggered CI Data pane. For details, see "Triggered CI Data" in the *RTSM Data Flow Management Guide*.
- 2 Add variables to the Input TQL. For details, access **Admin > RTSM Administration > Data Flow Management > Adapter Management**. Locate the Triggered CI Data pane. For details, see the Variables field in "Triggered CI Data" in the *RTSM Data Flow Management Guide*.



Example of Replacing Variables with Actual Data:

In this example, variables replace the `IpAddress` CI data with actual values that exist on real `IpAddress` CI instances in your system.


The Triggered CI data for the `IpAddress` CI includes a `fileName` variable. This variable enables the replacement of the `CONFIGURATION_DOCUMENT` node in the Input TQL with the actual values of the configuration file located on a host:

Triggered CI data	
Name	Value
Protocol	\${SOURCE.credentials_id}
credentialsId	\${SOURCE.credentials_id}
fileName	\${CONFIGURATION_DOCUMENT.name}
hostID	\${HOST.root_id}
ip_address	\${SOURCE.ip_address}
path	\${CONFIGURATION_DOCUMENT.resource_path}

The Trigger CI data is uploaded to the Probe with all variables replaced by actual values. The adapter script includes a command to use the DFM Framework to retrieve the actual values of the defined variables:

```
Framework.getTriggerCIData ('ip_address')
```

The fileName and path variables use the data_name and document_path attributes from the **CONFIGURATION_DOCUMENT** node (defined in the Input Query Editor – see previous example).

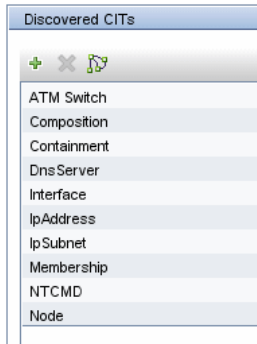


The Protocol, credentialsId, and ip_address variables use the root_class, credentials_id, and application_ip attributes:


Key	Name	Display Name	Type	Description	Default Value	Visible
	ack_cleared_time	ack_cleared_time	long			
	ack_id	ack_id	string			
	BODY_ICON	BODY_ICON	string		host	
	city	City	string	City location		✓
	codepage	CodePage	string	System su...		
	contextmenu	Context Menu	string_list	Context me...	itCIs	
	country	Country	string	Country loc...		✓
	credentials_id	Reference to the cre...	string	Reference ...		
	data_adminstate	Admin State	adminstate...	Admin State	Managed	

Define Adapter Output

The output of the adapter is a list of discovered CIs (**Admin > RTSM Administration > Data Flow Management > Adapter Management > Adapter Definition tab > Discovered CITs**) and the links between them:



You can also view the CITs as a topology map, that is, the components and the way in which they are linked together (click the **View Discovered CITs as Map** button):



The discovered CIs are returned by the DFM code (that is, the Jython script) in the format of BSM's `ObjectStateHolderVector`. For details, see "Results Generation by the Jython Script" on page 71.

Example of Adapter Output:

In this example, you define which CITs are to be part of the IP CI output.

- 1** Access **Admin > RTSM Administration > Data Flow Management > Adapter Management**.
- 2** In the Resources pane, select **Network > Adapters > NSLOOKUP_on_Probe**.
- 3** In the Adapter Definition tab, locate the Discovered CITs pane.
- 4** The CITs that are to be part of the adapter output are listed. Add CITs to, or remove from, the list. For details, see "Discovered CITs Pane" in the *RTSM Data Flow Management Guide*.

Override Adapter Parameters

To configure an adapter for more than one job, you can override adapter parameters. For example, the adapter `SQL_NET_Dis_Connection` is used by both the MSSQL Connection by SQL and the Oracle Connection by SQL jobs.

Example of Overriding an Adapter Parameter:

This example illustrates overriding an adapter parameter so that one adapter can be used to discover both Microsoft SQL Server and Oracle databases.

- 1** Access **Admin > RTSM Administration > Data Flow Management > Adapter Management**.
- 2** In the Resources pane, select **Database Basic > Adapters > SQL_NET_Dis_Connection**.
- 3** In the Adapter Definition tab, locate the **Discovery Pattern Parameters** pane. The `protocolType` parameter has a value of **all**:

Name	Value
protocolType	all

- 4** Right-click the `SQL_NET_Dis_Connection_MsSql` adapter and choose **Go to Discovery Job > MSSQL Connection by SQL**.
- 5** Display the Properties tab. Locate the Parameters pane:

Override	Name	Value
<input checked="" type="checkbox"/>	protocolType	MicrosoftSQLServer

The `all` value is overwritten with the `MicrosoftSQLServer` value.

Note: The **Oracle Connection by SQL** job includes the same parameter but the value is overwritten with an Oracle value.


For details on adding, deleting, or editing parameters, see "Adapter Parameters Pane" in the *RTSM Data Flow Management Guide*.

DFM begins looking for Microsoft SQL Server instances according to this parameter.

Step 2: Assign a Job to the Adapter

Each adapter has one or more associated jobs that define the execution policy. Jobs enable scheduling the same adapter differently over different set of Triggered CIs and also enable supplying different parameters for each set.

The jobs appear in the Discovery Modules tree, and this is the entity that the user activates.



The screenshot shows the Discovery Modules tree on the left and the configuration window on the right. The tree includes categories like Cluster and Load Balancing Solutions, Custom Jobs, Database, and Network Discovery. The configuration window has tabs for Details, Properties, and Dependencies. It contains a Parameters table and a Trigger Queries table.

Override	Name	Value
<input checked="" type="checkbox"/>	discoverDisks	true
<input checked="" type="checkbox"/>	discoverInstalledSoftware	false
<input checked="" type="checkbox"/>	discoverProcesses	false
<input checked="" type="checkbox"/>	discoverServices	false
<input checked="" type="checkbox"/>	discoverUsers	true

Query Name	Probe Limit
snmp	<<All Probes>>

Choose a Trigger TQL


Each job is associated with Trigger TQLs. These Trigger TQLs publish results that are used as Input Trigger CIs for the adapter of this job.

A Trigger TQL can add constraints to an Input TQL. For example, if an input TQL's results are IPs connected to SNMP, a trigger TQL's results can be IPs connected to SNMP within the range 195.0.0.0-195.0.0.10.

Note: A trigger TQL must refer to the same objects that the input TQL refers to. For example, if an input TQL queries for IPs running SNMP, you cannot define a trigger TQL (for the same job) to query for IPs connected to a host, because some of the IPs may not be connected to an SNMP object, as required by the input TQL.

Set Scheduling Information

The scheduling information for the Probe specifies when to run the code on Trigger CIs. If the **Invoke on new triggered CIs Immediately** check box is selected, the code also runs once on each Trigger CI when it reaches the Probe, regardless of future schedule settings.



For each scheduled occurrence for each job, the Probe runs the code against all Trigger CIs accumulated for that job. For details, see "Discovery Scheduler Dialog Box" in the *RTSM Data Flow Management Guide*.

Override Parameters

When configuring a job you can override the adapter parameters. For details, see "Override Adapter Parameters" on page 48.

Step 3: Create Jython Code

Run-time Service Model uses Jython scripts for adapter-writing. For example, the `SNMP_Connection.py` script is used by the `SNMP_NET_Dis_Connection` adapter to try and connect to machines using SNMP. Jython is a language based on Python and powered by Java.

For details on how to work in Jython, you can refer to these Web sites:

- <http://www.jython.org>
- <http://www.python.org>

For details, see "How to Create Jython Code" on page 65.

2

Discovery Content Migration Guidelines

This chapter includes:

Concepts

- ▶ Discovery Content Migration Guidelines Overview on page 54
- ▶ Version 9.x New Infrastructure Features on page 54
- ▶ Guidelines for Developing Cross-Data Model Scripts on page 59
- ▶ Implementation Tips on page 59

Tasks

- ▶ How to Access BTO Data Model Documentation Online on page 60

Reference

- ▶ Package Migration Utility on page 61

Troubleshooting and Limitations on page 61

Concepts

Discovery Content Migration Guidelines Overview

In Business Service Management version 9.x, the data model has significantly evolved, forcing correlated changes in the former Discovery and Dependency Mapping (DDM) content code. Consequently, some core mechanisms of the DDM content have changed. Thus, content developed for BSM prior to version 9.x has to be upgraded to correspond with the 9.x data model (BDM: BTO Data Model). This section guides you through the process of adopting DDM content and aligning it with BDM.

Version 9.x New Infrastructure Features

Note: For details on accessing the BDM documentation online, see "How to Access BTO Data Model Documentation Online" on page 60.

This section includes:

- "BTO Data Model (BDM)" on page 55
- "Differences Between BSM 8.0x Class Model and BSM 9.x Data Model" on page 55
- "New CIT Identification Mechanism" on page 55
- "Running Software Mechanism" on page 56
- "Probe Side Identification" on page 57
- "Transformation Layer" on page 58

BTO Data Model (BDM)

- For details on the BTO Data Model (BDM) see the Conceptual Data Model document. This document is a map of the concepts being modeled, as well as the scope of the model. This conceptual data model provides a starting point for understanding the semantics of the modeled domain.
- For details on the BDM classes, see the HP Software BTO Data Model Reference document. This document covers all BDM classes, including class description and attribute, qualifier, and hierarchy information.

Differences Between BSM 8.0x Class Model and BSM 9.x Data Model

Changes made between the BSM version 8.0x class model and BDM are downloaded to the Probe in the following Discovery configuration file:

C:\hp\UCMDB\DataFlowProbe\runtime\probeManager\discoveryConfigFiles\flat-class-model-changes.xml.

bdm_changes.xml. This XML file holds information regarding changes made to class names, attributes names, removed classes, attributes, qualifiers, and so on.

- For details on mapping between the BSM version 8.0x class model and BDM, see the Mapping of BSM 9.x (BTO Data Model) to BSM 8.0x Class Model document.
- For details on changes to the class model between version 8.0x and 9.x, see the BSM Class Model Changes Report document.

New CIT Identification Mechanism

In BSM versions prior to version 9.x, key attributes are used to identify CIs. In BSM version 9.x, this concept has been generalized and the identification is now done in a server component named Reconciliation Engine. The Reconciliation Engine is capable of identifying CIs by logical rules called DDA (Data Definition Algorithm) rules.

This new mechanism is mostly useful for CITs where the related topology is important for their identification (for example, the Node CIT—Host in prior versions—is identified by its name and the related topology, such as the IP Address and Interface CITs). Some CITs are still identified by key attributes; for those CITs, a DDA rule is not defined.

For details about the Reconciliation Engine, see "Reconciliation Overview" in the *RTSM Data Flow Management Guide*.

Running Software Mechanism

The version 8.0x **Software Element** CI is called **Running Software** in version 9.x BDM. This CIT is identified in version 9.x by a DDA rule and not by key attributes.

Say you have added a custom CIT derived from the **Running Software** CIT. In previous versions this custom CIT was identified by its key attributes. However, in version 9.x it is identified by an inherited DDA rule, and thus defined key attributes are ignored.

So if you add a derived CIT, consider the following:

- ▶ To identify the new CIT by the same DDA rule as all the Running Software CITs, you should keep the current configuration.
- ▶ To identify the new CIT by key attributes, you should create a new DDA rule, defining the identification by key attributes. Following is an example for such a DDA rule, defined for the **object** CIT:

```
<identification-config type="object">
  <identification-criteria>
    <identification-criterion targetType="root">
      <key-attributes-condition/>
    </identification-criterion>
  </identification-criteria>
</identification-config>
```


Probe Side Identification

DDM_ID_ATTRIBUTE. The version 9.x Data Flow Probe identifies CIs only by their key attributes (that is, **ID_ATTRIBUTE**). If a CIT includes a DDA rule (that is, a reconciliation rule), the CIT may not include a key attribute. In this case, the CIT main attributes are marked with a **DDM_ID_ATTRIBUTE** qualifier. Therefore, for the purposes of identifying a CI, the Probe considers all **DDM_ID_ATTRIBUTE** as well as **ID_ATTRIBUTE** qualifiers.

DDM_REQUIRED_TOPOLOGY. A DDA rule for a specific CIT may depend on different CIs reported in the same bulk, together with the examined CI. For example, J2EE Domain CIT identification is carried out not only by the domain name attribute but also by the J2EE Application Server CIT connected to it with a membership link.

To ensure that all the required CIs are reported with the examined CI, you should mark each one of the examined CIs with the **DDM_REQUIRED_TOPOLOGY** qualifier that contains a data item specifying the required link type. For example, in the above example, the J2EE Domain CIT is marked with the **DDM_REQUIRED_TOPOLOGY** qualifier and with a member link data item, so that when Discovery reports a J2EE domain, the servers are also reported. Data item name which specifies link types is **LINK_TYPES**.

As an example, to identify the Node CIT by interfaces and IPs connected to it, then the following qualifier should be added to the Node CIT definition:

```
<Class-Qualifier name="DDM_REQUIRED_TOPOLOGY">
  <Data-Items>
    <Data-Item name="LINK_TYPES" type="string">containment,composition</Data-Item>
    <Data-Item name="LINK_ENDS" type="string">ip_address,interface</Data-Item>
    <Data-Item name="LINK_DIRECTIONS" type="string">OUT,OUT</Data-Item>
    <Data-Item name="APPLY_TO_CHILD_TYPES" type="string">>true</Data-Item>
  </Data-Items>
</Class-Qualifier>
```

where

- **LINK_TYPES** (mandatory) indicates the link types of the current CIT topology.

- ▶ **LINK_ENDS** (optional) provides definitions for the CITs on the opposite ends of the specified link types according to their appearance in the **LINK_TYPES** list. These “opposite” ends are always applied hierarchically.

Omitting the **LINK_ENDS** data item or leaving one end as an empty string in the list means that the opposite end can be of any CIT.

- ▶ **LINK_DIRECTION** (optional) indicates the link direction, “OUT”, “IN” or “BOTH”, to check from the current CIT.

Omitting the **LINK_DIRECTION** data item or leaving an empty entry in the list means that both directions are checked.

- ▶ **APPLY_TO_CHILD_TYPES** (optional) indicates that the qualifier will be applied recursively to all children of the current CIT.

Omitting the **APPLY_TO_CHILD_TYPES** data item means that the qualifier is applied only to the current CIT.

If **DDM_REQUIRED_TOPOLOGY** is defined for a specific CIT, this will override the qualifier defined for its parent.

For details on qualifiers, see "Qualifiers Page" in the *RTSM Modeling Guide*.

Transformation Layer

To ensure backward compatibility, a new transformation mechanism is introduced in version 9.x on the Probe. The new mechanism is capable of converting version 8.0x topologies to 9.x topologies at runtime. It enables the Probe to continue running tasks, such as Jython scripts, which report topologies compatible with version 8.0x.

The new transformation mechanism uses the data kept in the **bdm_changes.xml** file, and performs the required changes (class and attributes name changes, attribute removal, hierarchy changes, and so on) to make the 8.0x topologies compatible with the BDM. Concurrently (and independently of the topologies reported by the tasks executed by the Probe), the BSM Server receives topologies compatible with BDM.

Guidelines for Developing Cross-Data Model Scripts

The following guidelines are applicable for both version 8.0x and 9.x.

Discovery Scripts API Library

The Discovery API library is fully backward compatible and therefore all version 8.0x libraries and APIs are supported. For details, see "Jython Libraries and Utilities" on page 112.

The 9.x API includes more elements and methods. For example, a Jython script now reports an error code (integer) instead of a string error message, thus enabling localized discovery error messages. For details, see "Error-Writing Conventions" on page 119.

Implementation Tips

- Use the **modeling** module for creating a **Running Software** CIT or any descendant for which the relevant method is present.
- Use **HostBuilder** for creating CIT of type **Node**.
- Use the **modeling.createOshByCmdbldString** to restore OSH by its ID.
- Use the **ShellUtils** instance of the **shellutils** module for all shell-based connections.
- Use the built-in mechanism to retrieve the BSM version: `logger.Version().getVersion(framework)`. For example, if an additional attribute `application_ip` is added only for BSM version 9.x or later:

```
versionAsDouble = logger.Version().getVersion(Framework)
if versionAsDouble >= 9:
    appServerOSH.setAttribute('application_ip', ip)
```

- Use **wmiutils** for creating a WMI-based discovery.
- Use **snmputils** for creating a SNMP-based discovery.

Tasks

How to Access BTO Data Model Documentation Online

To access the BDM documentation:

- 1 Log on to Business Service Management.
- 2 Click **Help > Documentation Library**.
- 3 Click the **Modeling** link under **System Administration** to access the **Modeling** portal.
- 4 Click the **Class Model** tab.

Reference

Package Migration Utility

The BSM 9.x installation includes an external Package Migration Utility that enables content developers to convert a content package from the 8.0x class model to the 9.x data model. The Package Migration Utility converts package resources, subsystem by subsystem, so that they are compatible with the new class model. CIT definitions, queries, jobs, adapters, and modules are transformed according to the data held in the **bdm_changes.xml** file. As a result, they can be deployed and used by a BSM 9.x Server.

Package Migration Utility Limitations

- ▶ Jython scripts are not upgraded by the Package Migration Utility. For supporting scripts that are designed to correspond with the BSM version 8.0x class model, a new **Transformation layer** module is introduced in BSM 9.x. For details, see "Transformation Layer" on page 58.
- ▶ Discovery Adapters of type Integration are not upgraded by the Package Migration Utility and thus should be upgraded manually.
- ▶ The Layer 2 Topology discovery job (and its corresponding resources, such as Discovery Adapter, TQL, and so on) has significantly changed and is removed by the Package Migration Utility instead of being upgraded.

Troubleshooting and Limitations

- ▶ The **ip_address** value is not passed by default to the pattern. It should be added explicitly to the pattern as Trigger CI Data.
- ▶ If a non-out-of-the-box Jython script requires an external jar or resource in the classpath, it should be located in the relevant package under a sub-folder named **discoveryResources**.

- ▶ While working with attributes of type **List** such as **StringVector** and **IntegerVector** (inherited from **BaseVector**), you cannot use both the **add element** and **remove element** operations on the same list object.

3

Developing Jython Adapters

This chapter includes:

Concepts

- ▶ HP Data Flow Management API Reference on page 64

Tasks

- ▶ How to Create Jython Code on page 65
- ▶ How to Support Localization in Jython Adapters on page 79
- ▶ How to Work with Discovery Analyzer on page 90
- ▶ How to Run Discovery Analyzer from Eclipse on page 99
- ▶ How to Record DFM Code on page 109

Reference

- ▶ Jython Libraries and Utilities on page 112

Concepts

HP Data Flow Management API Reference

For full documentation on the available APIs, see *HP Universal CMDB Data Flow Management API Reference*. These files are located in the following folder:

```
\\<BSM Gateway Server root directory>\AppServer\  
webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\DDM_JavaDoc\ind  
ex.html
```

Tasks

How to Create Jython Code

Run-time Service Model uses Jython scripts for adapter-writing. For example, the `SNMP_Connection.py` script is used by the `SNMP_NET_Dis_Connection` adapter to try and connect to machines using SNMP. Jython is a language based on Python and powered by Java.

For details on how to work in Jython, you can refer to these Web sites:

- <http://www.jython.org>
- <http://www.python.org>

The following section describes the actual writing of Jython code inside the DFM Framework. This section specifically addresses those contact points between the Jython script and the Framework that it calls, and also describes the Jython libraries and utilities that should be used whenever possible.

Note:

- Scripts written for DFM should be compatible with Jython version 2.1.
 - For full documentation on the available APIs, see the *HP Universal CMDB Data Flow Management API Reference*.
-

This section includes the following topics:

- "Use External Java JAR Files within Jython" on page 66
- "Execution of the Code" on page 66
- "Modifying Out-of-the-Box Scripts" on page 66
- "Structure of the Jython File" on page 68
- "Results Generation by the Jython Script" on page 71
- "Framework Instance" on page 73

- ▶ "Finding the Correct Credentials (for Connection Adapters)" on page 77
- ▶ "Handling Exceptions from Java" on page 79

Use External Java JAR Files within Jython

When developing new Jython scripts, external Java Libraries (JAR files) or third-party executable files are sometimes needed as either Java utility archives, connection archives such as JDBC Driver JAR files, or executable files (for example, **nmap.exe** is used for credential-less discovery).

These resources should be bundled in the package under the **External Resources** folder. Any resource put in this folder is automatically sent to any Probe that connects to your RTSM server.

In addition, when discovery is launched, any JAR file resource is loaded into the Jython's classpath, making all the classes within it available for import and use.

Execution of the Code

After a job is activated, a task with all the required information is downloaded to the Probe.

The Probe starts running the DFM code using the information specified in the task.

The Jython code flow starts running from a main entry in the script, executes code to discover CIs, and provides results of a vector of discovered CIs.

Modifying Out-of-the-Box Scripts

When making out-of-the-box script modifications, make only minimal changes to the script and place any necessary methods in an external script. You can track changes more efficiently and, when moving to a newer Runtime Service Model version, your code is not overwritten.

For example, the following single line of code in an out-of-the-box script calls a method that calculates a Web server name in an application-specific way:

```
serverName = iplanet_cspecific.PlugInProcessing(serverName, transportHN,
mam_utils)
```

The more complex logic that decides how to calculate this name is contained in an external script:

```
# implement customer specific processing for 'servername' attribute of httpplugin
#
def PlugInProcessing(servername, transportHN, mam_utils_handle):
    # support application-specific HTTP plug-in naming
    if servername == "appsrv_instance":
        # servername is supposed to match up with the j2ee server name,
        however some groups do strange things with their
        # iPlanet plug-in files. this is the best work-around we could find. this join
        can't be done with IP address:port
        # because multiple apps on a web server share the same IP:port for
        multiple websphere applications
        logger.debug('httpcontext_webapplicationserver attribute has been
        changed from [' + servername + '] to [' + transportHN[:5] + '] to facilitate websphere
        enrichment')
        servername = transportHN[:5]
    return servername
```

Save the external script in the External Resources folder. For details, see "Resources Pane" in the *RTSM Data Flow Management Guide*. If you add this script to a package, you can use this script for other jobs, too. For details on working with Package Manager, see "Package Manager" in the *RTSM Administration Guide*.

During upgrade, the change you make to the single line of code is overwritten by the new version of the out-of-the-box script, so you will need to replace the line. However, the external script is not overwritten.

Structure of the Jython File

The Jython file is composed of three parts in a specific order:

- 1 Imports
- 2 Main Function - DiscoveryMain
- 3 Functions definitions (optional)

The following is an example of a Jython script:

```
# imports section
from appilog.common.system.types import ObjectStateHolder
from appilog.common.system.types.vectors import ObjectStateHolderVector

# Function definition
def foo:
    # do something

# Main Function
def DiscoveryMain(Framework):
    OSHVResult = ObjectStateHolderVector()

    ## Write implementation to return new result CIs here...

    return OSHVResult
```

Imports

Jython classes are spread across hierarchical namespaces. In version 7.0 or later, unlike in previous versions, there are no implicit imports, and so every class you use must be imported explicitly. (This change was made for performance reasons and to enable an easier understanding of the Jython script by not hiding necessary details.)

- To import a Jython script:

```
import logger
```

- To import a Java class:

```
from appilog.collectors.clients import ClientsConsts
```

Main Function – DiscoveryMain

Each Jython runnable script file contains a main function: `DiscoveryMain`.

The `DiscoveryMain` function is the main entry into the script; it is the first function that runs. The main function may call other functions that are defined in the scripts:

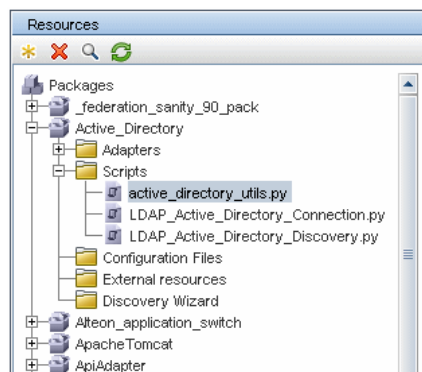
```
def DiscoveryMain(Framework):
```

The `Framework` argument must be specified in the main function definition. This argument is used by the main function to retrieve information that is required to run the scripts (such as information on the Trigger CI and parameters) and can also be used to report on errors that occur during the script run.

You can create a Jython script without any main method. Such scripts are used as library scripts that are called from other scripts.

Functions Definition

Each script can contain additional functions that are called from the main code. Each such function can call another function, which either exists in the current script or in another script (use the `import` statement). Note that to use another script, you must add it to the `Scripts` section of the package:



Example of a Function Calling Another Function:

In the following example, the main code calls the doQueryOSUsers(..) method which calls an internal method doOSUserOSH(..):

```
def doOSUserOSH(name):
    sw_obj = ObjectStateHolder('winosuser')

    sw_obj.setAttribute('data_name', name)
    # return the object
    return sw_obj

def doQueryOSUsers(client, OSHVResult):
    _hostObj = modeling.createHostOSH(client.getIpAddress())
    data_name_mib = '1.3.6.1.4.1.77.1.2.25.1.1,1.3.6.1.4.1.77.1.2.25.1.2,string'
    resultSet = client.executeQuery(data_name_mib)
    while resultSet.next():
        UserName = resultSet.getString(2)
        ##### send object #####
        OSUserOSH = doOSUserOSH(UserName)
        OSUserOSH.setContainer(_hostObj)
        OSHVResult.add(OSUserOSH)

def DiscoveryMain(Framework):
    OSHVResult = ObjectStateHolderVector()
    try:
        client =
Framework.getClientFactory(ClientsConsts.SNMP_PROTOCOL_NAME).createClient()
    except:
        Framework.reportError('Connection failed')
    else:
        doQueryOSUsers(client, OSHVResult)
        client.close()
    return OSHVResult
```

If this script is a global library that is relevant to many adapters, you can add it to the list of scripts in the jythonGlobalLibs.xml configuration file, instead of adding it to each adapter (**Admin > RTSM Administration > Adapter Management > Resources Pane > AutoDiscoveryContent > Configuration Files**).

Results Generation by the Jython Script

Each Jython script runs on a specific Trigger CI, and ends with results that are returned by the return value of the `DiscoveryMain` function.

The script result is actually a group of CIs and links that are to be inserted or updated in the RTSM. The script returns this group of CIs and links in the format of `ObjectStateHolderVector`.

The `ObjectStateHolder` class is a way to represent an object or link defined in the RTSM. The `ObjectStateHolder` object contains the CIT name and a list of attributes and their values. The `ObjectStateHolderVector` is a vector of `ObjectStateHolder` instances.

ObjectStateHolder Syntax

This task describes how to build the DFM results into a BSM model.

Example of Setting Attributes on the CIs:

The `ObjectStateHolder` class describes the DFM result graph. Each CI and link (relationship) is placed inside an instance of the `ObjectStateHolder` class as in the following Jython code sample:

```
# siebel application server
1 appServerOSH = ObjectStateHolder('siebelappserver' )
2 appServerOSH.setStringAttribute('data_name', sblsvrName)
3 appServerOSH.setStringAttribute ('application_ip', ip)
4 appServerOSH.setContainer(appServerHostOSH)
```

- Line 1 creates a CI of type **siebelappserver**.
- Line 2 creates an attribute called **data_name** with a value of **sblsvrName** which is a Jython variable set with the value discovered for the server name.
- Line 3 sets a non-key attribute that is updated in the RTSM.
- Line 4 is the building of containment (the result is a graph). It specifies that this application server is contained inside a host (another `ObjectStateHolder` class in the scope).

Note: Each CI being reported by the Jython script must include values for all the key attributes of the CI's CI Type.

Example of Relationships (Links):

The following link example explains how the graph is represented:

```
1 linkOSH = ObjectStateHolder('route')
2 linkOSH.setAttribute('link_end1', gatewayOSH)
3 linkOSH.setAttribute('link_end2', appServerOSH)
```

- ▶ Line 1 creates the link (that is also of the `ObjectStateHolder` class. The only difference is that `route` is a link CI Type).
- ▶ Lines 2 and 3 specify the nodes at the end of each link. This is done using the **end1** and **end2** attributes of the link which must be specified (because they are the minimal key attributes of each link). The attribute values are `ObjectStateHolder` instances. For details on End 1 and End 2, see "Link" in the *RTSM Data Flow Management Guide*.

Caution: A link is directional. You should verify that End 1 and End 2 nodes correspond to valid CITs at each end. If the nodes are not valid, the result object fails validation and is not reported correctly. For details, see "CI Type Relationships" in the *RTSM Modeling Guide*.

Example of Vector (Gathering CIs):

After creating objects with attributes, and links with objects at their ends, you must now group them together. You do this by adding them to an `ObjectStateHolderVector` instance, as follows:

```
oshvMyResult = ObjectStateHolderVector()
oshvMyResult.add(appServerOSH)
oshvMyResult.add(linkOSH)
```

For details on reporting this composite result to the Framework so it can be sent to the RTSM server, see the `sendObjects` method.

Once the result graph is assembled in an `ObjectStateHolderVector` instance, it must be returned to the DFM Framework to be inserted into the RTSM. This is done by returning the `ObjectStateHolderVector` instance as the result of the `DiscoveryMain()` function.

Note: For details on creating **OSH** for common CITs, see `modeling.py` in "Jython Libraries and Utilities" on page 112.

Framework Instance

The Framework instance is the only argument that is supplied in the main function in the Jython script. This is an interface that can be used to retrieve information required to run the script (for example, information on trigger CIs and adapter parameters), and is also used to report on errors that occur during the script run. For details, see "HP Data Flow Management API Reference" on page 64.

The correct usage of Framework instance is to pass it as argument to each method that uses it.

Example:

```
def DiscoveryMain(Framework):
    OSHVResult = helperMethod (Framework)
    return OSHVResult

def helperMethod (Framework):
    ....
    probe_name = Framework.getDestinationAttribute('probe_name')
    ...
    return result
```

This section describes the most important Framework usages:


- "Framework.getTriggerCIData(String attributeName)" on page 74
- "Framework.createClient(credentialsId, props)" on page 74
- "Framework.getParameter (String parameterName)" on page 76
- "Framework.reportError(String message) and Framework.reportWarning(String message)" on page 76

Framework.getTriggerCIData(String attributeName)

This API provides the intermediate step between the Trigger CI data defined in the adapter and the script.

Example of Retrieving Credential Information:

You request the following Trigger CI data information:



Name	Value
Protocol	\${SOURCE.credentials_id}
credentialsId	\${SOURCE.credentials_id}
fileName	\${CONFIGURATION_DOCUMENT.name}
hostID	\${HOST.root_id}
ip_address	\${SOURCE.ip_address}
path	\${CONFIGURATION_DOCUMENT.resource_path}

To retrieve the credential information from the task, use this API:

```
credId = Framework.getTriggerCIData('credentialsId')
```

Framework.createClient(credentialsId, props)

You make a connection to a remote machine by creating a client object and executing commands on that client. To create a client, retrieve the ClientFactory class. The getClientFactory() method receives the type of the requested client protocol. The protocol constants are defined in the ClientsConsts class. For details on credentials and supported protocols, see "Domain Credential References" in the *RTSM Data Flow Management Guide*.

Example of Creating a Client Instance for the Credentials ID:

To create a Client instance for the credentials ID:

```
properties = Properties()
codePage = Framework.getCodePage()
properties.put( BaseAgent.ENCODING, codePage)
client = Framework.createClient(credentialsID ,properties)
```

You can now use the Client instance to connect to the relevant machine or application.

Example of Creating a WMI Client and Running a WMI Query:

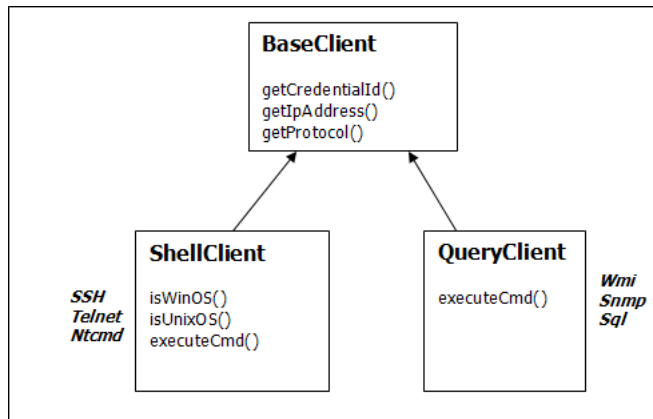
To create a WMI client and run a WMI query using the client:

```
wmiClient = Framework.createClient(credential)
resultSet = wmiClient.executeQuery("SELECT TotalPhysicalMemory
                                   FROM Win32_LogicalMemoryConfiguration")
```

Note: To make the `createClient()` API work, add the following parameter to the Trigger CI data parameters: **credentialsId = \${SOURCE.credentials_id}** in the Triggered CI Data pane. Or you can manually add the credentials ID when calling the function:

wmiClient = clientFactory().createClient(credentials_id).

The following diagram illustrates the hierarchy of the clients, with their commonly-supported APIs:



For details on the clients and their supported APIs, see `BaseClient`, `ShellClient`, and `QueryClient` in the *HP Universal CMDB Data Flow Management API Reference*.

Framework.getParameter (String parameterName)

In addition to retrieving information on the Trigger CI, you often need to retrieve an adapter parameter value. For example:

Parameters		
Override	Name	Value
<input checked="" type="checkbox"/>	protocolType	MicrosoftSQLServer

Example of Retrieving the Value of the protocolType Parameter:

To retrieve the value of the protocolType parameter from the Jython script, use the following API:

```
protocolType = Framework.getParameterValue('protocolType')
```

Framework.reportError(String message) and Framework.reportWarning(String message)

Some errors (for example, connection failure, hardware problems, timeouts) can occur during a script run. When such errors are detected, Framework can report on the problem. The message that is reported reaches the server and is displayed for the user.

Example of a Report Error and Message:

The following example illustrates the use of the reportError(<Error Msg>) API:

```
try:
    client =
Framework.getClientFactory(ClientsConsts.SNMP_PROTOCOL_NAME)
    createClient()
except:
    strException = str(sys.exc_info()[1]).strip()
    Framework.reportError ('Connection failed: %s' % strException)
```

You can use either one of the APIs—`Framework.reportError(String message)`, `Framework.reportWarning(String message)`—to report on a problem. The difference between the two APIs is that when reporting an error, the Probe saves a communication log file with the entire session's parameters to the file system. In this way you are able to track the session and better understand the error.

Finding the Correct Credentials (for Connection Adapters)

An adapter trying to connect to a remote system needs to try all possible credentials. One of the parameters needed when creating a client (through `ClientFactory`) is the credentials ID. The connection script gains access to possible credential sets and tries them one by one using the `clientFactory.getAvailableProtocols()` method. When one credential set succeeds, the adapter reports a CI connection object on the host of this trigger CI (with the credentials ID that matches the IP) to the RTSM. Subsequent adapters can use this connection object CI directly to connect to the credential set (that is, the adapters do not have to try all possible credentials again).

The following example shows how to obtain all entries of the SNMP protocol. Note that here the IP is obtained from the Trigger CI data (# Get the Trigger CI data values).

The connection script requests all possible protocol credentials (# Go over all the protocol credentials) and tries them in a loop until one succeeds (resultVector). For details, see the **two-phase connect paradigm** entry in "Separating Adapters" on page 35.

```
import logger
from appilog.collectors.clients import ClientsConsts
from appilog.common.system.types.vectors import ObjectStateHolderVector

def mainFunction(Framework):
    resultVector = ObjectStateHolderVector()

    # Get the Trigger CI data values
    ip_address = Framework.getDestinationAttribute('ip_address')
    ip_domain = Framework.getDestinationAttribute('ip_domain')

    # Create the client factory for SNMP
    clientFactory = framework.getClientFactory(ClientsConsts.SNMP_PROTOCOL_NAME)
    protocols = clientFactory.getAvailableProtocols(ip_address, ip_domain)
```

```
connected = 0
# Go over all the protocol credentials
for credentials_id in protocols:
    client = None
    try:
        # try to connect to the snmp agent
        client = clientFactory.createClient(credentials_id)

        // Query the agent
        ....

        # connection succeed
        connected = 1
    except:
        if client != None:
            client.close()
if (not connected):
    logger.debug('Failed to connect using all credentials')
else:
    // return the results as OSHV
    return resultVector
```

Handling Exceptions from Java

Some Java classes throw an exception upon failure. It is recommended to catch the exception and handle it, otherwise it causes the adapter to terminate unexpectedly.

When catching a known exception, in most cases you should print its stack trace to the log and issue a proper message to the UI, for example:

```
try:
    client = Framework.getClientFactory().createClient()
except Exception, msg:
    Framework.reportError('Connection failed')
    logger.debugException('Exception while connecting: %s' % (msg))
return
```

If the exception is not fatal and the script can continue, you should omit the call for the `reportError()` method and enable the script to continue.

How to Support Localization in Jython Adapters

The multi-lingual locale feature enables DFM to work across different operating system (OS) languages, and to enable appropriate customizations at runtime.

Previously, before Content Pack 3.00, DFM used statically-specified encoding to treat output from all network targets. However, this approach does not suit a multi-lingual IT network: to discover hosts with different OS languages, Probe administrators had to re-run DFM jobs manually several times with different job parameters each time. This procedure produced a serious overhead on network load but, even more, it avoided several key features of DFM, such as immediate job invocation on a trigger CI or automatic data refreshing in RTSM by the Schedule Manager.

The following locale languages are supported by default: Japanese, Russian, and German. The default locale is English.

This section includes:

- ▶ "Add Support for a New Language" on page 80
- ▶ "Change the Default Language" on page 82
- ▶ "Determine the Character Set for Encoding" on page 82
- ▶ "Define a New Job to Operate With Localized Data" on page 83
- ▶ "Decode Commands Without a Keyword" on page 84
- ▶ "Work with Resource Bundles" on page 85
- ▶ "API Reference" on page 86



Add Support for a New Language

This task describes how to add support for a new language.

This task includes the following steps:

- ▶ "Add a Resource Bundle (*.properties Files)" on page 80
- ▶ "Declare and Register the Language Object" on page 81

1 Add a Resource Bundle (*.properties Files)

Add a resource bundle according to the job that is to be run. The following table lists the DFM jobs and the resource bundle that is used by each job:

Job	Base Name of Resource Bundle
File Monitor by Shell	langFileMonitoring
Host Resources and Applications by Shell	langHost_Resources_By_TTY, langTCP
Hosts by Shell using NSLOOKUP in DNS Server	langNetwork
Host Connection by Shell	langNetwork

Job	Base Name of Resource Bundle
Collect Network Data by Shell or SNMP	langTCP
Host Resources and Applications by SNMP	langTCP
Microsoft Exchange Connection by NTCMD, Microsoft Exchange Topology by NTCMD	msExchange
MS Cluster by NTCMD	langMsCluster

For details on bundles, see "Work with Resource Bundles" on page 85.

2 Declare and Register the Language Object

To define a new language, add the following two lines of code to the `shellutils.py` script, that currently contains the list of all supported languages. The script is included in the `AutoDiscoveryContent` package. To view the script, access the Adapter Management window. For details, see "Adapter Management Window" in the *RTSM Data Flow Management Guide*.

- a** Declare the language, as follows:

```
LANG_RUSSIAN = Language(LOCALE_RUSSIAN, 'rus', ('Cp866', 'Cp1251'),
(1049,), 866)
```

For details on class `Language`, see "API Reference" on page 86. For details on the Class `Locale` object, see <http://download.oracle.com/javase/1.5.0/docs/api/java/util/Locale.html>. You can use an existing locale or define a new locale.

- b** Register the language by adding it to the following collection:

```
LANGUAGES = (LANG_ENGLISH, LANG_GERMAN, LANG_SPANISH,
LANG_RUSSIAN, LANG_JAPANESE)
```

Change the Default Language

If the OS language cannot be determined, the default one is used. The default language is specified in the `shellutils.py` file.

```
#default language for fallback
DEFAULT_LANGUAGE = LANG_ENGLISH
```

To change the default language, you initialize the `DEFAULT_LANGUAGE` variable with a different language. For details, see "Add Support for a New Language" on page 80.

Determine the Character Set for Encoding

The suitable character set for decoding command output is determined at runtime. The multi-lingual solution is based on the following facts and assumptions:

- 1** It is possible to determine the OS language in a locale-independent way, for example, by running the `chcp` command on Windows or the `locale` command on Linux.
- 2** Relation Language-Encoding is well known and can be defined statically. For example, the Russian language has two of the most popular encoding: Cp866 and Windows-1251.
- 3** One character set for each language is preferable, for example, the preferable character set for Russian language is Cp866. This means that most of the commands produce output in this encoding.
- 4** Encoding in which the next command output is provided is unpredictable, but it is one of the possible encoding for a given language. For example, when working with a Windows machine with a Russian locale, the system provides the `ver` command output in Cp866, but the `ipconfig` command is provided in Windows-1251.
- 5** A known command produces known key words in its output. For example, the `ipconfig` command contains the translated form of the **IP-Address** string. So the `ipconfig` command output contains **IP-Address** for the English OS, **IP-Адрес** for the Russian OS, **IP-Adresse** for the German OS, and so on.

Once it is discovered in which language the command output is produced (# 1), possible character sets are limited to one or two (# 2). Furthermore, it is known which key words are contained in this output (# 5).

The solution, therefore, is to decode the command output with one of the possible encoding by searching for a key word in the result. If the key word is found, the current character set is considered the correct one.



Define a New Job to Operate With Localized Data

This task describes how to write a new job that can operate with localized data.

Jython scripts usually execute commands and parse their output. To receive this command output in a properly decoded manner, you use the API for the **ShellUtils** class. For details, see "HP Universal CMDB Web Service API Overview" on page 308.

This code usually takes the following form:

```
client = Framework.createClient(protocol, properties)
shellUtils = shellutils.ShellUtils(client)
languageBundle = shellutils.getLanguageBundle('langNetwork', shellUtils.osLanguage,
Framework)
strWindowsIPAddress = languageBundle.getString('windows_ipconfig_str_ip_address')
ipconfigOutput = shellUtils.executeCommandAndDecode('ipconfig /all',
strWindowsIPAddress)
#Do work with output here
```

1 Create a client:

```
client = Framework.createClient(protocol, properties)
```

2 Create an instance of the **ShellUtils** class and add the operating system language to it. If the language is not added, the default language is used (usually English):

```
shellUtils = shellutils.ShellUtils(client)
```

During object initialization, DFM automatically detects machine language and sets preferable encoding from the predefined `Language` object. Preferable encoding is the first instance appearing in the encoding list.

- 3 Retrieve the appropriate resource bundle from `shellclient` using the `getLanguageBundle` method:

```
languageBundle = shellutils.getLanguageBundle ('langNetwork',
shellUtils.osLanguage, Framework)
```

- 4 Retrieve a keyword from the resource bundle, suitable for a particular command:

```
strWindowsIPAddress =
languageBundle.getString('windows_ipconfig_str_ip_address')
```

- 5 Invoke the `executeCommandAndDecode` method and pass the keyword to it on the `ShellUtils` object:

```
ipconfigOutput = shellUtils.executeCommandAndDecode('ipconfig /all',
strWindowsIPAddress)
```

The `ShellUtils` object is also needed to link a user to the API reference (where this method is described in detail).

- 6 Parse the output as usual.

Decode Commands Without a Keyword

The current approach for localization uses a keyword to decode all of the command output. For details, see step 4 on page 84 in "Define a New Job to Operate With Localized Data" on page 83.

However, another approach uses a keyword to decode the first command output only, and then decodes further commands with the character set used to decode the first command. To do this, you use the `getCharsetName` and `useCharset` methods of the `ShellUtils` object.

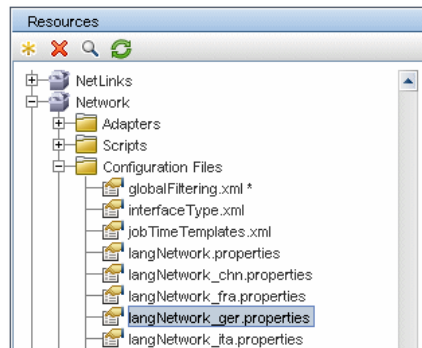
The regular use case works as follows:

- 1 Invoke the `executeCommandAndDecode` method once.
- 2 Obtain the most recently used character set name through the `getCharsetName` method.
- 3 Make `shellUtils` use this character set by default, by invoking the `useCharset` method on the `ShellUtils` object.
- 4 Invoke the `execCmd` method of `ShellUtils` one or more times. The output is returned with the character set specified in step 3. No additional decoding operations occur.

Work with Resource Bundles

A resource bundle is a file that takes a properties extension (`*.properties`). A properties file can be considered a dictionary that stores data in the format of `key = value`. Each row in a properties file contains one `key = value` association. The main functionality of a resource bundle is to return a value by its key.

Resource bundles are located on the Probe machine: `C:\hp\UCMDB\DataFlowProbe\runtime\probeManager\discoveryConfigFiles`. They are downloaded from the RTSM Server as any other configuration file. They can be edited, added, or removed, in the Resources window. For details, see "Configuration File Pane" in the *RTSM Data Flow Management Guide*.



When discovering a destination, DFM usually needs to parse text from command output or file content. This parsing is often based on a regular expression. Different languages require different regular expressions to be used for parsing. For code to be written once for all languages, all language-specific data must be extracted to resource bundles. There is a resource bundle for each language. (Although it is possible that a resource bundle contain data for different languages, in DFM one resource bundle always contains data for one language.)

The Jython script itself does not include hard coded, language-specific data (for example, language-specific regular expressions). The script determines the language of the remote system, loads the proper resource bundle, and obtains all language-specific data by a specific key.

In DFM, resource bundles take a specific name format: `<base_name>_<language_identifier>.properties`, for example, `langNetwork_spa.properties`. (The default resource bundle takes the following format: `<base_name>.properties`, for example, `langNetwork.properties`.)

The `base_name` format reflects the intended purpose of this bundle. For example, **langMsCluster** means the resource bundle contains language-specific resources used by the MS Cluster jobs.

The `language_identifier` format is a 3-letter acronym used to identify the language. For example, `rus` stands for the Russian language and `ger` for the German language. This language identifier is included in the declaration of the Language object.

API Reference

This section includes:

- "Language Class" on page 87
- "executeCommandAndDecode Method" on page 88
- "getCharsetName Method" on page 88
- "useCharset Method" on page 89
- "getLanguageBundle Method" on page 89
- "osLanguage Field" on page 89

Language Class

This class encapsulates information about the language, such as resource bundle postfix, possible encoding, and so on.

Fields

Name	Description
locale	Java object which represents locale.
bundlePostfix	Resource bundle postfix. This postfix is used in resource bundle file names to identify the language. For example, the langNetwork_ger.properties bundle includes a ger bundle postfix.
charsets	Character sets used to encode this language. Each language can have several character sets. For example, the Russian language is commonly encoded with the Cp866 and Windows-1251 encoding.
wmiCodes	The list of WMI codes used by the Microsoft Windows OS to identify the language. All possible codes are listed at http://msdn.microsoft.com/en-us/library/aa394239(VS.85).aspx (the OSLanguage section). One of the methods for identifying the OS language is to query the WMI class OS for the OSLanguage property.
codepage	Code page used with a specific language. For example, 866 is used for Russian machines and 437 for English machines. One of the methods for identifying the OS language is to retrieve its default codepage (for example, by the chcp command).

executeCommandAndDecode Method

This method is intended to be used by business logic Jython scripts. It encapsulates the decoding operation and returns a decoded command output.

Arguments

Name	Description
cmd	The actual command to be executed.
keyword	The keyword to be used for the decoding operation.
framework	The Framework object passed to every executable Jython script in DFM.
timeout	The command timeout.
waitForTimeout	Specifies if client should wait when timeout is exceeded.
useSudo	Specifies if sudo should be used (relevant only for UNIX machine clients).
language	Enables specifying the language directly instead of automatically detecting a language.

getCharsetName Method

This method return the name of the most recently used character set.

useCharset Method

This method sets the character set on the ShellUtils instance, which uses this character set for initial data decoding.

Arguments

Name	Description
charsetName	The name of the character set, for example, windows-1251 or UTF-8.

See also "getCharsetName Method" on page 88.

getLanguageBundle Method

This method should be used to obtain the correct resource bundle. This replaces the following API:

```
Framework.getEnvironmentInformation().getBundle(...)
```

Arguments

Name	Description
baseName	The name of the bundle without the language suffix, for example, langNetwork.
language	The language object. The ShellUtils.osLanguage should be passed here.
framework	The Framework, common object which is passed to every executable Jython script in DFM.

osLanguage Field

This field contains an object that represents the language.

How to Work with Discovery Analyzer

The Discovery Analyzer tool is intended for debugging purposes when developing packages, scripts, or any other content. The tool runs a job against a remote destination and returns logs containing information, warning, and error details and results of discovered CIs.

Note that results are not always reported to the UI. This is because the results are reported in two ways and only one of them is supported. Also, the communication log is not supported from Eclipse.

When executing the tool from Eclipse, the **DiscoveryProbe.properties** file (**C:\hp\UCMDB\DataFlowProbe\conf\DiscoveryProbe.properties**) must contain the following parameter set to **true**:

```
appilog.agent.local.discoveryAnalyzerFromEclipse = true
```

For details, see "How to Run Discovery Analyzer from Eclipse" on page 99.

In all other cases (when the tool is executed from the **cmd** file or while the Probe is running) this flag must be set to **false**:

```
appilog.agent.local.discoveryAnalyzerFromEclipse = false
```

Tasks and Records

A task file contains data regarding a task to be executed. The task consists of information such as the job's name and required parameters that define the trigger CI, for example, the remote destination address.

A record file contains task information as well as the results of a specific execution, that is, the detailed communication (including a response) between the Probe or Discovery Analyzer (whichever module executed the task) and the remote destination.

A task that is defined by a task file can be executed against a remote destination, whereas a task that is defined by a record file (that contains extra data regarding a specific execution) can be executed and can also be played back (that is, can reproduce the same execution documented in the record file).

Logs

Logs provide information about the latest run, as follows:

- ▶ **General Log.** This log includes all information data, errors, and warnings that occurred during the run.
- ▶ **Communication Log.** This log contains the detailed communication between the Discovery Analyzer and the remote destination (including its response). After the execution, the log can be saved as a record file.
- ▶ **Results Log.** Displays a list of discovered CIs. The appearance time of each CI depends on the design of the adapters and scripts.

You can save all logs together or each log separately. When you save all the logs, they are saved together under one name.

If you replay a record file, the same data is displayed in the communication log, the only difference being the time of execution.

Limitation: The Communication and Results logs are not available when running Discovery Analyzer through Eclipse.

This section includes the following steps:

- ▶ "Prerequisites" on page 92
- ▶ "Access Discovery Analyzer" on page 92
- ▶ "Define a Task" on page 93
- ▶ "Define a New Task" on page 94
- ▶ "Retrieve a Record" on page 95
- ▶ "Open a Task File" on page 95
- ▶ "Import a Task from the Database" on page 95
- ▶ "Edit a Task" on page 95
- ▶ "Save the Task and Logs" on page 96
- ▶ "Run the Task" on page 96

- ▶ "Send a Task Result to the Server" on page 97
- ▶ "Import Settings" on page 97
- ▶ "Breakpoints" on page 98
- ▶ "Configure Eclipse" on page 98

1 Prerequisites

- ▶ The Probe must be installed. (The Discovery Analyzer is installed as part of the Probe installation process and shares resources with it.)
- ▶ The Probe does not need to be running while you are working with Discovery Analyzer.

However, if the Probe has already run against an RTSM Server, all the required resources are already downloaded to the file system. If the Probe has not run, you can upload resources needed by Discovery Analyzer through the Settings menu. For details, see "Import Settings" on page 97.

- ▶ The RTSM Server does not need to be installed.

2 Access Discovery Analyzer

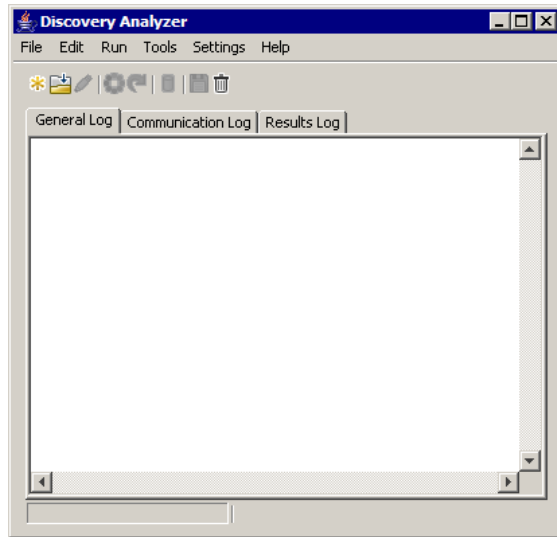
You access Discovery Analyzer either:

- ▶ When working with Eclipse.

The Probe installation comes with a default Eclipse workspace located at **C:\hp\UCMDB\DataFlowProbe\tools\discoveryAnalyzerWorkspace**. This workspace includes a Jython script to start Discovery Analyzer (**startDiscoveryAnalyzerScript.py**) as well as a link to all DFM scripts. If you start the tool in this way, you can locate breakpoints within the Jython scripts for debugging purposes.

- ▶ Directly, by double-clicking the file in the following folder: **C:\hp\UCMDB\DataFlowProbe\tools\discoveryAnalyzer.cmd**. For details, see the following section.

The Discovery Analyzer window opens:



3 Define a Task

You define a task using one of the following methods:


- ▶ By defining a new task. For details, see "Define a New Task" on page 94.
- ▶ By importing a task from a record file. For details, see "Retrieve a Record" on page 95.
- ▶ By importing a saved task from a task file. For details, see "Open a Task File" on page 95.
- ▶ By retrieving a job from the Probe's internal database. For details, see "Import a Task from the Database" on page 95.

4 Define a New Task



- a** Display the Task Editor: click the **New Task** button.

The Task Editor displays a list of jobs that currently exist in the file system. This list is updated each time the Probe receives tasks from the server, or packages are deployed manually from the Settings menu.



- b** Select a job.
- c** Enter values for all parameters.

The parameters displayed here are DFM adapter parameters. They can be viewed in the Discovery Pattern Parameters pane in the Pattern Signature tab.

All fields are mandatory (unless a job's script demands that the field be empty).

For parameters that require an ID or credentials ID input value, you can use randomly created IDs: right-click the value box and select **Generate random CMDB ID** or **Credential Chooser**.

The task is now active and the name of the open task is displayed in the title bar:



- d** Continue with the procedure for defining a task. For details, see "Save the Task and Logs" on page 96.

5 Retrieve a Record

You can define a task by opening a record file containing data regarding a specific execution. If a task is defined in this way, you can reproduce the specific execution by selecting the playback option. (If a task is replayed, responses are received from the data stored in the record file and not from the remote destination.)

Select **File > Open Record**. Browse to the folder where you saved the record. The record is now active and the name of the task is displayed in the title bar.

For details on acquiring a record file, see "How to Record DFM Code" on page 109.

6 Open a Task File

You can define a task from a task file: Select **File > Open Task**.

7 Import a Task from the Database

You can retrieve a task from the Probe database on condition that the Probe has already run and has active tasks in its internal database. You can use the parameter values to define the task.

- a** Select **File > Import Task from Probe Database**.
- b** In the dialog box that opens, select the task to run and click **OK**.
- c** Continue with the procedure for defining a task. For details, see "Save the Task and Logs" on page 96.

8 Edit a Task

After a task is defined, the name of the task (or the file) is displayed in the title bar. Now the file can be edited.

- a** Select **Edit > Edit Task**.
- b** Make any changes to the task and click **OK**.

9 Save the Task and Logs

You can save task parameters: Select **File > Save Task**.

The following options are available only after a task is executed.

- ▶ Save a record of the task. You can save the task parameters and the results of the task run: Select **File > Save Record**.
- ▶ Save a log of the task: Select **File > Save General Log**.
- ▶ Save results: Select **File > Save Results**.

10 Run the Task

The next step in the procedure is to run the task you created.

- a** Import the credentials/ranges configuration file. For details, see "Import Settings" on page 97.
- b** To execute the task only against a remote destination, click the **Run Task** button.

Discovery Analyzer executes the job and displays information in the three log files: **General**, **Communication**, and **Results**.

- c** You can save the log files, either together or separately: Select **File > Save General Log**, **Save Record**, **Save Results**, or **Save All Logs**. For details on the log files, see "Logs" on page 91.
- d** If a task is retrieved from a record file, the execution that is documented in this file can be reproduced by clicking the **Playback** button. The same Communication log is displayed, but the execution time is updated.

11 Send a Task Result to the Server

If a task's execution ends with results (that is, the Results Log tab displays a list of discovered CIs), you can send the results to the RTSM Server. This is useful if, for example, you were previously testing a script when the server was down.

Note: You can send results only to a an RTSM Server that receives tasks from the Probe that is installed on the same machine as Discovery Analyzer.

12 Import Settings

To run tasks or the playback record file, you must import the **domainScopeDocument.bin** file. During import, you enter a password.

- a** Launch a Web browser and enter the following URL:
http://localhost:8080/jmx-console. You may have to log in with a user name and password.
- b** Click **UCMDB:service=DiscoveryManager** to open the JMX MBEAN View page.
- c** Locate the **exportCredentialsAndRangesInformation** operation. Do the following:
 - Enter the customer ID (the default is **1**).
 - Enter a name for the exported file.
 - Enter the password.
 - Set **isEncrypted** to **False**.

- d** Click **Invoke** to export the **domainScopeDocument.bin** file.

When the export process completes successfully, the file is saved to the following location:

C:\hp\UCMDB\UCMDBServer\conf\discovery\<customer_dir>.

- e** Copy the **domainScopeDocument.bin** file to the Data Flow Probe file system and import it by selecting: **Settings > Import domainScopeDocument.**

Note: During the **domainScopeDocument** file import, you are requested to provide a password. This request is also displayed following each Discovery Analyzer restart and before the first task or record is executed.

13 Breakpoints

If you run Discovery Analyzer from the Python script, you can add breakpoints to your script.

14 Configure Eclipse

For details on running your Jython scripts in debug mode, see "How to Run Discovery Analyzer from Eclipse" on page 99.

How to Run Discovery Analyzer from Eclipse

This task explains how to configure Eclipse so that you can run your Jython scripts in debug mode, thus enabling better visibility to job threads, trigger CIs, and results.

This section includes the following steps:

- "Prerequisites" on page 99
- "Unpack Eclipse and start it" on page 99
- "Configure the default workspace" on page 100
- "Configure the PyDev Extensions" on page 101
- "Configure the Discovery Analyzer Workspace" on page 103
- "Configure the classpath and interpreter" on page 106
- "Run Discovery Analyzer" on page 109

1 Prerequisites


- Install the latest Eclipse version on your computer. The application is available at www.eclipse.org.
- Verify that the Data Flow Probe is installed on the same computer.
- Verify that the `appilog.agent.local.discoveryAnalyzerFromEclipse` parameter in the `DiscoveryProbe.properties` file is set to `true`.

2 Unpack Eclipse and start it

Unpack Eclipse and run the Eclipse executable file.


3 Configure the default workspace

Configure the default workspace where Eclipse saves and stores all projects and the related data.



4 Configure the PyDev Extensions


- a** Access **Help > Install New Software**, click **Add**, type a name for the PyDev plugin and in the Location field, add the URL of the site where pydev can be downloaded: <http://pydev.org/updates>. Click **OK**.



Note: PyDev and PyDev Extensions are now merged into one plugin since PyDev Extensions are now open source. For additional information visit <http://pydev.org>.

- b** In the window that opens, select **Pydev**. The second plugin is a plugin for task-focused UI. Click **Next**, check the installation details, and click **Next** again.

- c Accept the license agreement and click **Next**.
- d Pydev is installed. If you are asked to install unsigned content, confirm by clicking **OK**.

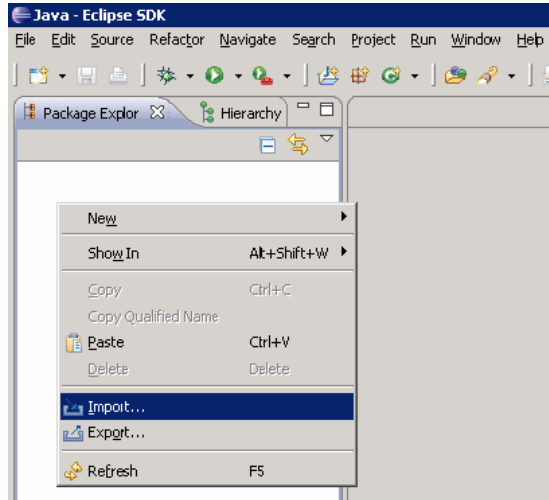


- e Restart Eclipse.


PyDev is now installed in your Eclipse IDE. You have new perspectives in Eclipse and the IDE is able to interpret Python scripts (text highlighting, additional configuration options, and so on).

5 Configure the Discovery Analyzer Workspace

- a Import necessary files: Right-click in the white area in Package Explorer and click **Import** to import the pre-configured **discoveryAnalyzerWorkspace**, included with the Probe installation.





- b** Under **General**, select **Existing projects into Workspace** to import the project into the Eclipse workspace.



- c** Under **Select root directory**, select the Analyzer workspace, usually located under:
`C:\hp\UCMDB\DataFlowProbe\tools\discoveryAnalyzerWorkspace`
- d** Select **Copy projects into workspace** to create a real copy of the existing workspace. This is an important step: In case of failure, you can re-import the original `discoveryAnalyserWorkspace`.

- e Click **Finish** to start the import.






6 Configure the classpath and interpreter

- a Right-click `discoveryAnalyzerWorkspace` and select **Properties** to display the Project specific settings.
- b Go to **Pydev > Interpreter/Grammar** and click **Please configure an interpreter in the related preferences before proceeding**.

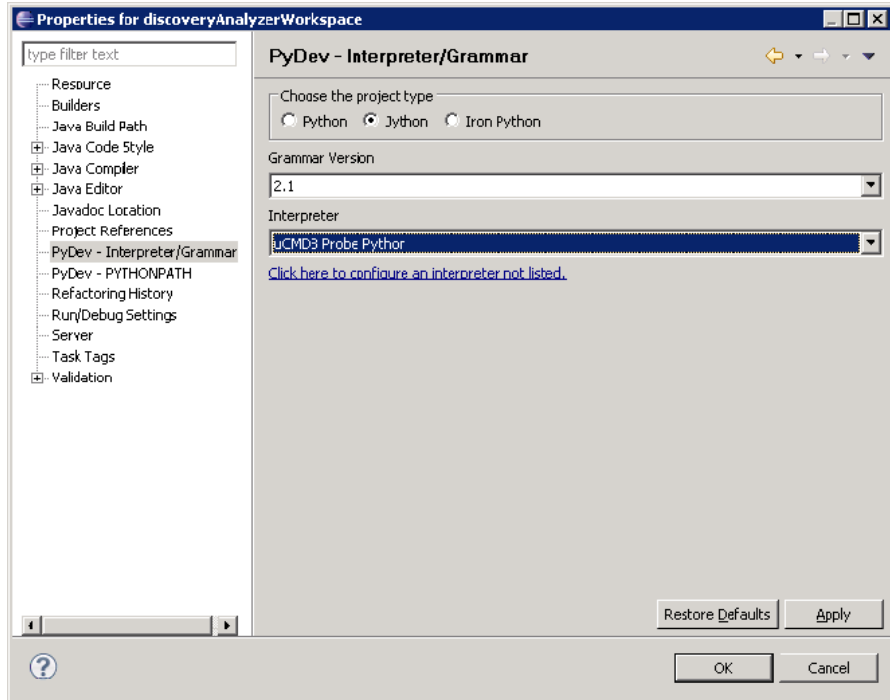
This step configures the same Jython interpreter as the Probe is using to ensure that scripts are not interpreted by a different Jython version.

- c Click **New**, type a name for the interpreter, and select the file from the following folder: **C:\hp\UCMDB\DataFlowProbe\jython\jython.jar**



- d Click **OK**. If a window is displayed, asking you to select the folders that should be imported into your Python system path, do not change anything (should be **C:\hp\UCMDB\DataFlowProbe\jython** and **C:\hp\UCMDB\DataFlowProbe\jython\lib**) and click **OK**.
- e Click **Apply** and then **OK**.

- f Click **Interpreter** and select the interpreter just created.

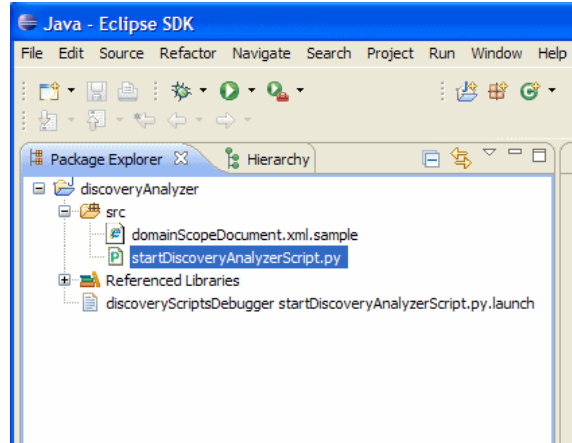


- g Click **Apply** and then **OK**.

The Jython interpreter is now the same as the one the Probe is using.

7 Run Discovery Analyzer

- a Add a breakpoint in the Jython script to be debugged.
- b To start Discovery Analyzer, select **startDiscoveryAnalyzerScript.py** in the **discoveryAnalyzerWorkspace\src** project. Right-click the file and choose **Debug as > Jython run**.




How to Record DFM Code

It can be very useful to record an entire execution, including all parameters, for example, when debugging and testing code. This task describes how to record an entire execution with all relevant variables. Furthermore, you can view extra debug information that is usually not printed to log files even at the debug level.

To record DFM code:

- 1 Access **Admin > RTSM Administration > Data Flow Management > Discovery Control Panel**. Right-click the job whose run must be logged and select **Edit adapter** to open the Adapter Management application.

2 Locate the **Execution Options** pane in the **Adapter Configuration** tab:



3 Change the **Create communication log** box to **Always**. For details on setting logging options, see "Execution Options Pane" in the *RTSM Data Flow Management Guide*.

The following example is the XML log file that is created when the Host Connection by Shell job is run and the **Create communication logs** box is set to **Always** or **On Failure**:

	Job name	Trigger CI data

```

- <execution jobId="Host Connection by Shell" destinationid="0e9787433d65e4a68839bfa8b224c92d">
- <destination>
  <destinationData name="ip_domain">DefaultDomain</destinationData>
  <destinationData name="hostId" />
  <destinationData name="ip_address">16.59.63.34</destinationData>
  <destinationData name="id">0e9787433d65e4a68839bfa8b224c92d</destinationData>
</destination>
    
```

The following example shows the message and stacktrace parameters:

Stacktrace

```
- <exec start="18:41:55" duration="2062" type="ssh" credentialsId="f464999bdfe5a1e1407b479b6f730d5b">
  <cmd>[CDATA: client_connect]</cmd>
  <result IS_NULL="Y" />
- <error class="com.hp.ucmdb.discovery.probe.services.dynamic.agents.SSHAgentException">
  <message>[CDATA: Failed to connect: Error connecting: Connection refused: connect]</message>
  - <stacktrace>
    <frame class="com.hp.ucmdb.discovery.probe.services.dynamic.agents.SSHAgent" method="connect" file="SSHAgent.java" />
    <frame class="com.hp.ucmdb.discovery.probe.clients.shell.SSHClient" method="createWrapper" file="SSHClient.java" />
    <frame class="com.hp.ucmdb.discovery.probe.clients.BaseClient" method="initPrivate" file="BaseClient.java" />
  </stacktrace>
</error>
</exec>
```

Reference

Jython Libraries and Utilities

Several utility scripts are used widely in adapters. These scripts are part of the AutoDiscovery package and are located under:

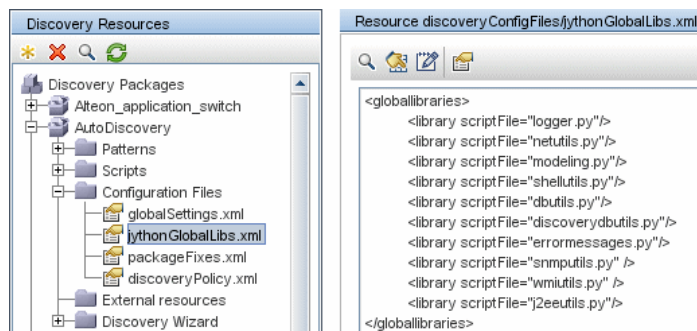
C:\hp\UCMDB\DataFlowProbe\runtime\probeManager\discoveryScripts with the other scripts that are downloaded to the Probe.

Note: The discoveryScript folder is created dynamically when the Probe begins working.

To use one of the utility scripts, add the following import line to the import section of the script:

```
import <script name>
```

The AutoDiscovery Python library contains Jython utility scripts. These library scripts are considered DFM's external library. They are defined in the jythonGlobalLibs.xml file (located in the **Configuration Files** folder).



Each script that appears in the `jythonGlobalLibs.xml` file is loaded by default at Probe startup, so there is no need to use them explicitly in the adapter definition.

This section includes the following topics:

- "logger.py" on page 113
- "modeling.py" on page 114
- "netutils.py" on page 114
- "shellutils.py" on page 115

logger.py

The `logger.py` script contains log utilities and helper functions for error reporting. You can call its `debug`, `info`, and `error` APIs to write to the log files. Log messages are recorded in `C:\hp\UCMDB\DataFlowProbe\runtime\log`.

Messages are entered in the log file according to the debug level defined for the `PATTERNS_DEBUG` appender in the

`C:\hp\UCMDB\DataFlowProbe\conf\log\probeMgrLog4j.properties` file. (By default, the level is `DEBUG`.) For details, see "Error Severity Levels" on page 122.

```
#####
#####          PATTERNS_DEBUG log          #####
#####
log4j.category.PATTERNS_DEBUG=DEBUG, PATTERNS_DEBUG
log4j.appender.PATTERNS_DEBUG=org.apache.log4j.RollingFileAppender
log4j.appender.PATTERNS_DEBUG.File=C:\hp\UCMDB\DataFlowProbe\runtime\log\pr
obeMgr-patternsDebug.log
log4j.appender.PATTERNS_DEBUG.Append=true
log4j.appender.PATTERNS_DEBUG.MaxFileSize=15MB
log4j.appender.PATTERNS_DEBUG.Threshold=DEBUG
log4j.appender.PATTERNS_DEBUG.MaxBackupIndex=10
log4j.appender.PATTERNS_DEBUG.layout=org.apache.log4j.PatternLayout
log4j.appender.PATTERNS_DEBUG.layout.ConversionPattern=<%d> [%-5p] [%t] -
%m%n
log4j.appender.PATTERNS_DEBUG.encoding=UTF-8
```

The info and error messages also appear in the Command Prompt console.

There are two sets of APIs:

- `logger.<debug/info/warn/error>`
- `logger.<debugException/infoException/warnException/errorException>`

The first set issues the concatenation of all its string arguments at the appropriate log level and the second set issues the concatenation as well as issuing the stack trace of the most recently-thrown exception, to provide more information, for example:

```
logger.debug('found the result')
logger.errorException('Error in discovery')
```

modeling.py

The **modeling.py** script contains APIs for creating hosts, IPs, process CIs, and so on. These APIs enable the creation of common objects and make the code more readable. For example:

```
ipOSH= modeling.createIpOSH(ip)
host = modeling.createHostOSH(ip_address)
member1 = modeling.createLinkOSH('member', ipOSH, networkOSH)
```

netutils.py

The **netutils.py** library is used to retrieve network and TCP information, such as retrieving operating system names, checking if a MAC address is valid, checking if an IP address is valid, and so on. For example:

```
dnsName = netutils.getHostName(ip, ip)
isValidIp = netutils.isValidIp(ip_address)
address = netutils.getHostAddress(hostName)
```

shellutils.py

The **shellutils.py** library provides an API for executing shell commands and retrieving the end status of an executed command, and enables running multiple commands based on that end status. The library is initialized with a Shell Client, and uses the client to run commands and retrieve results. For example:

```
ttyClient = clientFactory.createClient(Props)
clientShUtils = shellutils.ShellUtils(ttyClient)
if (clientShUtils.isWinOs()):
    logger.debug ('discovering Windows..')
```


4

Developing Generic Database Adapters

This chapter includes:

Concepts

- ▶ Generic Database Adapter Overview on page 119
- ▶ TQL Queries for the Generic Database Adapter on page 119
- ▶ Reconciliation on page 120
- ▶ Hibernate as JPA Provider on page 121

Tasks

- ▶ How to Prepare for Adapter Creation on page 124
- ▶ How to Prepare the Adapter Package on page 129
- ▶ How to Upgrade the Generic DB Adapter from 9.00 or 9.01 to 9.02 and Later on page 133
- ▶ How to Configure the Adapter – Minimal Method on page 134
- ▶ How to Configure the Adapter – Advanced Method on page 138
- ▶ How to Implement a Plugin on page 144
- ▶ How to Deploy the Adapter on page 148
- ▶ How to Edit the Adapter on page 148
- ▶ How to Create an Integration Point on page 148
- ▶ How to Create a View on page 149
- ▶ How to Calculate the Results on page 150
- ▶ How to View the Results on page 150
- ▶ How to View Reports on page 150

- ▶ How to Enable Log Files on page 150
- ▶ How to Use Eclipse to Map Between CIT Attributes and Database Tables on page 151

Reference

- ▶ Adapter Configuration Files on page 161
- ▶ Out-of-the-Box Converters on page 186
- ▶ Plugins on page 191
- ▶ Configuration Examples on page 191
- ▶ Adapter Log Files on page 203
- ▶ External References on page 205

Troubleshooting and Limitations on page 205

Concepts

Generic Database Adapter Overview

The purpose of the generic database adapter platform is to create adapters that can integrate with relational database management systems (RDBMS) and run TQL queries and population jobs against the database. The RDBMS supported by the generic database adapter are Oracle, Microsoft SQL Server, and MySQL.

This version of the database adapter implementation is based on a JPA (Java Persistence API) standard with the Hibernate ORM library as the persistence provider.


TQL Queries for the Generic Database Adapter

For population jobs, every required layout of a CI must be checked in the Layout Settings Dialog Box in the Modeling Studio. For details, see "Query Node/Relationship Properties Dialog Box" in the *RTSM Modeling Guide*. It is important to note that a CI might require an attribute to be identified, and without those attributes the CI will fail to be added to RTSM.


The following limitations exist on the TQL queries calculated by the Generic Database Adapter only:

- Subgraphs are not supported
- Compound relationships are not supported
- Cycles or cycle parts are not supported

The following TQL query is an example of a cycle:



- ▶ Function layout is not supported.
- ▶ 0..0 cardinality is not supported.
- ▶ The Join relationship is not supported.
- ▶ Qualifier conditions are not supported.
- ▶ To connect between two CIs, a relationship in the form of a table or foreign key must exist in the external database source.



Reconciliation

Reconciliation is carried out as part of the TQL calculation on the adapter side. For reconciliation to occur, the RTSM side is mapped to a federated entity called reconciliation CIT.

Mapping. Each attribute in the RTSM is mapped to a column in the data source.

Although mapping is done directly, transformation functions on the mapping data are also supported. You can add new functions through the Java code (for example, lowercase, uppercase). The purpose of these functions is to enable value conversions (values that are stored in the RTSM in one format and in the federated database in another format).

Note:

- ▶ To connect the RTSM and external database source, an appropriate association must exist in the database. For details, see "Prerequisites" on page 124.
 - ▶ Reconciliation with the RTSM id is also supported.
-

Hibernate as JPA Provider

Hibernate is an object-relational (OR) mapping tool, which enables mapping Java classes to tables over several types of relational databases (for example, Oracle and Microsoft SQL Server). For details, see "Functional Limitations" on page 206.

In an elementary mapping, each Java class is mapped to a single table. More advanced mapping enables inheritance mapping (as can occur in the RTSM database).

Other supported features include mapping a class to several tables, support for collections, and associations of types one-to-one, one-to-many, and many-to-one. For details, see "Associations" on page 123.

For our purposes, there is no need to create Java classes. The mapping is defined from the RTSM class model CITs to the database tables.

This section also includes the following topics:

- ▶ "Examples of Object-Relational Mapping" on page 122
- ▶ "Associations" on page 123


➤ "Usability" on page 123

Examples of Object-Relational Mapping


The following examples describe object-relational mapping:

Example of 1 RTSM Class Mapped to 1 Database Table:

Class M1, with attributes A1, A2, and A3, is mapped to table 1 columns c1, c2, and c3. This means that any M1 instance has a matching row in table 1.

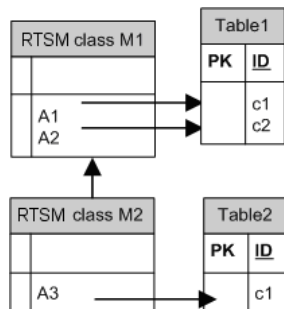


Example of 1 RTSM Class Mapped to 2 Database Tables:



Example of Inheritance:

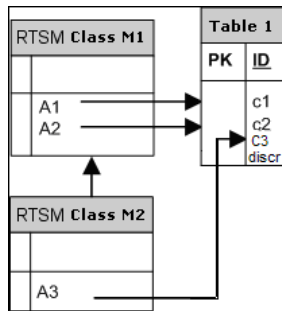
This case is used in the RTSM, where each class has its own database table.



Example of Single Table Inheritance with Discriminator:

An entire hierarchy of classes is mapped to a single database table, whose columns comprise a super-set of all attributes of the mapped classes. The table also contains an additional column (Discriminator), whose value indicates which specific class should be mapped to this entry.

When you use discriminator capabilities, you cannot skip a class in the hierarchy; that is, since C3 inherits from C2 and C2 inherits from C1, you cannot just define C1 and C3, you must define all three classes.



Associations

There are three types of associations: one-to-many, many-to-one and many-to-many. To connect between the different database objects, one of these associations must be defined by using a foreign key column (for the one-to-many case) or a mapping table (for the many-to-many case).

Usability

As the JPA schema is very extensive, a streamlined XML file is provided to ease definitions.

The use case for using this XML file is as follows: Federated data is modeled into one federated class. This class has many-to-one relations to a non-federated RTSM class. In addition, there is only one possible relation type between the federated class and the non-federated class.

Tasks

How to Prepare for Adapter Creation

This task describes the preparations that are necessary for creating an adapter.

Note: You can view samples for the Generic DB adapter in the UCMDB API. Specifically, the DDMi Adapter sample contains a complicated **orm.xml** file, as well as the implementations for some plugin interfaces.

This task includes the following steps:

- ▶ "Prerequisites" on page 124
- ▶ "Create a CI Type" on page 126
- ▶ "Create a Relationship" on page 127

1 Prerequisites

To validate that you can use the database adapter with your database, check the following:

- ▶ The reconciliation classes and their attributes (also known as multinodes) exist in the database. For example, if the reconciliation is run by node name, verify that there is a table that contains a column with node names. If the reconciliation is run according to node `cmdb_id`, verify that there is a column with RTSM IDs that matches the RTSM IDs of the nodes in the RTSM. For details on reconciliation, see "Reconciliation" on page 120.

ID	NAME	IP_ADDRESS
31	BABA	16.59.33.60
33	ext3.devlab.ad	16.59.59.116

ID	NAME	IP_ADDRESS
46	LABM1MAM15	16.59.58.188
72	cert-3-j2ee	16.59.57.100
102	labm1sun03.devlab.ad	16.59.58.45
114	LABM2PCOE73	16.59.66.79
116	CUT	16.59.41.214
117	labm1hp4.devlab.ad	16.59.60.182

- To correlate two CITs with a relationship, there must be correlation data between the CIT tables. The correlation can be either by a foreign key column or by a mapping table. For example, to correlate between node and ticket, there must be a column in the ticket table that contains the node ID, a column in the node table with the ticket ID that is connected to it, or a mapping table whose end1 is the node ID and end2 is the ticket ID. For details on correlation data, see "Hibernate as JPA Provider" on page 121.

The following table shows the foreign key NODE_ID column:

NODE_ID	CARD_ID	CARD_TYPE	CARD_NAME
2015	1	Serial Bus Controller	Intel © 82801EB USB Universal Host Controller
3581	2	System	Intel © 631xESB/6321ESB/3100 Chipset LPC
3581	3	Display	ATI ES1000
3581	4	Base System Peripheral	HP ProLiant iLO 2 Legacy Support Function

- Each CIT can be mapped to one or more tables. To map one CIT to more than one table, check that there is a primary table whose primary key exists in the other tables, and is a unique value column.

For example, a ticket is mapped to two tables: `ticket1` and `ticket2`. The first table has columns `c1` and `c2` and the second table has columns `c3` and `c4`. To enable them to be considered as one table, both must have the same primary key. Alternatively, the first table primary key can be a column in the second table.

In the following example, the tables share the same primary key called `CARD_ID`:

CARD_ID	CARD_TYPE	CARD_NAME
1	Serial Bus Controller	Intel ® 82801EB USB Universal Host Controller
2	System	Intel ® 631xESB/6321ESB/3100 Chipset LPC
3	Display	ATI ES1000
4	Base System Peripheral	HP ProLiant iLO 2 Legacy Support Function

CARD_ID	CARD_VENDOR
1	Hewlett-Packard Company
2	(Standard USB Host Controller)
3	Hewlett-Packard Company
4	(Standard system devices)
5	Hewlett-Packard Company

2 Create a CI Type

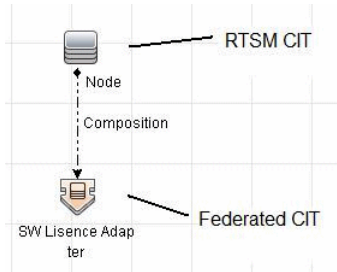
In this step you create a CIT that represents the data in the RDBMS (the external data source).

- a** In BSM, access the CI Type Manager and create a new CI Type. For details, see "How to Create a CI Type" in the *RTSM Modeling Guide*.
- b** Add the necessary attributes to the CIT, such as last access time, vendor, and so on. These are the attributes that the adapter will retrieve from the external data source and bring into RTSM views.

3 Create a Relationship

In this step you add a relationship between the BSM CIT and the new CIT that represents the data from the external data source.

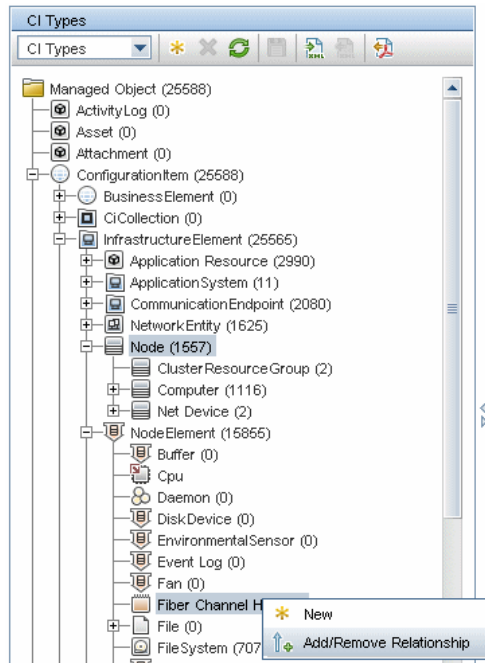
Add appropriate, valid relationships to the new CIT. For details, see "Add/Remove Relationship Dialog Box" in the *RTSM Modeling Guide*.



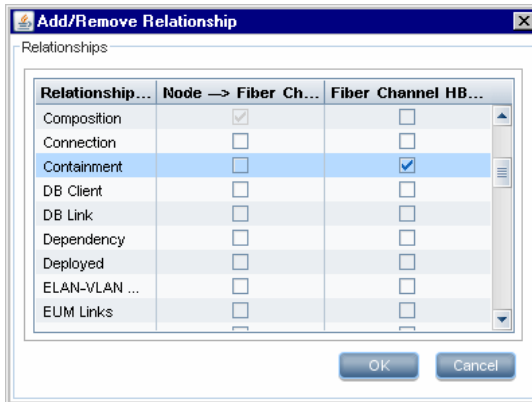
Note: At this stage, you cannot yet view the federated data or populate the external data, as you have not yet defined the method for bringing in the data.

Example of Creating a Containment Relationship:

1 In the CIT Manager, select the two CITs:



2 Create a **Containment** relationship between the two CITs:



How to Prepare the Adapter Package

In this step, you locate and configure the Generic DB adapter package.

- 1 Locate the **db-adapter.zip** package in the <HP BSM root directory>\odb\content\adapters folder.
- 2 Extract the package to a local temporary directory.
- 3 Edit the adapter XML file:
 - Open the **discoveryPatterns\db_adapter.xml** file in a text editor.
 - Locate the **adapter id** attribute and replace the name:

```
<pattern id="MyAdapter" displayLabel="My Adapter"
xsi:noNamespaceSchemaLocation="../../../Patterns.xsd" description="Discovery
Pattern Description"
  schemaVersion="9.0" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance" displayName="UCMDB API Population">
```

If the adapter supports population, the following capability should be added to the <adapter-capabilities> element:

```
<support-replicatioin-data>
  <source>
    <changes-source/>
  </source>
</support-replicatioin-data>
```

The display label or ID appears in the list of adapters in the Integration Point pane in Business Service Management.

If the plug-in for **FcmdbPluginForSyncGetChangesTopology** has not been implemented, only the following should be added:

```
<support-replicatioin-data>
  <source>
    <!--<changes-source/-->
  </source>
</support-replicatioin-data>
```

This will return the full topology and perform auto-delete according to the returned CIs.

For details about populating the RTSM with data, see "Integration Studio Page" in the *RTSM Data Flow Management Guide*.

- ▶ If the adapter is using the mapping engine from version 8.x (meaning that it is not using the new reconciliation mapping engine), replace the following element:

```
<default-mapping-engine/>
```

with

```
<default-mapping-engine>com.hp.ucmdb.federation.mappingEngine.AdapterMappingEngine</default-mapping-engine>
```

To revert to the new mapping engine, return the element to the following value:

```
<default-mapping-engine/>
```

- ▶ Locate the **category** definition:

```
<category>Generic</category>
```

Change the **Generic** category name to the category of your choice.

Note: Adapters whose categories are specified as **Generic** are not listed in the Integration Studio when you create a new integration point.

- ▶ The connection to the database can be described using a user name (schema), password, database type, database host machine name, and database name or SID.

For this type of connection, parameters have the following elements in the **parameter** section of the adapter's XML file:

```
<parameters>
  <!--The description attribute may be written in simple text or HTML.-->
  <!--The host attribute is treated as a special case by UCMDB-->
  <!--and will automatically select the probe name (if possible)-->
  <!--according to this attribute's value.-->
  <!--Display name and description may be overwritten by I18N values-->
  <parameter name="host" display-name="Hostname/IP" type="string"
description="The host name or IP address of the remote machine"
mandatory="false" order-index="10"/>
  <parameter name="port" display-name="Port" type="integer"
description="The remote machine's connection port" mandatory="false"
order-index="11"/>
  <parameter name="dbtype" display-name="DB Type" type="string"
description="The type of database"
valid-values="Oracle;SQLServer;MySQL;BO" mandatory="false"
order-index="13">Oracle</parameter>
  <parameter name="dbname" display-name="DB Name/SID" type="string"
description="The name of the database or its SID (in case of Oracle)"
mandatory="false" order-index="13"/>
  <parameter name="credentialsId" display-name="Credentials ID"
type="integer" description="The credentials to be used" mandatory="true"
order-index="12"/>
</parameters>
```

Note: This is the default configuration. Therefore, the **db_adapter.xml** file, already contains this definition.

There are situations in which the connection to the database cannot be configured in this way. For example, connecting to Oracle RAC or connecting using a database driver other than the one supplied with the RTSM.

For these situations, you can describe the connection using user name (schema), password, and a connection URL string.

To define this, edit the adapter's XML parameters section as follows:

```
<parameters>
  <!--The description attribute may be written in simple text or HTML.-->
  <!--The host attribute is treated as a special case by CMDBRTSM-->
  <!--and will automatically select the probe name (if possible)-->
  <!--according to this attribute's value.-->
  <!--Display name and description may be overwritten by I18N values-->
  <parameter name="url" display-name="Connection String" type="string"
description="The connection string to connect to the database" mandatory="true"
order-index="10"/>
  <parameter name="credentialsId" display-name="Credentials ID"
type="integer" description="The credentials to be used" mandatory="true"
order-index="12"/>
</parameters>
```

An example of a URL that connects to an Oracle RAC using the out-of-the-box Data Direct driver is: **`jdbc:mercury:oracle://labm3amdb17:1521;ServiceName=RACQA;AlternateServers=(labm3amdb18:1521);LoadBalancing=true`**.

- 4 In the temporary directory, open the **adapterCode** folder and rename **GenericDBAdapter** to the value of **adapter id** that was used in step 3.

This folder contains the adapter's configuration, for example, the adapter name, the queries and classes in the RTSM, and the fields in the RDBMS that the adapter supports.

- 5 Configure the adapter as required. For details, see "How to Configure the Adapter – Minimal Method" on page 134.
- 6 Create a *.zip file with the same name as you gave to the **adapter id** attribute, as described in step 3 on page 129.

Note: The **descriptor.xml** file is a default file that exists in every package.

- 7 Save the new package that you created in the previous step. The default directory for adapters is:


```
<HP BSM root directory>\odb\content\adapters.
```

How to Upgrade the Generic DB Adapter from 9.00 or 9.01 to 9.02 and Later

- 1 Copy your adapter package to a local temporary directory.
- 2 Extract the files.
- 3 Remove the following files from the `adapterCode\<Your Adapter Name>` folder:
 - `asm.jar`
 - `asm-attrs.jar`
 - `cglib.jar`
 - `db-adapter.jar`
 - `jboss-archive-browsing.jar`
 - `saxon-b.jar`
- 4 Recreate your adapter package.

Note: For any deployed Generic DB adapters that you may have, the UCMDB installer will remove the necessary files from the UCMDB and Probe file system. However, you still need to fix the package yourself, in order to re-deploy it when necessary.

How to Configure the Adapter – Minimal Method

The following procedure describes a method of mapping the class model in the RTSM to an RDBMS.

These configuration files are located in the **db-adapter.zip** package in the **<HP BSM root directory>\odb\content\adapters** folder that you extracted in step 2 of "How to Prepare the Adapter Package" on page 129.

Note: The **orm.xml** file that is automatically generated as a result of running this method is a good example that you can use when working with the advanced method.

You would use this minimal method when you need to:

- Federate/populate a single node such as a node attribute.
- Demonstrate the Generic Database Adapter capabilities.

This method:

- supports one-node federation\population only
- supports many-to-one virtual relationships only

This task includes the following steps:

- "Configure the adapter.conf File" on page 134
- "Configure the simplifiedConfiguration.xml File" on page 135

Configure the adapter.conf File

In this step, you change the settings in the **adapter.conf** file so that the adapter uses the simplified configuration method.

- 1** Open the **adapter.conf** file in a text editor.
- 2** Locate the following line: **use.simplified.xml.config=<true/false>**.
- 3** Change it to **use.simplified.xml.config=true**.

Configure the `simplifiedConfiguration.xml` File

In this step, you configure the `simplifiedConfiguration.xml` file by mapping the CIT in the RTSM to the fields in the RDBMS table.

1 Open the `simplifiedConfiguration.xml` file in a text editor.

This file includes a template that you use for each entity to be mapped.

Note: Do not edit the `simplifiedConfiguration.xml` file in any version of Notepad from Microsoft Corporation. Use Notepad++, UltraEdit, or some other third-party text editor.

2 Make changes to the following attributes:


- The CIT name in BSM (`cmdb-class-name`) and the corresponding table name in the RDBMS (`default-table-name`):

```
<cmdb-class cmdb-class-name="node" default-table-name="Device">
```

The `cmdb-class-name` attribute is taken from the node CIT:



The `default-table-name` attribute is taken from the Device table:



- The unique identifier in the RDBMS:

```
<primary-key column-name="Device_ID"/>
```

- The reconciliation rule (reconciliation-by-two-nodes):

```
<reconciliation-by-two-nodes connected-node-cmdb-class-name="ip_address"
cmdb-link-type="containment">
```

- The reconciliation attribute in BSM (`cmdb-attribute-name`) and in the RDBMS (`column-name`):

```
<connected-node-attribute cmdb-attribute-name="name"
column-name="[column_name]"/>
```


- The name of the CIT (`cmdb-class-name`) and the name of the corresponding table in the RDBMS (`default-table-name`). Also the RTSM relationship (`connected-cmdb-class-name`) and the CIT relationship (`link-class-name`):

```
<class cmdb-class-name="sw_sub_component"
default-table-name="SWSubComponent" connected-cmdb-class-name="node"
link-class-name="composition">
```

- The primary key and the foreign key:

```
<foreign-primary-key column-name="Device_ID"
cmdb-class-primary-key-column="Device_ID"/>
```

- The unique identifier in the RDBMS:

```
<primary-key column-name="Device_ID"/>
```

- The mapping between the RTSM attribute (`cmdb-attribute-name`) and the column name in the RDBMS (`column-name`):

```
<attribute cmdb-attribute-name="last_access_time"
column-name="SWSubComponent_LastAccess TimeStamp"/>
```

- 3 Save the file.

How to Configure the Adapter – Advanced Method

These configuration files are located in the **db-adapter.zip** package in the <HP BSM root directory>\odb\content\adapters folder that you extracted in step 2 of "How to Prepare the Adapter Package" on page 129.

This task includes the following steps:

- "Configure the orm.xml File" on page 138
- "Configure the reconciliation_types.txt File" on page 143
- "Configure the reconciliation_rules.txt File" on page 143

Configure the orm.xml File

In this step, you map the CITs and relationships in the RTSM to the tables in the RDBMS.

- 1 Open the **orm.xml** file in a text editor.


This file, by default, contains a template that you use to map as many CITs and relationships as needed.

Note: Do not edit the **orm.xml** file in any version of Notepad from Microsoft Corporation. Use Notepad++, UltraEdit, or some other third-party text editor.

- 2 Make changes to the file according to the data entities to be mapped. For details, see the following examples.

The following types of relationships may be mapped in the **orm.xml** file:


► One to one:



The code for this type of relationship is:

```
<one-to-one name="end1" target-entity="node">
  <join-column name="Device_ID" />
</one-to-one>
<one-to-one name="end2" target-entity="sw_sub_component">
  <join-column name="Device_ID" />
  <join-column name="Version_ID" />
</one-to-one>
```


► Many to one:



The code for this type of relationship is:

```
<many-to-one name="end1" target-entity="node">
  <join-column name="Device_ID" />
</many-to-one>
<one-to-one name="end2" target-entity="sw_sub_component">
  <join-column name="Device_ID" />
  <join-column name="Version_ID" />
</one-to-one>
```

- Many to many:



The code for this type of relationship is:

```
<many-to-one name="end1" target-entity="node">
  <join-column name="Device_ID" />
</many-to-one>
<many-to-one name="end2" target-entity="sw_sub_component">
  <join-column name="Device_ID" />
  <join-column name="Version_ID" />
</many-to-one>
```

For details about naming conventions, see "Naming Conventions" on page 171.

Example of Entity Mapping Between the Data Model and the RDBMS:

Note: Attributes that do not have to be configured are omitted from the following examples.

- The class of the RTSM CIT:

```
<entity class="generic_db_adapter.node">
```

- The name of the table in the RDBMS:

```
<table name="Device"/>
```

- The column name of the unique identifier in the RDBMS table:

```
<column name="Device ID"/>
```

- The name of the attribute in the RTSM CIT:

```
<basic name="name">
```

- The name of the table field in the external data source:

```
<column name="Device_Name"/>
```

- The name of the new CIT you created in "Create a CI Type" on page 126:

```
<entity class="generic_db_adapter.MyAdapter">
```

- The name of the corresponding table in the RDBMS:

```
<table name="SW_License"/>
```

- The unique identity in the RDBMS:

```
<id name="id1">
  <column updatable="false" insertable="false" name="Device_ID"/>
  <generated-value strategy="TABLE"/>
</id>
<id name="id2">
  <column updatable="false" insertable="false" name="Version_ID"/>
  <generated-value strategy="TABLE"/>
</id>
```

- The attribute name in the RTSM CIT and the name of the corresponding attribute in the RDBMS:

```
<basic name="license_required">
  <column updatable="false" insertable="false"
name="MyAdapter_LicenseRequired"/>
```

Example of Relationship Mapping Between the Data Model and the RDBMS:

- ▶ The class of the RTSM relationship:

```
<entity class="generic_db_adapter.node_containment_MyAdapter">
```

- ▶ The name of the RDBMS table where the relationship is performed:

```
<table name="MyAdapter"/>
```

- ▶ The unique ID in the RDBMS:

```
<id name="id1">  
  <column updatable="false" insertable="false" name="Device_ID"/>  
  <generated-value strategy="TABLE"/>  
</id>  
<id name="id2">  
  <column updatable="false" insertable="false" name="Version_ID"/>  
  <generated-value strategy="TABLE"/>  
</id>
```

- ▶ The relationship type and the RTSM CIT:

```
<many-to-one target-entity="node" name="end1">
```

- ▶ The primary key and foreign key fields in the RDBMS:

```
<join-column updatable="false" insertable="false"  
referenced-column-name="[column_name]" name="Device_ID"/>
```

Configure the reconciliation_types.txt File

Open the `reconciliation_types.txt` file in a text editor.

For details, see "reconciliation_types.txt file" on page 179.

Configure the reconciliation_rules.txt File

In this step you define the rules by which the adapter reconciles the RTSM and the RDBMS (only if Mapping Engine is used, for backward compatibility with version 8.x):

- 1 Open `META-INF\reconciliation_rules.txt` in a text editor.
- 2 Make changes to the file according to the CIT you are mapping. For example, to map a node CIT, use the following expression:

```
multinode[node] ordered expression[^name]
```

Note:

- If the data in the database is case sensitive, do not delete the control character (^).
 - Check that each opening square bracket has a matching closing bracket.
-

For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.

How to Implement a Plugin

This task describes how to implement and deploy a Generic DB Adapter with plugins.

Note: Before writing a plugin for an adapter, make sure you have completed all the necessary steps in "How to Prepare the Adapter Package" on page 129.

- 1 Copy the following jar files from the UCMDB server installation directory to your development class path:
 - Copy the **db-interfaces.jar** file and **db-interfaces-javadoc.jar** file from the **tools\adapter-dev-kit\db-adapter-framework** folder.
 - Copy the **federation-api.jar** file and **federation-api-javadoc.jar** file from the **\tools\adapter-dev-kit\SampleAdapters\production-lib** folder.

Note: More information about developing a plugin can be found in the **db-interfaces-javadoc.jar** and **federation-api-javadoc.jar** files and in the online documentation at:

- **C:\hp\UCMDB\UCMDBServer\deploy\ucmdb-docs\docs\eng\doc_lib\DevRef_guide\DBAdapterFramework_JavaAPI\index.html**
 - **C:\hp\UCMDB\UCMDBServer\deploy\ucmdb-docs\docs\eng\doc_lib\DevRef_guide\Federation_JavaAPI\index.html**
-

- 2 Write a Java class implementing the plugin's Java interface. The interfaces are defined in the **db-interfaces.jar** file. The table below specifies the interface that must be implemented for each plugin:

Plugin Type	Interface Name	Method
Synchronize Full Topology	FcmdbPluginForSyncGetFull Topology	getFullTopology
Synchronize Changes	FcmdbPluginForSyncGet ChangesTopology	getChangesTopology
Synchronize Layout	FcmdbPluginForSyncGet Layout	getLayout
Retrieve Supported Queries	FcmdbPluginForSyncGet SupportedQueries	getSupportedQueries
Alter TQL query definition and results	FcmdbPluginGetTopology CmdbFormat	getTopologyCmdbFormat
Alter layout request for CIs	FcmdbPluginGetCIsLayout	getCisLayout
Alter layout request for links	FcmdbPluginGetRelations Layout	getRelationsLayout

The plugin's class must have a public default constructor. Also, all of the interfaces expose a method called `initPlugin`. This method is guaranteed to be called before any other method and is used to initialize the adapter with the containing adapter's environment object.

If **FcmdbPluginForSyncGetChangesTopology** is implemented, there are two different ways to report the changes:

- **Report the entire root topology at all times.** According to this topology, the auto-delete function finds which CIs should be removed. In this case, the auto-delete function should be enabled by using the following:

```
<autoDeleteCITs isEnabled="true">  
  <CIT>link</CIT>  
  <CIT>object</CIT>  
</autoDeleteCITs>
```

- **Report each CI instance that was removed/updated.** In this case the auto-delete mechanism should be disabled by using the following:

```
<autoDeleteCITs isEnabled="false">  
  <CIT>link</CIT>  
  <CIT>object</CIT>  
</autoDeleteCITs>
```

- 3** Make sure you have the Federation SDK JAR and the Generic DB Adapter JARs in your class path before compiling your Java code. The Federation SDK is the **federation_api.jar** file, which can be found in the **<HP BSM root directory>\odb\lib** directory.
- 4** Pack your class into a jar file and put it under the adapterCode**<Your Adapter Name>** folder in the adapter package, prior to deploying it.
The plug-ins are configured using the **plugins.txt** file, located in the **\META-INF** folder of the adapter.

The following is an example of the file from the DDMi adapter:

```
# mandatory plugin to sync full topology
[getFullTopology]
com.hp.ucmdb.adapters.ed.plugins.replication.EDReplicationPlugin

# mandatory plugin to sync changes in topology
[getChangesTopology]
com.hp.ucmdb.adapters.ed.plugins.replication.EDReplicationPlugin

# mandatory plugin to sync layout
[getLayout]
com.hp.ucmdb.adapters.ed.plugins.replication.EDReplicationPlugin

# plugin to get supported queries in sync. If not defined return all tqIs names
[getSupportedQueries]

# internal not mandatory plugin to change tqI definition and tqI result
[getTopologyCmdbFormat]

# internal not mandatory plugin to change layout request and CIs result
[getCisLayout]

# internal not mandatory plugin to change layout request and relations result
[getRelationsLayout]
```

Legend:



- A comment line.

[<Adapter Type>] – Start of the definition section for a specific adapter type.

There can be an empty line under each [<Adapter Type>], meaning that there is no plugin class associated, or the fully qualified name of your plugin class can be listed.

- 5 Pack your adapter with the new jar file and the updated **plugins.xml** file. The remainder of the files in the package should be the same as in any adapter based on the Generic DB adapter.

How to Deploy the Adapter

- 1** In BSM, access the Package Manager. For details, see "Package Manager Page" in the *RTSM Administration Guide*.
-  **2** Click the **Deploy Packages to Server (from local disk)** icon and browse to your adapter package. Select the package and click **Open**, then click **Deploy** to display the package in the Package Manager.
-  **3** Select your package in the list and click the **View package resources** icon to verify that the package contents are recognized by Package Manager.

How to Edit the Adapter

Once you have created and deployed the adapter, you can then edit it within BSM. For details, see "Adapter Management" on page 169.

How to Create an Integration Point

In this step you check that the federation is working, that is, that the connection is valid and that the XML file is valid. However, this check does not verify that the XML is mapping to the correct fields in the RDBMS.


- 1** In BSM, access the Integration Studio (**Admin > RTSM Administration > Data Flow Management > Integration Studio**).
- 2** Create an integration point. For details, see "New Integration Point/Edit Integration Point Dialog Box" in the *RTSM Data Flow Management Guide*.

The Federation tab displays all CITs that can be federated using this integration point. For details, see "Federation Tab" in the *RTSM Data Flow Management Guide*.

How to Create a View

In this step you create a view that enables you to view instances of the CIT.

- 1 In BSM, access the Modeling Studio (**Admin > RTSM Administration > Modeling > Modeling Studio**).
- 2 Create a view. For details, see "How to Create a Pattern View" in the *RTSM Modeling Guide*.
- 3 You can add conditions to the TQL, for example, the last access time is greater than six months:



The screenshot shows the SWLicenseAdapter Modeling Studio interface. The main workspace displays a hierarchical diagram with a 'Node' icon connected to an 'SW License Adapter' icon by a 'Container link'. Below the workspace is the 'Information Pane' with the following tabs: Attributes, Cardinality, Qualifiers, Selected Identities, Grouping, and Details. The 'Attributes' tab is active, showing the following TQL query:

```
sw_license_required Equal true
AND last_access_time Less than or equal Sun Feb 3 2008 06:11 PM IST
```

An 'Edit' button is located at the bottom right of the Information Pane.

How to Calculate the Results

In this step you check the results.

- 1 In BSM, access the Modeling Studio (**Admin > RTSM Administration > Modeling > Modeling Studio**).
- 2 Open a view.
- 3 Calculate results by clicking the **Calculate Query Result Count** button.
- 4 Click the **Preview** button to view the CIs in the view.



How to View the Results

In this step you view the results and debug problems in the procedure. For example, if nothing is shown in the view, check the definitions in the **orm.xml** file; remove the relationship attributes and reload the adapter.

- 1 In BSM, access the IT Universe Manager (**Admin > RTSM Administration > Modeling > IT Universe Manager**).
- 2 Select a CI. The Properties tab displays the results of the federation.

How to View Reports

In this step you view Topology reports. For details, see "Topology Reports Overview" in the *RTSM Modeling Guide*.

How to Enable Log Files

To understand the calculation flows, adapter lifecycle, and to view debug information, you can consult the log files. For details, see "Adapter Log Files" on page 203.

How to Use Eclipse to Map Between CIT Attributes and Database Tables

Caution: This procedure is intended for users with an advanced knowledge of content development. For any questions, contact HP Software Support.

This task describes how to install and use the JPA plugin, provided with the J2EE edition of Eclipse, to:

- Enable graphical mapping between RTSM class attributes and database table columns.
- Enable manual editing of the mapping file (`orm.xml`), while providing correctness. The correctness check includes a syntax check as well as verification that the class attributes and mapped database table columns are stated correctly.
- Enable deployment of the mapping file to the RTSM server and to view the errors, as a further correctness check.
- Define a sample query on the RTSM server and run it directly from Eclipse, to test the mapping file.

Version 1.1 of the plugin is compatible with UCMDB version 9.01 or later and Eclipse IDE for Java EE Developers, version 1.2.2.20100217-2310 or later.

This task includes the following steps:

- "Prerequisites" on page 152
- "Installation" on page 152
- "Prepare the Work Environment" on page 153
- "Create an Adapter" on page 154
- "Configure the RTSM Plugin" on page 154
- "Import the BSM Class Model" on page 154
- "Build the ORM File – Map BSM Classes to Database Tables" on page 155

- "Map IDs" on page 155
- "Map Attributes" on page 156
- "Map a Valid Link" on page 156
- "Build the ORM File – Use Secondary Tables" on page 157
- "Define a Secondary Table" on page 157
- "Map an Attribute to a Secondary Table" on page 158
- "Use an Existing ORM File as a Base" on page 158
- "Importing an Existing ORM File from an Adapter" on page 158
- "Check the Correctness of the orm.xml File – Built-in Correctness Check" on page 159
- "Create a New Integration Point" on page 159
- "Deploy the ORM File to the RTSM" on page 159
- "Run a Sample TQL Query" on page 160

1 Prerequisites

Install the latest update for **Java Runtime Environment (JRE) 6** on the machine where you will run Eclipse from the following site:
<http://java.sun.com/javase/downloads/index.jsp>.

2 Installation

- a** Download and extract **Eclipse IDE for Java EE Developers** from <http://www.eclipse.org/downloads> to a local folder, for example, **C:\Program Files\eclipse**.
- b** Copy **com.hp.plugin.import_cmdb_model_1.0.jar** from **C:\hp\UCMDB\UCMDBServer\tools\db-adapter-eclipse-plugin\bin** to **C:\Program Files\Eclipse\plugins**.
- c** Launch **C:\Program Files\Eclipse\eclipse.exe**. If a message is displayed that the Java virtual machine is not found, launch **eclipse.exe** with the following command line:

```
"C:\Program Files\eclipse\eclipse.exe" -vm "<JRE installation folder>\bin"
```


3 Prepare the Work Environment

In this step, you set up the workspace, database, connections, and driver properties.

- a Extract the file **workspaces_gdb.zip** from **C:\hp\UCMDB\UCMDBServer\tools\db-adapter-eclipse-plugin\workspace** into **C:\Documents and Settings\All Users**.

Note: You must use the exact folder path. If you unzip the file to the wrong path or leave the file unzipped, the procedure will not work.

- b In Eclipse, choose **File > Switch Workspace > Other:**

If you are working with:

- SQL Server, select the following folder: **C:\Documents and Settings\All Users\workspace_gdb_sqlserver**.
- MySQL, select the following folder: **C:\Documents and Settings\All Users\workspace_gdb_mysql**.
- Oracle, select the following folder: **C:\Documents and Settings\All Users\workspace_gdb_oracle**.

- c Click **OK**.
- d In Eclipse, display the Project Explorer view and select **<Active project> > JPA Content > persistence.xml > <active project name> > orm.xml**.
- e In the Data Source Explorer view (the bottom left pane), right-click the database connection and select the **Properties** menu.
- f In the **Properties for <Connection name>** dialog box, select **Common** and select the **Connect every time the workbench is started** check box. Select **Driver Properties** and fill in the connection properties. Click **Test Connection** and verify that the connection is working. Click **OK**.

- g** In the Data Source Explorer view, right-click the database connection and click **Connect**. A tree containing the database schemas and tables is displayed under the database connection icon.

4 Create an Adapter

Create an adapter using the guidelines in "Step 1: Create an Adapter" on page 40.

5 Configure the RTSM Plugin

- a** In Eclipse, click **UCMDB > Settings** to open the **CMDB Settings** dialog box.
- b** If not already selected, select the newly created JPA project as the Active project.
- c** Enter the RTSM host name, for example, **localhost** or **labm1.itdep1**. There is no need to include the port number or **http://** prefix in the address.
- d** Fill in the user name and password for accessing the RTSM API, usually **admin/admin**.
- e** Make sure that the **C:\hp** folder on the RTSM server is mapped as a network drive.
- f** Select the base folder of the relevant adapter under **C:\hp**. The base folder is the one that contains the **dbAdapter.jar** file and the **META-INF** subfolder. Its path should be **<HP BSM root directory>\odb\ runtime\fcmdb\CodeBase\<adapter name>**. Verify that there is no backslash (****) at the end.

6 Import the BSM Class Model

In this step, you select the CITs to be mapped as JPA entities.

- a** Click **UCMDB > Import CMDB Class Model** to open the **CI Types Selection** dialog box.
- b** Select the CI types that you intend to map as JPA entities. Click **OK**. The CI types are imported as Java classes. Verify that they appear under the **src** folder of the active project.

7 Build the ORM File – Map BSM Classes to Database Tables

In this step, you map the Java classes (that you imported in the previous step) to the database tables.

- a** Make sure the DB connection is connected. Right-click the active project (called myProject by default) in Project Explorer. Select the JPA view, select the **Override default schema from connection** check box, and select the relevant database schema. Click **OK**.
- b** Map a CIT: In the JPA Structure view, right-click the **Entity Mappings** branch and select **Add Class**. The **Add Persistent Class** dialog box opens. Do not change the **Map as** field (**Entity**).
- c** Click **Browse** and select the BSM class to be mapped (all BSM classes belong to the **generic_db_adapter** package).
- d** Click **OK** in both dialog boxes. The selected class is displayed under the **Entity Mappings** branch in the JPA Structure view.

Note: If the entity appears without an attribute tree, right-click the active project in the Project Explorer view. Choose **Close** and then **Open**.

- e** In the JPA Details view, select the primary database table to which the BSM class should be mapped. Leave all other fields unchanged.

8 Map IDs

According to JPA standards, each persistent class must have at least one ID attribute. For BSM classes, you can map up to three attributes as IDs. Potential ID attributes are called **id1**, **id2**, and **id3**. To map an ID attribute:

- a** Expand the corresponding class under the **Entity Mappings** branch in the JPA Structure view, right-click the relevant attribute (for example, **id1**), and select **Add Attribute to XML and Map...**
- b** The **Add Persistent Attribute** dialog box opens. Select **Id** in the **Map as** field and click **OK**.

- c In the JPA Details view, select the database table column to which the ID field should be mapped.

9 Map Attributes

In this step, you map attributes to the database columns.

- a Expand the corresponding class under the **Entity Mappings** branch in the JPA Structure view, right-click the relevant attribute (for example, **host_hostname**), and select **Add Attribute to XML and Map....**
- b The **Add Persistent Attribute** dialog box opens. Select **Basic** in the **Map as** field and click **OK**.
- c In the JPA Details view, select the database table column to which the attribute field should be mapped.

10 Map a Valid Link

Perform the steps described in step b on page 155 for mapping a BSM class denoting a valid link. The name of each such class takes the following structure: **<end1 entity name>_<link name>_<end 2 entity name>**. For example, a **Contains** link between a host and a location is denoted by a Java class whose name is **generic_db_adapter.host_contains_location**. For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.

- a Map the ID attributes of the link class as described in "Map IDs" on page 155. For each ID attribute, expand the **Details** check box group in the JPA Details view and clear the **Insertable** and **Updateable** check boxes.
- b Map the **end1** and **end2** attributes of the link class as follows: For each of the **end1** and **end2** attributes of the link class:
 - Expand the corresponding class under the **Entity Mappings** branch in the JPA Structure view, right-click the relevant attribute (for example, **end1**), and select **Add Attribute to XML and Map....**
 - In the **Add Persistent Attribute** dialog box, select **Many to One** or **One to One** in the **Map as** field.

- ▶ Select **Many to One** if the specified **end1** or **end2** CI can have multiple links of this type. Otherwise, select **One to One**. For example, for a **host_contains_ip** link the **host** end should be mapped as **Many to One**, since one host can have multiple IPs, and the **ip** end should be mapped as **One to One**, since one IP can have only a single host.
- ▶ In the JPA Details view, select **Target entity**, for example, **generic_db_adapter.host**.
- ▶ In the **Join Columns** section of the JPA Details view, check **Override Default**. Click **Edit**. In the **Edit Join Column** dialog box, select the foreign key column of the link database table that points to an entry in the **end1/end2** target entity's table. If the referenced column name in the **end1/end2** target entity's table is mapped to its ID attribute, leave the **Referenced Column Name** unchanged. Otherwise, select the name of the column to which the foreign key column points. Clear the **Insertable** and **Updatable** check boxes and click **OK**.
- ▶ If the **end1/end2** target entity has more than one ID, click the **Add** button to add additional join columns and map them in the same way as described in the previous step.

11 Build the ORM File – Use Secondary Tables

JPA enables a Java class to be mapped to more than one database table. For example, **Host** can be mapped to the **Device** table to enable persistence of most of its attributes and to the **NetworkNames** table to enable persistence of **host_hostName**. In this case, **Device** is the primary table and **NetworkNames** is the secondary table. Any number of secondary tables can be defined. The only condition is that there must be a one-to-one relationship between the entries of the primary and secondary tables.

12 Define a Secondary Table

Select the appropriate class in the JPA Structure view. In the **JPA Details** view, access the **Secondary Tables** section and click **Add**. In the **Add Secondary Table** dialog box, select the appropriate secondary table. Leave the other fields unchanged.

If the primary and the secondary table do not have the same primary keys, configure the join columns in the **Primary Key Join Columns** section of the **JPA Details** view.

13 Map an Attribute to a Secondary Table

You map a class attribute to a field of a secondary table as follows:

- a Map the attribute as described in "Map Attributes" on page 156.
- b In the **Column** section of the JPA Details view, select the secondary table name in the **Table** field, to replace the default value.

14 Use an Existing ORM File as a Base

To use an existing **orm.xml** file as a basis for the one you are developing, perform the following steps:

- a Verify that all CITs mapped in the existing **orm.xml** file are imported into the active Eclipse project.
- b Select and copy all or part of the entity mappings from the existing file.
- c Select the **Source** tab of the **orm.xml** file in the Eclipse JPA perspective.
- d Paste all copied entity mappings under the **<entity-mappings>** tag of the edited **orm.xml** file, beneath the **<schema>** tag. Make sure that the schema tag is configured as described in step b on page 155. All pasted entities now appear in the JPA Structure view. From now on, mappings can be edited both graphically and manually through the xml code of the **orm.xml** file.
- e Click **Save**.

15 Importing an Existing ORM File from an Adapter

If an adapter already exists, the Eclipse Plugin can be used to edit its ORM file graphically. Import the **orm.xml** file into Eclipse, edit it using the plugin and then deploy it back to the BSM machine. To import the ORM file, press the button on the Eclipse toolbar. A confirmation dialog is displayed. Click **OK**. The ORM file is copied from the BSM machine to the active Eclipse project and all relevant classes are imported from the BSM class model.

If the relevant classes do not appear in the JPA Structure view, right-click the active project in the Project Explorer view, choose **Close** and then **Open**.

From now on, the ORM file can be edited graphically using Eclipse, and then deployed back to the BSM machine as described in "Deploy the ORM File to the RTSM" on page 159.

16 Check the Correctness of the orm.xml File – Built-in Correctness Check

The Eclipse JPA plugin checks if any errors are present and marks them in the **orm.xml** file. Both syntax (for example, wrong tag name, unclosed tag, missing ID) and mapping errors (for example, wrong attribute name or database table field name) are checked. If there are errors, their description appears in the **Problems** view:

17 Create a New Integration Point

If no integration point exists in the RTSM for this adapter, you can create it in the Integration Studio. For details, see "Integration Studio" in the *RTSM Data Flow Management Guide*.

Fill in the integration point name in the dialog box that opens. The **orm.xml** file is copied to the adapter folder. An integration point is created with all the imported CI types as its supported classes, except for multinode CITs, if they are configured in the **reconciliation_rules.txt** file. For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.

18 Deploy the ORM File to the RTSM

Save the **orm.xml** file and deploy it to the BSM server by clicking **UCMDB > Deploy ORM**. The **orm.xml** file is copied to the adapter folder and the adapter is reloaded. The operation result is shown in an **Operation Result** dialog box. If any error occurs during the reload process, the Java exception stack trace is displayed in the dialog box. If no integration point has yet been defined using the adapter, no mapping errors are detected upon deployment.

19 Run a Sample TQL Query

- a** Define a query (not a view) in the Modeling Studio. For details, see "Modeling Studio" in the *RTSM Modeling Guide*.
- b** Create an integration point using the adapter that you created in step 17 on page 159. For details, see "New Integration Point/Edit Integration Point Dialog Box" in the *RTSM Data Flow Management Guide*.
- c** During the creation of the adapter, verify that the CI types that should participate in the query are supported by this integration point.
- d** When configuring the RTSM plugin, use this sample query name in the Settings dialog box. For details, see "Configure the RTSM Plugin" on page 154.
- e** Click the **Run TWL** button to run a sample TQL and verify whether it returns the required results using the newly created **orm.xml** file.

Reference

Adapter Configuration Files

The files discussed in this section are located in the **db-adapter.zip** package in the <HP BSM root directory>\odb\content\adapters folder.

This section includes the following topics:

- "General Configuration" on page 161
- "Advanced Configuration" on page 161
- "Hibernate Configuration" on page 162
- "Simple Configuration" on page 162

General Configuration

- **adapter.conf**. The adapter configuration file. For details, see "Adapter.conf File" on page 162.

Advanced Configuration

- **orm.xml**. The object-relational mapping file in which you map between RTSM CITs and database tables. For details, see "orm.xml File" on page 166.
- **reconciliation_types.txt**. Contains the rules that are used to configure the reconciliation types. For details, see "reconciliation_types.txt file" on page 179.
- **reconciliation_rules.txt**. Contains the reconciliation rules. For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.
- **transformations.txt**. Transformations file in which you specify the converters to apply to convert from the RTSM value to the database value, and vice versa. For details, see "transformations.txt File" on page 181.

- ▶ **Discriminator.properties.** This file maps each supported CI type to a comma-separated list of possible corresponding values. For details, see "discriminator.properties File" on page 184.
- ▶ **Replication_config.txt.** This file contains a comma-separated list of CI and relationship types whose property conditions are supported by the replication plugin. For details, see "replication_config.txt File" on page 186.
- ▶ **Fixed_values.txt.** This file enables you to configure fixed values for specific attributes of certain CITs. For details, see "fixed_values.txt File" on page 186.

Hibernate Configuration

- ▶ **persistence.xml.** Used to override out-of-the-box Hibernate configurations. For details, see "persistence.xml File" on page 182.

Simple Configuration

- ▶ **simplifiedConfiguration.xml.** Configuration file that replaces **orm.xml**, **transformations.txt**, and **reconciliation_rules.txt** with less capabilities. For details, see "simplifiedConfiguration.xml File" on page 163.

Adapter.conf File

This file contains the following settings:

- ▶ **use.simplified.xml.config=false.** **true:** uses **simplifiedConfiguration.xml**.

Note: Usage of this file means that **orm.xml**, **transformations.txt**, and **reconciliation_rules.txt** are replaced with fewer capabilities.

- ▶ **dal.ids.chunk.size=300.** Do not change this value.

- ▶ **dal.use.persistence.xml=false. true:** the adapter reads the Hibernate configuration from persistence.xml.

Note: It is not recommended to override the Hibernate configuration.

- ▶ **performance.memory.id.filtering=true.** When the GDBA executes TQLS, in some cases, a large number of IDs may be retrieved and sent back to the database using SQL. To avoid this excessive work and improve performance, the GDBA attempts to read the entire view/table and filters the results in-memory.
- ▶ **id.reconciliation.cmdb.id.type=string/bytes.** When mapping the Generic DB adapter using ID Reconciliation (for information, see step 1 in "How to Implement the Mapping Engine" on page 245), you can either map the **cmdb_id** to a **string** or **bytes/raw** column type by changing the **META-INF/ adapter.conf** property.

simplifiedConfiguration.xml File

This file is used for simple mapping of BSM classes to database tables. To access the template for editing the file, navigate to **Admin > RTSM Administration > Adapter Management > db-adapter > Configuration files.**

This section includes the following topics:

- ▶ "simplifiedConfiguration.xml File Template" on page 164
- ▶ "Limitations" on page 166

simplifiedConfiguration.xml File Template

The **CMDB-class-name** property is the multinode type (the node to which federated CITs connect in the TQL):

```
<?xml version="1.0" encoding="UTF-8"?>
<generic-DB-adapter-config xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="META-CONF/simplifiedConfiguration.xsd">
  <CMDB-class CMDB-class-name="node" default-table-name="[table_name]">
    <primary-key column-name="[column_name]"/>
  </CMDB-class>
</generic-DB-adapter-config>
```

reconciliation-by-two-nodes. Reconciliation can be done using one node or two nodes. In this case example, reconciliation uses two nodes.

connected-node-CMDB-class-name. The second class type needed in the reconciliation TQL.

CMDB-link-type. The relationship type needed in the reconciliation TQL.

link-direction. The direction of the relationship in the reconciliation TQL (from node to ip_address or from ip_address to node):

```
<reconciliation-by-two-nodes connected-node-CMDB-class-name="ip_address"
CMDB-link-type="containment" link-direction="main-to-connected">
```

The reconciliation expression is in the form of ORs and each OR includes ANDs.

is-ordered. Determines if reconciliation is done in order form or by a regular OR comparison.

```
<or is-ordered="true">
```

If the reconciliation property is retrieved from the main class (the multinode), use the **attribute** tag, otherwise use the **connected-node-attribute** tag.

ignore-case. true: when data in the BSM class model is compared with data in the RDBMS, case does not matter:

```
<attribute CMDB-attribute-name="name"
column-name="[column_name]" ignore-case="true"/>
```

The column name is the name of the foreign key column (the column with values that point to the multinode primary key column).

If the multinode primary key column is composed of several columns, there needs to be several foreign key columns, one for each primary key column.

```
<foreign-primary-key column-name="[column_name]"
CMDB-class-primary-key-column="[column_name]"/>
```

If there are few primary key columns, duplicate this column.

```
<primary-key column-name="[column_name]"/>
```

The **from-CMDB-converter** and **to-CMDB-converter** properties are Java classes that implement the following interfaces:

- com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.FcldbDalTransformerFromExternalDB
- com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.FcldbDalTransformerToExternalDB

Use these converters if the value in the RTSM and in the database are not the same.

In this example, `GenericEnumTransformer` is used to convert the enumerator according to the XML file that is written inside the parenthesis (**generic-enum-transformer-example.xml**):

```
<attribute CMDB-attribute-name="[CMDB_attribute_name]"
column-name="[column_name]"
from-CMDB-converter="com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.i
mpl.GenericEnumTransformer(generic-enum-transformer-example.xml)"
to-CMDB-converter="com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.
GenericEnumTransformer(generic-enum-transformer-example.xml)"/>
<attribute CMDB-attribute-name="[CMDB_attribute_name]"
column-name="[column_name]"/>
<attribute CMDB-attribute-name="[CMDB_attribute_name]"
column-name="[column_name]"/>
</class>
</generic-DB-adapter-config>
```

Limitations

- ▶ Can be used to map only TQL queries containing one node (in the database source). For example, you can run a `node > ticket` and a `ticket TQL` query. To bring the hierarchy of nodes from the database, you must use the advanced **orm.xml** file.
- ▶ Only one-to-many relations are supported. For example, you can bring one or more tickets on each node. You cannot bring tickets that belong to more than one node.
- ▶ You cannot connect the same class to different types of RTSM CITs. For example, if you define that `ticket` is connected to `node`, it cannot be connected to `application` as well.

orm.xml File

This file is used for mapping RTSM CITs to database tables.

A template to use for creating a new file is located in the **<HP BSM root directory>\odb\runtime\fcmbd\CodeBase\GenericDBAdapter\META-INF** directory.

To edit the XML file for a deployed adapter, navigate to **Admin > RTSM Administration > Adapter Management > db-adapter > Configuration files**.

This section includes the following topics:

- "orm.xml File Template" on page 167
- "Multiple ORM files" on page 171
- "Naming Conventions" on page 171
- "Using Inline SQL Statements Instead of Table Names" on page 172
- "orm.xml Schema" on page 173

orm.xml File Template

```
<?xml version="1.0" encoding="UTF-8"?>
<entity-mappings xmlns="http://java.sun.com/xml/ns/persistence/orm" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance" version="1.0" xsi:schemaLocation="http://
java.sun.com/xml/ns/persistence/orm http://java.sun.com/xml/ns/persistence/
orm_1_0.xsd">
  <description>Generic DB adapter orm</description>
```

Do not change the package name.

```
<package>generic_db_adapter</package>
```

entity. The RTSM CIT name. This is the multinode entity.

Make sure that **class** includes a **generic_db_adapter.** prefix.

```
<entity class="generic_db_adapter.node">
  <table name="[table_name]"/>
```

Use a secondary table if the entity is mapped to more than one table.

```
<secondary-table name=""/>
<attributes>
```

For a single table inheritance with discriminator, use the following code:

```
<inheritance strategy="SINGLE_TABLE"/>  
<discriminator-value>node</discriminator-value>  
<discriminator-column name="[column_name]"/>
```

Attributes with tag **id** are the primary key columns. Make sure that the naming convention for these primary key columns are **idX** (id1, id2, and so on) where **X** is the column index in the primary key.

```
<id name="id1">
```

Change only the column name of the primary key.

```
<column updatable="false" insertable="false" name="[column_name]"/>  
<generated-value strategy="TABLE"/>  
</id>
```

basic. Used to declare the RTSM attributes. Make sure to edit only **name** and **column_name** properties.

```
<basic name="name">  
<column updatable="false" insertable="false" name="[column_name]"/>  
</basic>
```


For a single table inheritance with discriminator, map the extending classes as follows:

```

<entity name="[cmdb_class_name]" class="generic_db_adapter.nt" name="nt">
  <discriminator-value>nt</discriminator-value>
  <attributes/>
</entity>
<entity class="generic_db_adapter.unix" name="unix">
  <discriminator-value>unix</discriminator-value>
  <attributes/>
</entity>
<entity name="[CMDB_class_name]"
class="generic_db_adapter.[CMDB[cmdb_class_name]]">
  <table name="[default_table_name]"/>
  <secondary-table name=""/>
  <attributes>
    <id name="id1">
      <column updatable="false" insertable="false" name="[column_name]"/>
      <generated-value strategy="TABLE"/>
    </id>
    <id name="id2">
      <column updatable="false" insertable="false" name="[column_name]"/>
      <generated-value strategy="TABLE"/>
    </id>
    <id name="id3">
      <column updatable="false" insertable="false" name="[column_name]"/>
      <generated-value strategy="TABLE"/>
    </id>
  </attributes>
</entity>

```

The following example shows a RTSM attribute name with no prefix:

```

<basic name="[CMDB_attribute_name]">
  <column updatable="false" insertable="false" name="[column_name]"/>
</basic>
<basic name="[CMDB_attribute_name]">
  <column updatable="false" insertable="false" name="[column_name]"/>
</basic>
<basic name="[CMDB_attribute_name]">
  <column updatable="false" insertable="false" name="[column_name]"/>
</basic>
</attributes>
</entity>

```

This is a relationship entity. The naming convention is **end1Type_linkType_end2Type**. In this example **end1Type** is **node** and the **linkType** is **composition**.

```
<entity name="node_composition_[CMDB_class_name]"
class="generic_db_adapter.node_composition_[CMDB_class_name]">
  <table name="[default_table_name]"/>
  <attributes>
    <id name="id1">
      <column updatable="false" insertable="false" name="[column_name]"/>
      <generated-value strategy="TABLE"/>
    </id>
```

The target entity is the entity that this property is pointing to. In this example, **end1** is mapped to **node** entity.

many-to-one. Many relationships can be connected to one node.

join-column. The column that contains **end1** IDs (the target entity IDs).

referenced-column-name. The column name in the target entity (**node**) that contain the IDs that are used in the join column.

```
<many-to-one target-entity="node" name="end1">
  <join-column updatable="false" insertable="false"
referenced-column-name="[column_name]" name="[column_name]"/>
</many-to-one>
```

one-to-one. One relationship can be connected to one **[CMDB_class_name]**.

```
<one-to-one target-entity="[CMDB_class_name]" name="end2">
  <join-column updatable="false" insertable="false"
referenced-column-name="" name="[column_name]"/>
</one-to-one>
</attributes>
</entity>
</entity-mappings>
```

node attribute. This is an example of how to add a node attribute.

```
<entity class="generic_db_adapter.host_node">
  <discriminator-value>host_node</discriminator-value>
  <attributes/>
</entity>
<entity class="generic_db_adapter.nt">
  <discriminator-value>nt</discriminator-value>
  <attributes>
    <basic name="nt_servicepack">
      <column updatable="false" insertable="false" name="specific_type_value"/>
    </basic>
  </attributes>
</entity>
```

Multiple ORM files

Multiple mapping files are supported. Each mapping file name should end with **orm.xml**. All mapping files should be placed under the META-INF folder of the adapter.

Naming Conventions

- ▶ In each entity, the class property must match the name property with the prefix of **generic_db_adapter**.
- ▶ Primary key columns must take names of the form **idX** where **X = 1, 2, ...**, according to the number of primary keys in the table.
- ▶ Attribute names must match class attribute names even as regards case.
- ▶ The relationship name takes the form **end1Type_linkType_end2Type**.
- ▶ RTSM CITs, which are also reserved words in Java, should be prefixed by **gdba_**. For example, for the RTSM CIT **goto**, the ORM entity should be named **gdba_goto**.

Using Inline SQL Statements Instead of Table Names

You can map entities to inline `select` clauses instead of to database tables. This is equivalent to defining a view in the database and mapping an entity to this view. For example:

```
<entity class="generic_db_adapter.node">
  <table name="(select d.id as id1, d.name as name , d.os as host_os from
Device d)"/>
```

In this example, the `node` attributes should be mapped to columns `id1`, `name`, and `host_os`, rather than `id`, `name`, and `os`.

The following limitations apply:

- ▶ The inline SQL statement is available only when using Hibernate as the JPA provider.
- ▶ Round brackets around the inline SQL select clause are mandatory.
- ▶ The `<schema>` element should not be present in the `orm.xml` file. In the case of Microsoft SQL Server 2005, this means that all table names should be prefixed with `dbo.`, rather than defining them globally by `<schema>dbo</schema>`.

orm.xml Schema

The following table explains the common elements of the **orm.xml** file. The complete schema can be found at http://java.sun.com/xml/ns/persistence/orm_1_0.xsd. The list is not complete, and it mostly explains the specific behavior of the standard Java Persistence API for the Generic Database Adapter.

Element		Attributes
Name and Path	Description	
entity-mappings	The root element for the entity mapping document. This element should be exactly the same as the one given in the GDBA sample files.	
description (entity-mappings)	A free text description of the entity mapping document. Optional.	

Element		Attributes
Name and Path	Description	
package (entity-mappings)	The name of the Java package that will contain the mapping classes. Should always contain the text <code>generic_db_adapter</code> .	<p>Name. name</p> <p>Description. The name of the UCMDDB CI type to which this entity is mapped. If this is entity is mapped to a link in the CMDB, the name of the entity should be in the format <code><end_1>_<link_name>_<end_2></code>. For example, <code>node_composition_cpu</code> defines an entity that will be mapped to the composition link between a node and a CPU. If the name of the CI type is the same as the name of the Java class without the package prefix, this field can be omitted.</p> <p>Is required. Optional</p> <p>Type. String</p>
		<p>Name. class</p> <p>Description. The fully qualified name of the Java class that will be created for this DB entity. The name of the Java class' package should be the same as the name given in the <code>package</code> element. You may not use Java reserved words, such as <code>interface</code> or <code>switch</code>, as the class name. Instead, add the prefix <code>gdba_</code> to the name (so <code>interface</code> will be <code>generic_db_adapter.gdba_interface</code>).</p> <p>Is required. Required</p> <p>Type. String</p>

Element		Attributes
Name and Path	Description	
table (entity-mappings > entity)	This element defines the primary table of the DB entity. Can only appear once. Required.	<p>Name. name</p> <p>Description. The name of the primary table. If the name of the table does not contain the schema to which it belongs, the table will be searched only in the schema of the user that was used to create the integration point. This can also be any a valid SELECT statement. If this is a SELECT statement, it must be encapsulated with parentheses.</p> <p>Is required. Required</p> <p>Type. String</p>
secondary-table (entity-mappings > entity)	This element may be used to define a secondary table for the DB entity. This table must be connected to the primary table with a 1-to-1 relationship. You may define more than one secondary table. Optional.	<p>Name. name</p> <p>Description. The name of the secondary table. If the name of the table does not contain the schema to which it belongs, the table will be searched only in the schema of the user that was used to create the integration point. This can also be any a valid SELECT statement. If this is a SELECT statement, it must be encapsulated with parentheses.</p> <p>Is required. Required</p> <p>Type. String</p>
primary-key-join-column (entity-mappings > entity > secondary-table)	If the secondary table and primary table are not connected using fields with the same name, this element defines the name of the primary key field in the secondary table that needs to be connected to the primary key field of the primary table.	<p>Name. name</p> <p>Description. The name of the primary key field in the secondary table. If this element does not exist, it is assumed that the primary key field has the same name as the primary key field of the primary table.</p> <p>Is required. Optional</p> <p>Type. String</p>

Element		Attributes
Name and Path	Description	
inheritance (entity-mappings > entity)	If the current entity is the parent entity for a family of DB entities, then use this element to mark it as such. Optional.	<p>Name. strategy</p> <p>Description. Defines the way the inheritance is implemented in your DB.</p> <p>Is required. Required</p> <p>Type. One of the following values:</p> <ul style="list-style-type: none"> ▶ SINGLE_TABLE – This entity and all child entities exist in the same table. ▶ JOINED – The child entities are in joined tables. ▶ TABLE_PER_CLASS – Each entity is completely defined by a separate table.
discriminator-column (entity-mappings > entity)	If the inheritance is of type SINGLE_TABLE, this element is used to define the name of the field used to determine the type of entity for each row.	<p>Name. name</p> <p>Description. The name of the discriminator column.</p> <p>Is required. Required</p> <p>Type. String</p>
discriminator-value (entity-mappings > entity)	This element defines the type of the specific entity in the inheritance tree. This name needs to be the same as the name defined in the discriminator.properties file for the value group of this specific entity type.	
attributes (entity-mappings > entity)	The root element for all of the attribute mappings for an entity.	

Element		Attributes
Name and Path	Description	
id (entity-mappings > entity attributes)	This element defines the key field for the entity. There must be at least one id field defined. If more than one id element exists, its fields create a compound key for the entity. You should try and avoid compound keys for CI entities (not for links).	<p>Name. name</p> <p>Description. A string of type idX, where X is a number between 1 and 9. The first id should be marked as id1, the second as id2 and so on. This is NOT the name of the key attribute in UCMDB.</p> <p>Is required. Required</p> <p>Type. String</p>
basic (entity-mappings > entity attributes)	This element defines a mapping between a field in the table, which is not part of the table's primary key, and a UCMDB attribute.	<p>Name. name</p> <p>Description. The name of the UCMDB attribute to which the field is mapped. This attribute must exist in the UCMDB CI type to which the current entity is mapped.</p> <p>Is required. Required</p> <p>Type. String</p>
column (entity-mappings > entity > attributes > id -OR- (entity-mappings > entity > attributes > basic)	Defines the name of the column in the table for basic mapping or an id field.	<p>Name. name</p> <p>Description. The name of the field.</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. table</p> <p>Description. The name of the table to which the field belongs. This must be either the primary table or one of the secondary tables defined for the entity. If this attribute is omitted, it is assumed that the field belongs to the primary table.</p> <p>Is required. Optional</p> <p>Type. String</p>

Element		Attributes
Name and Path	Description	
one-to-one (entity-mappings > entity > attributes)	Defines a column whose value is in another table, and the two tables are connected using a one-to-one relationship. This element is only supported for link entity mappings and not for other CI types. This is the only way to define a mapping between a table and a UCMDB link.	<p>Name. name</p> <p>Description. Which of the two ends this field represents.</p> <p>Is required. Required</p> <p>Type. Either end1 or end2</p>
		<p>Name. target-entity</p> <p>Description. The name of the entity to which the end refers.</p> <p>Is required. Required</p> <p>Type. One of the entity names defined in the entity mapping document</p>
join-column (entity-mappings > entity attributes > one-to-one)	Defines the way to join the target-entity defined in the parent one-to-one element and the current entity.	<p>Name. name</p> <p>Description. The name of the field in the current table that will be used to perform the one-to-one join.</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. name</p> <p>Description. The name of a field in the joint entity by which to perform the join. If this attribute is omitted, it is assumed that the joint table has a column with the same name as the field defined in the name attribute.</p> <p>Is required. Optional</p> <p>Type. String</p>

reconciliation_types.txt file

This file is used to configure the reconciliation types.

Each row in the file represents a RTSM CIT that is connected to a federated database CIT in the TQL query.

reconciliation_rules.txt File (for backwards compatibility)

This file is used to configure the reconciliation rules if you want to perform reconciliation when the DBMappingEngine is configured in the adapter. If you do not use the DBMappingEngine, the generic RTSM reconciliation mechanism is used and there is no need to configure this file.

Each row in the file represents a rule. For example:

```
multinode[node] expression[^node.name OR ip_address.name] end1_type[node]
end2_type[ip_address] link_type[containment]
```

The multinode is filled with the multinode name (the RTSM CIT that is connected to the federated database CIT in the TQL query).

This expression includes the logic that decides whether two multinodes are equal (one multinode in the RTSM and the other in the database source).

The expression is composed of ORs or ANDs.

The convention regarding attribute names in the expression part is [className].[attributeName]. For example, attributeName in the ip_address class is written ip_address.name.

For an ordered match (if the first OR sub-expression returns an answer that the multinodes are not equal, the second OR sub-expression is not compared), then use ordered expression instead of expression.

To ignore case during a comparison, use the control (^) sign.

The parameters end1_type, end2_type and link_type are used only if the reconciliation TQL query contains two nodes and not just a multinode. In this case, the reconciliation TQL query is end1_type > (link_type) > end2_type.

There is no need to add the relevant layout as it is taken from the expression.

Types of Reconciliation Rules

Reconciliation rules take the form of OR and AND conditions. You can define these rules on several different nodes (for example, node is identified by name from node AND/OR name from ip_address).

The following options find a match:

- ▶ **Ordered match.** The reconciliation expression is read from left to right. Two OR sub-expressions are considered equal if they have values and they are equal. Two OR sub-expressions are considered not equal if both have values and they are not equal. For any other case there is no decision, and the next OR sub-expression is tested for equality.

name from node OR from ip_address. If both the RTSM and the data source include name and they are equal, the nodes are considered as equal. If both have name but they are not equal, the nodes are considered not equal without testing the name of ip_address. If either the RTSM or the data source is missing name of node, the name of ip_address is checked.

- ▶ **Regular match.** If there is equality in one of the OR sub-expressions, the RTSM and the data source are considered equal.

name from node OR from ip_address. If there is no match on name of node, name of ip_address is checked for equality.

For complex reconciliations, where the reconciliation entity is modeled in the class model as several CITs with relationships (such as node), the mapping of a superset node includes all relevant attributes from all modeled CITs.

Note: As a result, there is a limitation that all reconciliation attributes in the data source should reside in tables that share the same primary key.

Another limitation states that the reconciliation TQL query should have no more than two nodes. For example, the `node > ticket` TQL query has a node in the RTSM and a ticket in the data source.

To reconcile the results, `name` must be retrieved from the node and/or `ip_address`.

If the name in the RTSM is in the format of `*.m.com`, a converter can be used from RTSM to the federated database, and vice versa, to convert these values.

The `node_id` column in the database ticket table is used to connect between the entities (the defined association can also be made in a node table):

DB Node		DB IP_Address	
PK	node_id	PK	ip_id
	name		name

DB Ticket	
PK	ticket_id
	node_id

Note: The three tables must be part of the federated RDBMS source and not the RTSM database.



transformations.txt File

This file contains all the converter definitions.

The format is that each line contains a new definition.

transformations.txt File Template

```
entity[[CMDB_class_name]] attribute[[CMDB_attribute_name]]
to_DB_class[com.mercury.topaz.fcldb.adapters.dbAdapter.dal.
transform.impl.GenericEnumTransformer(generic-enum-transformer-example.xml)]
from_DB_class[com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.impl.
GenericEnumTransformer(generic-enum-transformer-example.xml)]
```

entity. The entity name as it appears in the orm.xml file.

attribute. The attribute name as it appears in the orm.xml file.

to_DB_class. The full, qualified name of a class that implements the interface

com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.FcldbDalTransformerToExternalDB. The elements in the parenthesis are given to this class constructor. Use this converter to transform RTSM values to database values, for example, to append the suffix of **.com** to each node name.

from_DB_class. The full, qualified name of a class that implements the

com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.

FcldbDalTransformerFromExternalDB interface. The elements in the parenthesis are given to this class constructor. Use this converter to transform database values to RTSM values, for example, to append the suffix of **.com** to each node name.

For details, see "Out-of-the-Box Converters" on page 186.

persistence.xml File

This file is used to override the default Hibernate settings and to add support for database types that are not out of the box (OOB database types are Oracle Server, Microsoft MSSQL Server, and MySQL).

If you need to support a new database type, make sure that you supply a connection pool provider (the default is c3p0) and a JDBC driver for your database (put the *.jar files in the adapter folder).

To see all available Hibernate values that can be changed, check the **org.hibernate.cfg.Environment** class (for details, refer to <http://www.hibernate.org>.)

Example of the persistence.xml File:

```

<persistence xmlns="http://java.sun.com/xml/ns/persistence" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://java.sun.com/
xml/ns/persistence
http://java.sun.com/xml/ns/persistence/persistence_1_0.xsd" version="1.0">
  <!-- Don't change this value -->
  <persistence-unit name="GenericDBAdapter">
    <properties>
      <!-- Don't change this value -->
      <property name="hibernate.archive.autodetection" value="class, hbm"/>
      <!--The driver class name"/-->
      <property name="hibernate.connection.driver_class"
value="com.mercury.jdbc.MercOracleDriver"/>
      <!--The connection url"/-->
      <property name="hibernate.connection.url" value="jdbc:mercury:oracle://
artist:1521;sid=cmdb2"/>
      <!--DB login credentials"/-->
      <property name="hibernate.connection.username" value="CMDB"/>
      <property name="hibernate.connection.password" value="CMDB"/>
      <!--connection pool properties"/-->
      <property name="hibernate.c3p0.min_size" value="5"/>
      <property name="hibernate.c3p0.max_size" value="20"/>
      <property name="hibernate.c3p0.timeout" value="300"/>
      <property name="hibernate.c3p0.max_statements" value="50"/>
      <property name="hibernate.c3p0.idle_test_period" value="3000"/>
      <!--The dialect to use-->
      <property name="hibernate.dialect"
value="org.hibernate.dialect.OracleDialect"/>
    </properties>
  </persistence-unit>
</persistence>

```

discriminator.properties File

This file maps each supported CI type (that is also used as a discriminator value in `orm.xml`) to a comma-separated list of possible corresponding values of the discriminator column, or a condition to match possible values of the discriminator column.

If a condition is used, use the syntax: `like(condition)`, where `condition` is a string that can contain the following wildcards:

- ▶ `%` (percent sign) - allows you to match any string of any length (including a zero length string)
- ▶ `_` (underscore) - allows you to match a single character

For example, `like(%unix%)` will match `unix`, `linux`, `unix-aix`, and so on. Like conditions may only be applied to string columns.

You can also have a single discriminator value mapped to any value that does not belong to another discriminator by stating `'all-other'`.

If the adapter you are creating uses discriminator capabilities, you must define all the discriminator values in the **discriminator.properties** file.

Example of Discriminator Mapping:

For example, the adapter supports the CI types `node`, `nt`, and `unix`, and the database contains a single table named `t_nodes` that contains a column called **type**. If the type is `10001`, the row represents a node; if the type is `10004`, it represents a unix machine, and so on. The **discriminator.properties** file might look like this:

```
node=10001, 10005
nt=10002,10003
unix=2%
mainframe=all-other
```


The `orm.xml` file includes the following code:

```
<entity class="generic_db_adapter.node" >
  <table name="t_nodes"/>
  ...
  <inheritance strategy="SINGLE_TABLE"/>
  <discriminator-value>node</discriminator-value>
  <discriminator-column name="type"/>
  ...
</entity>
<entity class="generic_db_adapter.nt" name="nt">
  <discriminator-value>nt</discriminator-value>
  <attributes/>
</entity>
<entity class="generic_db_adapter.unix" name="unix">
  <discriminator-value>unix</discriminator-value>
  <attributes/>
</entity>
```

The `discriminator_column` attribute is then calculated as follows:

- ▶ If **type** contains 10002 or 10003 for a certain entry, the entry is mapped to the **nt** CIT.
- ▶ If **type** contains 10001 or 10005 for a certain entry, the entry is mapped to the **node** CIT.
- ▶ If **type** starts with 2 for a certain entry, the entry is mapped to the **unix** CIT.
- ▶ Any other value in the **type** column is mapped to the **mainframe** CIT.

Note: The **node** CIT is also the parent of **nt** and **unix**.

replication_config.txt File

This file contains a comma-separated list of CI and relationship types whose property conditions are supported by the replication plugin. For details, see "Plugins" on page 191.

fixed_values.txt File

This file enables you to configure fixed values for specific attributes of certain CITS. In this way, each of these attributes can be assigned a fixed value that is not stored in the database.

The file should contain zero or more entries of the following format:

```
entity[<entityName>] attribute[<attributeName>] value[<value>]
```

For example:

```
entity[ip_address] attribute[ip_domain] value[DefaultDomain]
```

The file also supports a list of constants. To define a constants list, use the following syntax:

```
entity[<entityName>] attribute[<attributeName>] value[{<Val1>, <Val2>, <Val3>, ... }]
```

Out-of-the-Box Converters

You can use the following converters (transformers) to convert federated queries and replication jobs to and from database data.

This section includes the following topics:

- ▶ "enum-transformer Converter" on page 187
- ▶ "SuffixTransformer Converter" on page 189
- ▶ "PrefixTransformer Converter" on page 190
- ▶ "BytesToStringTransformer Converter" on page 190

enum-transformer Converter

This converter uses an XML file that is given as an input parameter.

The XML file maps between hard-coded RTSM values and database values (enums). If one of the values does not exist, you can choose to return the same value, return null, or throw an exception.

The transformer performs a comparison between two strings using a case sensitive, or a case insensitive method. The default behavior is case sensitive. To define it as case insensitive use: `case-sensitive="false"` in the `enum-transformer` element.

Use one XML mapping file for each entity attribute.

Note: This converter can be used for both the `to_DB_class` and `from_DB_class` fields in the `transformations.txt` file.

Input File XSD:

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" elementFormDefault="qualified"
attributeFormDefault="unqualified">
  <xs:element name="enum-transformer">
    <xs:complexType>
      <xs:sequence>
        <xs:element ref="value" minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
      <xs:attribute name="db-type" use="required">
        <xs:simpleType>
          <xs:restriction base="xs:string">
            <xs:enumeration value="integer"/>
            <xs:enumeration value="long"/>
            <xs:enumeration value="float"/>
            <xs:enumeration value="double"/>
            <xs:enumeration value="boolean"/>
            <xs:enumeration value="string"/>
            <xs:enumeration value="date"/>
            <xs:enumeration value="xml"/>
            <xs:enumeration value="bytes"/>
          </xs:restriction>
        </xs:simpleType>
      </xs:attribute>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

```

</xs:attribute>
<xs:attribute name="cldb-type" use="required">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="integer"/>
      <xs:enumeration value="long"/>
      <xs:enumeration value="float"/>
      <xs:enumeration value="double"/>
      <xs:enumeration value="boolean"/>
      <xs:enumeration value="string"/>
      <xs:enumeration value="date"/>
      <xs:enumeration value="xml"/>
      <xs:enumeration value="bytes"/>
    </xs:restriction>
  </xs:simpleType>
</xs:attribute>
<xs:attribute name="non-existing-value-action" use="required">
  <xs:simpleType>
    <xs:restriction base="xs:string">
      <xs:enumeration value="return-null"/>
      <xs:enumeration value="return-original"/>
      <xs:enumeration value="throw-exception"/>
    </xs:restriction>
  </xs:simpleType>
</xs:attribute>
<xs:attribute name="case-sensitive" use="optional">
  <xs:simpleType>
    <xs:restriction base="xs:boolean">
    </xs:restriction>
  </xs:simpleType>
</xs:attribute>
</xs:complexType>
</xs:element>
<xs:element name="value">
  <xs:complexType>
    <xs:attribute name="cldb-value" type="xs:string" use="required"/>
    <xs:attribute name="external-db-value" type="xs:string" use="required"/>
    <xs:attribute name="is-cldb-value-null" type="xs:boolean" use="optional"/>
    <xs:attribute name="is-db-value-null" type="xs:boolean" use="optional"/>
  </xs:complexType>
</xs:element>
</xs:schema>

```

Example of Converting 'sys' Value to 'System' Value:

In this example, `sys` value in the RTSM is transformed into `System` value in the federated database, and `System` value in the federated database is transformed into `sys` value in the RTSM.

If the value does not exist in the XML file (for example, the string `demo`), the converter returns the same input value it receives.

```
<enum-transformer CMDB-type="string" DB-type="string"
non-existing-value-action="return-original" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance" xsi:noNamespaceSchemaLocation="..META-CONF/
generic-enum-transformer.xsd">
  <value CMDB-value="sys" external-DB-value="System"/>
</enum-transformer>
```

SuffixTransformer Converter

This converter is used to add or remove suffixes from the RTSM or federated database source value.

There are two implementations:

- **`com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.AdapterToCmdbAddSuffixTransformer`**. Adds the suffix (given as input) when converting from federated database value to RTSM value and removes the suffix when converting from RTSM value to federated database value.
- **`com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.AdapterToCmdbRemoveSuffixTransformer`**. Removes the suffix (given as input) when converting from federated database value to RTSM value and adds the suffix when converting from RTSM value to federated database value.

PrefixTransformer Converter

This converter is used to add or remove a prefix from the RTSM or federated database value.

There are two implementations:

- ▶ **com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.impl.AdapterToCmdbAddPrefixTransformer.** Adds the prefix (given as input) when converting from federated database value to RTSM value and removes the prefix when converting from RTSM value to federated database value.
- ▶ **com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.impl.AdapterToCmdbRemovePrefixTransformer.** Removes the prefix (given as input) when converting from federated database value to RTSM value and adds the prefix when converting from RTSM value to federated database value.

BytesToStringTransformer Converter

This converter is used to convert byte arrays in the RTSM to their string representation in the federated database source.

The converter is:

com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.impl.CmdbToAdapterBytesToStringTransformer.

StringDelimitedListTransformer Converter

This converter is used to transform a single string list to an integer/string list in RTSM.

The converter is:

com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.impl.StringDelimitedListTransformer.

Plugins

The generic database adapter supports the following plugins:

- An optional plugin for full topology synchronization.
- An optional plug-in for synchronizing changes in topology. If no plug-in for synchronizing changes is implemented, it is possible to perform a differential synchronization, but that synchronization will actually be a full one.
- An optional plugin for synchronizing layout.
- An optional plugin to retrieve supported queries for synchronization. If this plugin is not defined, all TQL names are returned.
- An internal, optional plugin to change the TQL definition and TQL result.
- An internal, optional plugin to change a layout request and CIs result.
- An internal, optional plugin to change a layout request and relationships result.

For details about implementing and deploying plugins, see "How to Implement a Plugin" on page 144.

Configuration Examples

This section gives examples of configurations.

This section includes the following topics:

- "Use Case" on page 192
- "Single Node Reconciliation" on page 193
- "Two Node Reconciliation" on page 195
- "Using a Primary Key that Contains More Than One Column" on page 199
- "Using Transformations" on page 201

Use Case

Use case. A TQL query is:

node > (composition) > card

where:

node is the RTSM entity

card is the federated database source entity

composition is the relationship between them

The example is run against the ED database. ED nodes are stored in the Device table and card is stored in the hwCards table. In the following examples, card is always mapped in the same manner.

Single Node Reconciliation

In this example the reconciliation is run against the name property.

Simplified Definition

The reconciliation is done by node and it is emphasized by the special tag **CMDB-class**.

```
<?xml version="1.0" encoding="UTF-8"?>
<generic-DB-adapter-config xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="../META-CONF/simplifiedConfiguration.xsd">
  <CMDB-class CMDB-class-name="node" default-table-name="Device">
    <primary-key column-name="Device_ID"/>
    <reconciliation-by-single-node>
      <or>
        <attribute CMDB-attribute-name="name" column-name="Device_Name"/>
      </or>
    </reconciliation-by-single-node>
  </CMDB-class>
  <class CMDB-class-name="card" default-table-name="hwCards"
connected-CMDB-class-name="node" link-class-name="composition">
    <foreign-primary-key column-name="Device_ID" CMDB-class-primary-key-column="Device_ID"/>
    <primary-key column-name="hwCards_Seq"/>
    <attribute CMDB-attribute-name="card_class" column-name="hwCardClass"/>
    <attribute CMDB-attribute-name="card_vendor" column-name="hwCardVendor"/>
    <attribute CMDB-attribute-name="card_name" column-name="hwCardName"/>
  </class>
</generic-DB-adapter-config>
```

Advanced Definition

orm.xml File

Pay attention to the addition of the relationship mapping. For details, see the definition section in "orm.xml File" on page 166.

Example of the orm.xml File:

```

<?xml version="1.0" encoding="UTF-8"?>
<entity-mappings xmlns="http://java.sun.com/xml/ns/persistence/orm" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://java.sun.com/xml/ns/
persistence/orm http://java.sun.com/xml/ns/persistence/orm_1_0.xsd" version="1.0">
  <description>Generic DB adapter orm</description>
  <package>generic_db_adapter</package>
  <entity class="generic_db_adapter.node" >
    <table name="Device"/>
    <attributes>
      <id name="id1">
        <column name="Device_ID" insertable="false" updatable="false"/>
        <generated-value strategy="TABLE"/>
      </id>
      <basic name="name">
        <column name="Device_Name"/>
      </basic>
    </attributes>
  </entity>
  <entity class="generic_db_adapter.card" >
    <table name="hwCards"/>
    <attributes>
      <id name="id1">
        <column name="hwCards_Seq" insertable="false" updatable="false"/>
        <generated-value strategy="TABLE"/>
      </id>
      <basic name="card_class">
        <column name="hwCardClass" insertable="false" updatable="false"/>
      </basic>
      <basic name="card_vendor">
        <column name="hwCardVendor" insertable="false" updatable="false"/>
      </basic>
      <basic name="card_name">
        <column name="hwCardName" insertable="false" updatable="false"/>
      </basic>
    </attributes>
  </entity>
  <entity class="generic_db_adapter.node_composition_card" >
    <table name="hwCards"/>
    <attributes>
      <id name="id1">
        <column name="hwCards_Seq" insertable="false" updatable="false"/>
        <generated-value strategy="TABLE"/>
      </id>
      <many-to-one name="end1" target-entity="node">

```

```

    <join-column name="Device_ID" insertable="false" updatable="false"/>
  </many-to-one>
  <one-to-one name="end2" target-entity="card">
    <join-column name="hwCards_Seq" referenced-column-name="hwCards_Seq"
insertable="false" updatable="false"/>
  </one-to-one>
</attributes>
</entity>
</entity-mappings>

```

reconciliation_types.txt File

For details, see "reconciliation_types.txt file" on page 179.

```
node
```

reconciliation_rules.txt File

For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.

```
multinode[node] expression[node.name]
```

transformation.txt File

This file remains empty as no values need to be converted in this example.

Two Node Reconciliation

In this example, reconciliation is calculated according to the name property of node and of ip_address with different variations.

The reconciliation TQL query is **node > (containment) > ip_address**.

Simplified Definition

The reconciliation is by name of node OR of ip_address:

```
<?xml version="1.0" encoding="UTF-8"?>
<generic-DB-adapter-config xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="../META-CONF/simplifiedConfiguration.xsd">
  <CMDB-class CMDB-class-name="node" default-table-name="Device">
    <primary-key column-name="Device_ID"/>
    <reconciliation-by-two-nodes connected-node-CMDB-class-name="ip_address"
CMDB-link-type="containment">
      <or>
        <attribute CMDB-attribute-name="name" column-name="Device_Name"/>
        <connected-node-attribute CMDB-attribute-name="name"
column-name="Device_PreferredIPAddress"/>
      </or>
    </reconciliation-by-two-nodes>
  </CMDB-class>
  <class CMDB-class-name="card" default-table-name="hwCards"
connected-CMDB-class-name="node" link-class-name="containment">
    <foreign-primary-key column-name="Device_ID" CMDB-class-primary-key-column="Device_ID"/>
    <primary-key column-name="hwCards_Seq"/>
    <attribute CMDB-attribute-name="card_class" column-name="hwCardClass"/>
    <attribute CMDB-attribute-name="card_vendor" column-name="hwCardVendor"/>
    <attribute CMDB-attribute-name="card_name" column-name="hwCardName"/>
  </class>
</generic-DB-adapter-config>
```

The reconciliation is name of node AND of ip_address:

```
<?xml version="1.0" encoding="UTF-8"?>
<generic-DB-adapter-config xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="..META-CONF/simplifiedConfiguration.xsd">
  <CMDB-class CMDB-class-name="node" default-table-name="Device">
    <primary-key column-name="Device_ID"/>
    <reconciliation-by-two-nodes connected-node-CMDB-class-name="ip_address"
CMDB-link-type="containment">
      <and>
        <attribute CMDB-attribute-name="name" column-name="Device_Name"/>
        <connected-node-attribute CMDB-attribute-name="name"
column-name="Device_PREFERREDIPAddress"/>
      </and>
    </reconciliation-by-two-nodes>
  </CMDB-class>
  <class CMDB-class-name="card" default-table-name="hwCards"
connected-CMDB-class-name="node" link-class-name="containment">
    <foreign-primary-key column-name="Device_ID" CMDB-class-primary-key-column="Device_ID"/>
    <primary-key column-name="hwCards_Seq"/>
    <attribute CMDB-attribute-name="card_class" column-name="hwCardClass"/>
    <attribute CMDB-attribute-name="card_vendor" column-name="hwCardVendor"/>
    <attribute CMDB-attribute-name="card_name" column-name="hwCardName"/>
  </class>
</generic-DB-adapter-config>
```

The reconciliation is by name of `ip_address`:

```
<?xml version="1.0" encoding="UTF-8"?>
<generic-DB-adapter-config xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="..../META-CONF/simplifiedConfiguration.xsd">
  <CMDB-class CMDB-class-name="node" default-table-name="Device">
    <primary-key column-name="Device_ID"/>
    <reconciliation-by-two-nodes connected-node-CMDB-class-name="ip_address"
CMDB-link-type="containment">
      <or>
        <connected-node-attribute CMDB-attribute-name="name"
column-name="Device_PREFERREDIPAddress"/>
      </or>
    </reconciliation-by-two-nodes>
  </CMDB-class>
  <class CMDB-class-name="card" default-table-name="hwCards"
connected-CMDB-class-name="node" link-class-name="containment">
    <foreign-primary-key column-name="Device_ID" CMDB-class-primary-key-column="Device_ID"/>
    <primary-key column-name="hwCards_Seq"/>
    <attribute CMDB-attribute-name="card_class" column-name="hwCardClass"/>
    <attribute CMDB-attribute-name="card_vendor" column-name="hwCardVendor"/>
    <attribute CMDB-attribute-name="card_name" column-name="hwCardName"/>
  </class>
</generic-DB-adapter-config>
```

Advanced Definition

orm.xml File

Since the reconciliation expression is not defined in this file, the same version should be used for any reconciliation expression.

reconciliation_types.txt File

For details, see "reconciliation_types.txt file" on page 179.

node

reconciliation_rules.txt File

For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.

```
multinode[node] expression[ip_address.name OR node.name] end1_type[node]
end2_type[ip_address] link_type[containment]
```

```
multinode[node] expression[ip_address.name AND node.name] end1_type[node]
end2_type[ip_address] link_type[containment]
```

```
multinode[node] expression[ip_address.name] end1_type[node] end2_type[ip_address]
link_type[containment]
```

transformation.txt File

This file remains empty as no values need to be converted in this example.

Using a Primary Key that Contains More Than One Column

If the primary key is composed of more than one column, the following code is added to the XML definitions:

Simplified Definition

There is more than one primary key tag and for each column there is a tag.

```
<class CMDB-class-name="card" default-table-name="hwCards"
connected-CMDB-class-name="node" link-class-name="containment">
  <foreign-primary-key column-name="Device_ID"
CMDB-class-primary-key-column="Device_ID"/>
  <primary-key column-name="Device_ID"/>
  <primary-key column-name="hwBusesSupported_Seq"/>
  <primary-key column-name="hwCards_Seq"/>
  <attribute CMDB-attribute-name="card_class" column-name="hwCardClass"/>
  <attribute CMDB-attribute-name="card_vendor"
column-name="hwCardVendor"/>
  <attribute CMDB-attribute-name="card_name" column-name="hwCardName"/>
</class>
```

Advanced Definition

orm.xml File

A new id entity is added that maps to the primary key columns. Entities that use this id entity must add a special tag.

If you use a foreign key (join-column tag) for such a primary key, you must map between each column in the foreign key to a column in the primary key.

For details, see "orm.xml File" on page 166.

Example of the orm.xml File:

```
< entity class="generic_db_adapter.card" >
  <table name="hwCards"/>
  <attributes>
    <id name="id1">
      <column name="Device_ID" insertable="false" updatable="false"/>
      <generated-value strategy="TABLE"/>
    </id>
    <id name="id2">
      <column name="hwBusesSupported_Seq" insertable="false" updatable="false"/>
      <generated-value strategy="TABLE"/>
    </id>
    <id name="id3">
      <column name="hwCards_Seq" insertable="false" updatable="false"/>
      <generated-value strategy="TABLE"/>
    </id>
  .
  .
  .
<entity class="generic_db_adapter.node_containment_card" >
  <table name="hwCards"/>
  <attributes>
    <id name="id1">
      <column name="Device_ID" insertable="false" updatable="false"/>
      <generated-value strategy="TABLE"/>
    </id>
    <id name="id2">
      <column name="hwBusesSupported_Seq" insertable="false" updatable="false"/>
      <generated-value strategy="TABLE"/>
    </id>
    <id name="id3">
      <column name="hwCards_Seq" insertable="false" updatable="false"/>
    </id>
```



```

    <generated-value strategy="TABLE"/>
  </id>
  <many-to-one name="end1" target-entity="node">
    <join-column name="Device_ID" insertable="false" updatable="false"/>
  </many-to-one>
  <one-to-one name="end2" target-entity="card">
    <join-column name="Device_ID" referenced-column-name="Device_ID" insertable="false"
updatable="false"/>
    <join-column name="hwBusesSupported_Seq"
referenced-column-name="hwBusesSupported_Seq" insertable="false" updatable="false"/>
    <join-column name="hwCards_Seq" referenced-column-name="hwCards_Seq"
insertable="false" updatable="false"/>
  </one-to-one>
</attributes>
</entity>
</entity-mappings>

```

Using Transformations

In the following example, the generic **enum** transformer is converted from values 1, 2, 3 to values a, b, c respectively in the name column.

The mapping file is `generic-enum-transformer-example.xml`.

```

<enum-transformer CMDB-type="string" DB-type="string"
non-existing-value-action="return-original" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance" xsi:noNamespaceSchemaLocation="..META-CONF/
generic-enum-transformer.xsd">
  <value CMDB-value="1" external-DB-value="a"/>
  <value CMDB-value="2" external-DB-value="b"/>
  <value CMDB-value="3" external-DB-value="c"/>
</enum-transformer>

```

Simplified Definition

```

<CMDB-class CMDB-class-name="node" default-table-name="Device">
  <primary-key column-name="Device_ID"/>
  <reconciliation-by-two-nodes connected-node-CMDB-class-name="ip_address"
CMDB-link-type="containment">
    <or>
      <attribute CMDB-attribute-name="name" column-name="Device_Name"
from-CMDB-converter="com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.i
mpl.GenericEnumTransformer(generic-enum-transformer-example.xml)"
to-CMDB-converter="com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.
GenericEnumTransformer(generic-enum-transformer-example.xml)"/>
      <connected-node-attribute CMDB-attribute-name="name"
column-name="Device_PreferredIPAddress"/>
    </or>
  </reconciliation-by-two-nodes>
</CMDB-class>
.
.
.

```

Advanced Definition

There is a change only to the **transformation.txt** file.

transformation.txt File

Make sure that the attribute names and entity names are the same as in the **orm.xml** file.

```

entity[node] attribute[name]
to_DB_class[com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.Generic
EnumTransformer(generic-enum-transformer-example.xml)]
from_DB_class[com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.Gene
ricEnumTransformer(generic-enum-transformer-example.xml)]

```

Adapter Log Files

To understand the calculation flows and adapter lifecycle, and to view debug information, you can consult the following log files.

This section includes the following topics:

- "Log Levels" on page 203
- "Log Locations" on page 204

Log Levels

You can configure the log level for each of the logs.

In a text editor, open the `<HP BSM root directory>\odb\conf\log\fcmdb.gdba.properties` file.

The default log level is **ERROR**:

```
#loglevel can be any of DEBUG INFO WARN ERROR FATAL
loglevel=ERROR
```

- To increase the log level for all log files, change **loglevel=ERROR** to **loglevel=DEBUG** or **loglevel=INFO**.
- To change the log level for a specific file, change the specific **log4j** category line accordingly. For example, to change the log level of `fcmdb.gdba.dal.sql.log` to **INFO**, change

```
log4j.category.fcmdb.gdba.dal.SQL=${loglevel},fcmdb.gdba.dal.SQL.appender
```

to:

```
log4j.category.fcmdb.gdba.dal.SQL=INFO,fcmdb.gdba.dal.SQL.appender
```

Log Locations

The log files are located in the <HP BSM root directory>\odb\runtime\log directory.

► **Fcmdb.gdba.log**

The adapter lifecycle log. Gives details about when the adapter started or stopped, and which CITs are supported by this adapter.

Consult for initiation errors (adapter load/unload).

► **fcmdb.log**

Consult for exceptions.

► **cmdb.log**

Consult for exceptions.

► **Fcmdb.gdba.mapping.engine.log**

The mapping engine log. Gives details about the reconciliation TQL query that the mapping engine uses, and the reconciliation topologies that are compared during the connect phase.

Consult this log when a TQL query gives no results even though you know there are relevant CIs in the database, or the results are unexpected (check the reconciliation).

► **Fcmdb.gdba.TQL.log**

The TQL log. Gives details about the TQL queries and their results.

Consult this log when a TQL query does not return results and the mapping engine log shows that there are no results in the federated data source.

► **Fcmdb.gdba.dal.log**

The DAL lifecycle log. Gives details about CIT generation and database connection details.

Consult this log when you cannot connect to the database or when there are CITs or attributes that are not supported by the query.

► **Fcmdb.gdba.dal.command.log**

The DAL operations log. Gives details about internal DAL operations that are called. (This log is similar to `cmdb.dal.command.log`).

► **Fcmdb.gdba.dal.SQL.log**

The DAL SQL queries log. Gives details about called JPAQLs (object oriented SQL queries) and their results.

Consult this log when you cannot connect to the database or when there are CITs or attributes that are not supported by the query.

► **Fcmdb.gdba.hibernate.log**

The Hibernate log. Gives details about the SQL queries that are run, the parsing of each JPAQL to SQL, the results of the queries, data regarding Hibernate caching, and so on. For details on Hibernate, see "Hibernate as JPA Provider" on page 121.

External References

For details on the JavaBeans 3.0 specification, see <http://jcp.org/aboutJava/communityprocess/final/jsr220/index.html>.

Troubleshooting and Limitations

This section describes troubleshooting and limitations for the generic database adapter.

General Limitations

- SQL Server NTLM authentication is not supported.
- When you update an adapter package, use Notepad++, UltraEdit, or some other third-party text editor rather than Notepad (any version) from Microsoft Corporation to edit the template files. This prevents the use of special symbols, which cause the deployment of the prepared package to fail.

JPA Limitations

- All tables must have a primary key column.
- RTSM class attribute names must follow the JavaBeans naming convention (for example, names must start with lower case letters).
- Two CIs that are connected with one relationship in the class model must have direct association in the database (for example, if `node` is connected to `ticket` there must be a foreign key or linkage table that connects them).
- Several tables that are mapped to the same CIT must share the same primary key table.

Functional Limitations

- You cannot create a manual relationship between the RTSM and federated CITs. To be able to define virtual relationships, a special relationship logic must be defined (it can be based on properties of the federated class).
- Federated CITs cannot be trigger CITs in an impact rule, but they can be included in an impact analysis TQL query.
- A federated CIT can be part of an enrichment TQL, but cannot be used as the node on which enrichment is performed (you cannot add, update, or delete the federated CIT).
- Using a class qualifier in a condition is not supported.
- Subgraphs are not supported.
- Compound relationships are not supported.
- The external CI RTSM id is composed from its primary key and not its key attributes.
- A column of type `bytes` cannot be used as a primary key column in Microsoft SQL Server.
- TQL query calculation fails if attribute conditions that are defined on a federated node have not had their names mapped in the `orm.xml` file.
- The Generic DB Adapter does not support Windows Authentication for SQL Server.

5

Developing Java Adapters

This chapter includes:

Concepts

- ▶ Federation Framework Overview on page 208
- ▶ Adapter and Mapping Interaction with the Federation Framework on page 214
- ▶ Federation Framework Flow for Federated TQL Queries on page 215
- ▶ Federation Framework Flow for Population on page 231
- ▶ Adapter Interfaces on page 233

Tasks

- ▶ How to Debug Adapter Resources on page 236
- ▶ How to Add an Adapter for a New External Data Source on page 236
- ▶ How to Implement the Mapping Engine on page 245
- ▶ How to Create a Sample Adapter on page 247

Reference

- ▶ XML Configuration Tags and Properties on page 249

Concepts

Federation Framework Overview

Note:

- The term **relationship** is equivalent to the term **link**.
 - The term **CI** is equivalent to the term **object**.
 - A graph is a collection of nodes and links.
 - For a glossary of definitions and terms, see "Glossary" in the *RTSM Administration Guide*.
-

The Federation Framework functionality uses an API to retrieve information from federated sources. The Federation Framework provides three main capabilities:

- **Federation on the fly.** All queries are run over original data repositories and results are built on the fly in the RTSM.
- **Population.** Populates data (topological data and CI properties) to the RTSM from an external data source.
- **Data Push.** Pushes data (topological data and CI properties) from the local RTSM to a remote data source.

All action types require an adapter for each data repository, which can provide the specific capabilities of the data repository and retrieve and/or update the required data. Every request to the data repository is made through its adapter.

This section also includes the following topics:

- "Federation on the Fly" on page 209
- "Data Push" on page 211
- "Population" on page 212


Federation on the Fly

Federated TQL queries enables data retrieval from any external data repository without replicating its data.

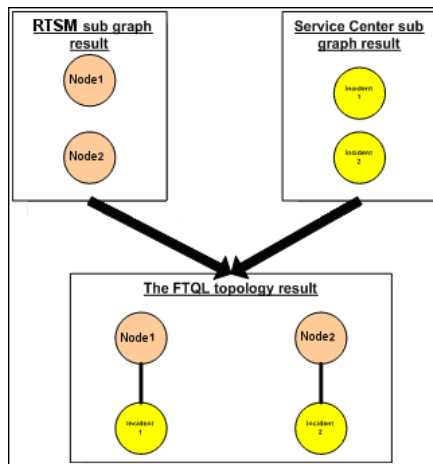
A federated TQL query uses adapters that represent external data repositories, to create appropriate external relationships between CIs from different external data repositories and the BSM CIs.

Example of Federation-on-the-Fly Flow:

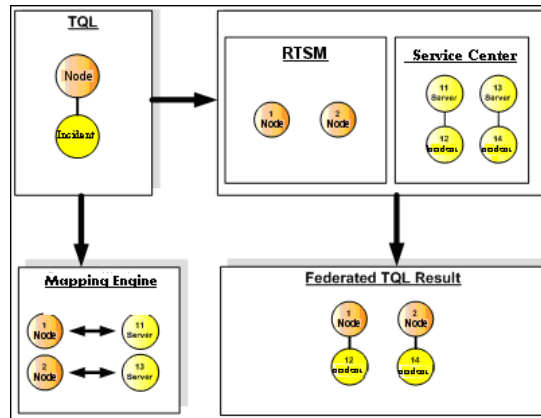
- 1 The Federation Framework splits a federated TQL query into several subgraphs, where all nodes in a subgraph refer to the same data repository. Each subgraph is connected to the other subgraphs by a virtual relationship (but itself contains no virtual relationships).



- 2 After the federated TQL query is split into subgraphs, the Federation Framework calculates each subgraph's topology and connects two appropriate subgraphs by creating virtual relationships between the appropriate nodes.



- After the federated TQL topology is calculated, the Federation Framework retrieves a layout for the topology result.



Data Push

You use the data push flow to synchronize data from your current local RTSM to a remote service or target data repository.

In data push, data repositories are divided into two categories: source (local RTSM) and target. Data is retrieved from the source data repository and updated to the target data repository. The data push process is based on query names, meaning that data is synchronized between the source (local RTSM) and target data repositories, and is retrieved by a TQL query name from the local RTSM.

The data push process flow includes the following steps:

- Retrieving the topology result with signatures from the source data repository.
- Comparing the new results with the previous results.
- Retrieving a full layout (that is, all CI properties) of CIs and relationships, for changed results only.
- Updating the target data repository with the received full layout of CIs and relationships. If any CIs or relationships are deleted in the source data repository and the query is exclusive, the replication process removes the CIs or relationships in the target data repository as well.

The RTSM has 2 hidden data sources (**hiddenRMIDataSource** and **hiddenChangesDataSource**), which are always the 'source' data source in data push flows. To implement a new adapter for data push flows, you only have to implement the 'target' adapter.

Population

You use the population flow to populate the RTSM with data from external sources.

The flow always uses one 'source' data source to retrieve the data, and pushes the retrieved data to the Probe in a similar process to the flow of a discovery job.

To implement a new adapter for population flows, you only have to implement the source adapter, since the Data Flow Probe acts as the target.

The adapter in the population flow is executed on the Probe. Debugging and logging should be done on the Probe and not on the RTSM.

The population flow is based on query names, that is, data is synchronized between the source data repository and the Data Flow Probe, and is retrieved by a query name in the source data repository. For example, in BSM, the query name is the name of the TQL query. However, in another data repository the query name can be a code name that returns data. The adapter is designed to correctly handle the query name.

Each job can be defined as an exclusive job. This means that the CIs and relationships in the job results are unique in the local RTSM, and no other query can bring them to the target. The adapter of the source data repository supports specific queries, and can retrieve the data from this data repository. The adapter of the target data repository enables the update of retrieved data on this data repository.

SourceDataAdapter Flow

- Retrieves the topology result with signatures from the source data repository.
- Compares the new results with the previous results.

- Retrieves a full layout (that is, all CI properties) of CIs and relationships, for changed results only.
- Updates the target data repository with the received full layout of CIs and relationships. If any CIs or relationships are deleted in the source data repository and the query is exclusive, the replication process removes the CIs or relationships in the target data repository as well.

SourceChangesDataAdapter Flow

- Retrieves the topology result that occurred since the last date given.
- Retrieves a full layout (that is, all CI properties) of CIs and relationships, for changed results only.
- Updates the target data repository with the received full layout of CIs and relationships. If any CIs or relationships are deleted in the source data repository and the query is exclusive, the replication process removes the CIs or relationships in the target data repository as well.

PopulateDataAdapter Flow

- Retrieves the full topology with requested layout result.
- Uses the topology chunk mechanism to retrieve the data in chunks.
- The probe filters out any data that was already brought in earlier runs
- Updates the target data repository with the received layout of CIs and relationships. If any CIs or relationships are deleted in the source data repository and the query is exclusive, the replication process removes the CIs or relationships in the target data repository as well.

PopulateChangesDataAdapter Flow

- Retrieves the topology with requested layout result that has changes since the last run.
- Uses the topology chunk mechanism to retrieve the data in chunks.
- The probe filters out any data that was already brought in earlier runs (including this flow).

- ▶ Updates the target data repository with the received layout of CIs and relationships. If any CIs or relationships are deleted in the source data repository and the query is exclusive, the replication process removes the CIs or relationships in the target data repository as well.

Adapter and Mapping Interaction with the Federation Framework

An adapter is an entity in BSM that represents external data (data that is not saved in BSM). In federated flows, all interactions with external data sources are performed through adapters. The Federation Framework interaction flow and adapter interfaces are different for replication and for federated TQL queries.

This section also includes the following topics:

- ▶ "Adapter Lifecycle" on page 214
- ▶ "Adapter assist Methods" on page 215

Adapter Lifecycle

An adapter instance is created for each external data repository. The adapter begins its lifecycle with the first action applied to it (such as, **calculate** TQL or **retrieve/update** data). When the **start** method is called, the adapter receives environmental information, such as the data repository configuration, logger, and so on. The adapter lifecycle ends when the data repository is removed from the configuration, and the **shutdown** method is called. This means that the adapter is stateful and can contain the connection to the external data repository if it is required.

Adapter assist Methods

The adapter has several **assist** methods that can add external data repository configurations. These methods are not part of the adapter lifecycle and create a new adapter each time they are called.

- ▶ The first method tests the connection to the external data repository for a given configuration. `testConnection` can be executed either on the BSM server or the Data Flow Probe, depending on the type of adapter.
- ▶ The second method is relevant only for the source adapter and returns the supported queries for replication. (This method is executed on the Probe only.)
- ▶ The third method is relevant only for federation and population flows, and returns supported external classes by the external data repository. (This method is executed on the BSM server.)

All these methods are used when you create or view integration configurations.

Federation Framework Flow for Federated TQL Queries

This section includes the following topics:

- ▶ "Definitions and Terms" on page 216
- ▶ "Mapping Engine" on page 217
- ▶ "Federated Adapter" on page 217
- ▶ "Flow Diagrams" on page 218

Definitions and Terms

Reconciliation data. The rule for matching CIs of the specified type that are received from the RTSM and the external data repository. The reconciliation rule can be of three types:

- ▶ **ID reconciliation.** This can be used only if the external data repository contains the RTSM ID of reconciliation objects.
- ▶ **Property reconciliation.** This is used when the matching can be done by properties of the reconciliation CI type only.
- ▶ **Topology reconciliation.** This is used when you need the properties of additional CITs (not only of the reconciliation CIT) to perform a match on reconciliation CIs. For example, you can perform reconciliation of the node type by the name property that belongs to the `ip_address` CIT.

Reconciliation object. The object is created by the adapter according to received reconciliation data. This object should refer to an external CI and is used by the Mapping Engine to connect between the external CIs and the RTSM CIs.

Reconciliation CI type. The type of CIs that represent reconciliation objects. These CIs must be stored in both the RTSM and in the external data repositories.

Mapping engine. A component that identifies relations between CIs from different data repositories that have a virtual relationship between them. The identification is performed by reconciling RTSM reconciliation objects and external CI reconciliation objects.

Mapping Engine

Federation Framework uses the Mapping Engine to calculate the federated TQL query. The Mapping Engine connects between CIs that are received from different data repositories and are connected by virtual relationships. The Mapping Engine also provides reconciliation data for the virtual relationship. One end of the virtual relationship must refer to the RTSM. This end is a reconciliation type. For the calculation of the two subgraphs, a virtual relationship can start from any end node.

Federated Adapter

The Federated adapter brings two kinds of data from external data repositories: external CI data and reconciliation objects that belong to external CIs.

- ▶ **External CI data.** The external data that does not exist in the RTSM. It is the target data of the external data repository.
- ▶ **Reconciliation object data.** The auxiliary data that is used by the federation framework to connect between RTSM CIs and external data. Each reconciliation object should refer to an External CI. The type of reconciliation object is the type (or subtype) of one of the virtual relationship ends from which data is retrieved. Reconciliation objects should fit the adapter received to reconciliation data. The reconciliation object can be one of three types: `IdReconciliationObject`, `PropertyReconciliationObject`, or `TopologyReconciliationObject`.

In the `DataAdapter`-based interfaces (`DataAdapter`, `PopulateDataAdapter`, and `PopulateChangesDataAdapter`), the reconciliation is requested as part of the query definition.

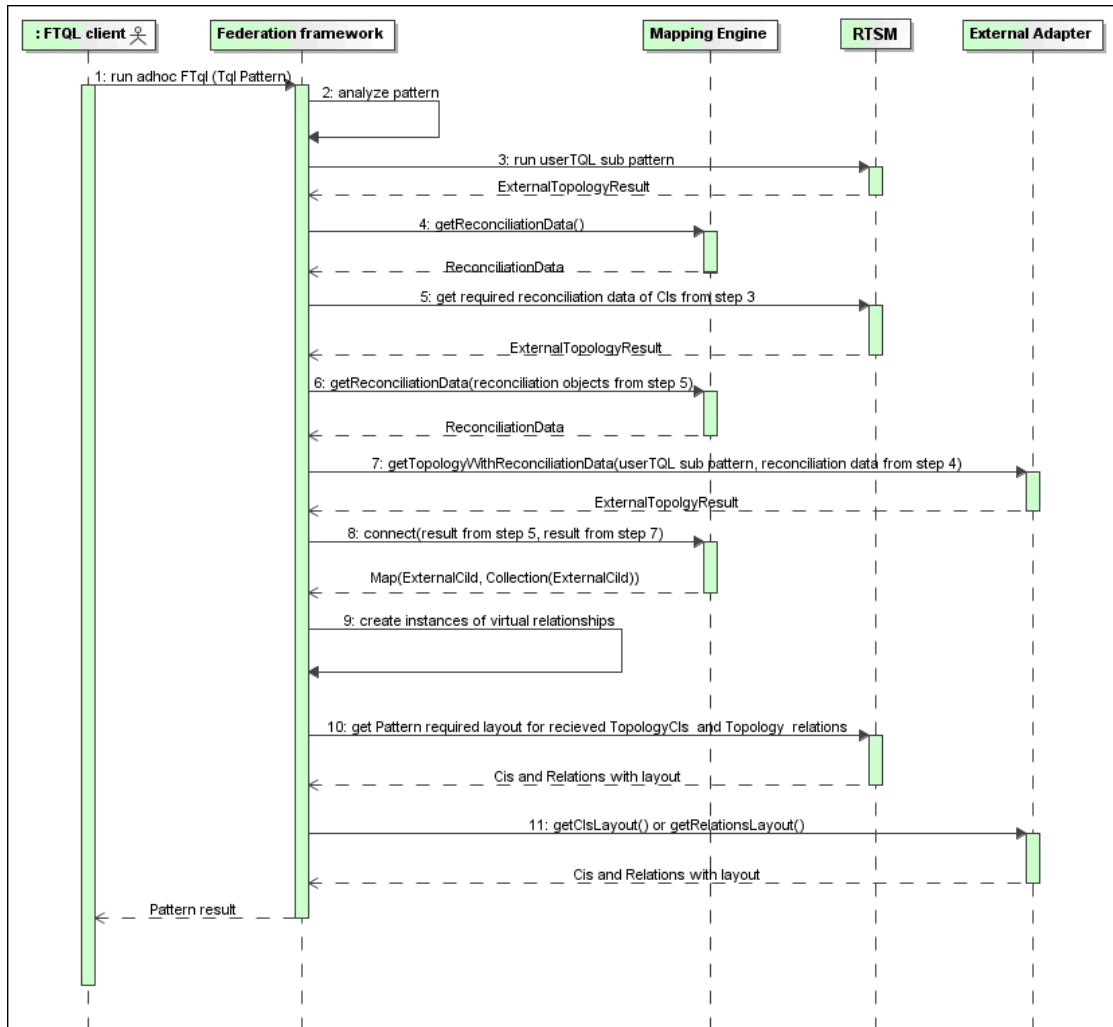
Flow Diagrams

The following diagrams illustrate the interactions between the Federation Framework, BSM, the adapter, and the Mapping Engine. The federated TQL query in the example diagrams has only one virtual relationship, so that only BSM and one external data repository are involved in the federated TQL query.

In the first diagram the calculation begins in BSM and in the second diagram in the external adapter. Each step in the diagram includes references to the appropriate method call of the adapter or mapping engine interface.

Calculation Starts at the Business Service Management End

The following sequence diagram illustrates the interaction between the Federation Framework, BSM, the adapter, and the Mapping Engine. The federated TQL query in the example diagram has only one virtual relationship, so that only BSM and one external data repository are involved in the federated TQL query.

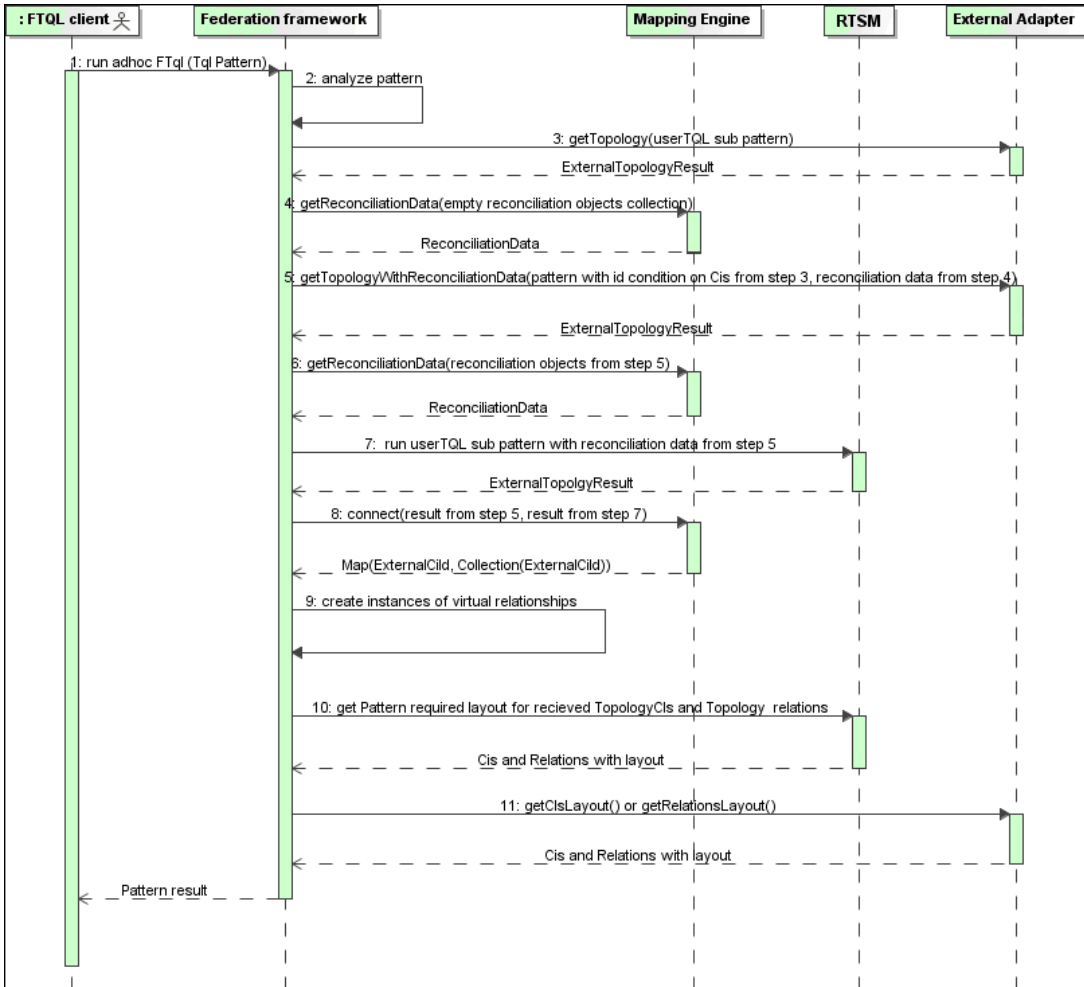


The numbers in this image are explained below:

Number	Explanation
1	The Federation Framework receives a call for a federated TQL calculation.
2	The Federation Framework analyzes the adapter, finds the virtual relationship, and divides the original TQL into two sub-adapters—one for BSM and one for the external data repository.
3	The Federation Framework requests the topology of the sub-TQL from BSM.
4	After receiving the topology results, the Federation Framework calls the appropriate Mapping Engine for the current virtual relationship and requests reconciliation data. The <code>reconciliationObject</code> parameter is empty at this stage, that is, no condition is added to reconciliation data in this call. The returned reconciliation data defines which data is needed to match the reconciliation CIs in BSM to the external data repository. The reconciliation data can be one of the following types: <ul style="list-style-type: none"> ▶ IdReconciliationData. CIs are reconciled according to their ID. ▶ PropertyReconciliationData. CIs are reconciled according to the properties of one of the CIs. ▶ TopologyReconciliationData. CIs are reconciled according to the topology (for example, to reconcile node CIs, the IP address of IP is required too).
5	The Federation Framework requests reconciliation data for the CIs of the virtual relationship ends that were received in step 3 from BSM.
6	The Federation Framework calls the Mapping Engine to retrieve the reconciliation data. In this state (by contrast with step 3), the Mapping Engine receives the reconciliation objects from step 5 as parameters. The Mapping Engine translates the received reconciliation object to the condition on the reconciliation data.
7	The Federation Framework requests the topology of the sub-TQL from the external data repository. The external adapter receives the reconciliation data from step 6 as a parameter.

Number	Explanation
8	The Federation Framework calls the Mapping Engine to connect between the received results. The <code>firstResult</code> parameter is the external topology result received from BSM in step 5 and the <code>secondResult</code> parameter is the external topology result received from the External Adapter in step 7. The Mapping Engine returns a map where External CI ID from the first data repository (BSM in this case) is mapped to the External CI IDs from the second (external) data repository.
9	For each mapping, the Federation Framework creates a virtual relationship.
10	After the calculation of the federated TQL query results (only at the topology stage), the Federation Framework retrieves the original TQL layout for the resulting CIs and relationships from the appropriate data repositories.

Calculation Starts at the External Adapter End



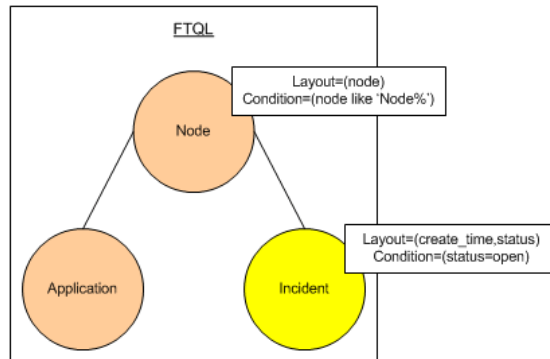
The numbers in this image are explained below:

Number	Explanation
1	The Federation Framework receives a call for an federated TQL calculation.
2	The Federation Framework analyzes the adapter, finds the virtual relationship, and divides the original TQL into two sub-adapters – one for BSM and one for the external data repository.
3	The Federation Framework requests the topology of the sub-TQL from the External Adapter. The returned <code>ExternalTopologyResult</code> is not supposed to contain any reconciliation object, since the reconciliation data is not part of the request.
4	<p>After receiving the topology results, the Federation Framework calls the appropriate Mapping Engine with the current virtual relationship and requests reconciliation data. The <code>reconciliationObjects</code> parameter is empty at this state, that is, no condition is added to the reconciliation data in this call. The returned reconciliation data defines what data is needed to match the reconciliation CIs in BSM to the external data repository. The reconciliation data can be one of three following types:</p> <ul style="list-style-type: none"> ▶ IdReconciliationData. CIs are reconciled according to their ID. ▶ PropertyReconciliationData. CIs are reconciled according to the properties of one of the CIs. ▶ TopologyReconciliationData. CIs are reconciled according to the topology (for example, to reconcile node CIs, the IP address of <code>IP</code> is required too).
5	The Federation Framework requests reconciliation objects for the CIs that were received in step 3 from the external data repository. The Federation Framework calls the <code>getTopologyWithReconciliationData()</code> method in the External Adapter, where the requested topology is a one-node topology with CIs received in step 3 as the ID condition and reconciliation data from step 4.

Number	Explanation
6	The Federation Framework calls the Mapping Engine to retrieve the reconciliation data. In this state (by contrast with step 3), the Mapping Engine receives the reconciliation objects from step 5 as parameters. The Mapping Engine translates the received reconciliation object to the condition on the reconciliation data.
7	The Federation Framework requests the topology of the sub-TQL with reconciliation data from step 6 from BSM.
8	The Federation Framework calls the Mapping Engine to connect between the received results. The <code>firstResult</code> parameter is the external topology result received from the External Adapter at step 5 and the <code>secondResult</code> parameter is the external topology result received from BSM at step 7. The Mapping Engine returns a map where the External CI ID from the first data repository (the external data repository in this case) is mapped to the External CI IDs from the second data repository (BSM).
9	For each mapping, the Federation Framework creates a virtual relationship.
10	After the calculation of the federated TQL query results (only at the topology stage), the Federation Framework retrieves the original TQL layout for the resulting CIs and relationships from the appropriate data repositories.

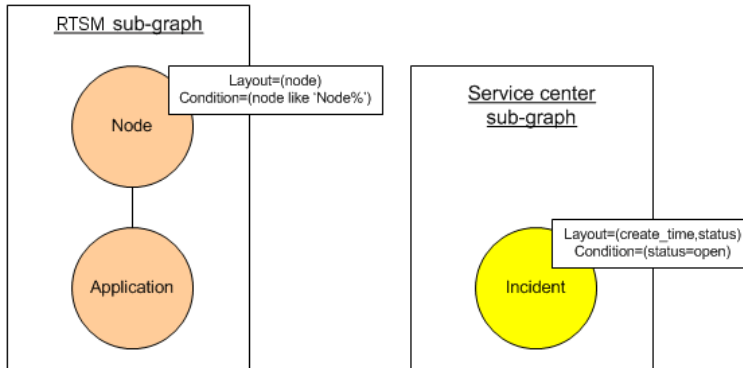
Example of Federation Framework Flow for Federated TQL Queries

This example explains how to view all open incidents on specific nodes. The ServiceCenter data repository is the external data repository. The node instances are stored in BSM, and the incident instances are stored in ServiceCenter. It is assumed that to connect the incident instances to the appropriate node, the `node` and `ip_address` properties of the host and IP are needed. These are reconciliation properties that identify the nodes from ServiceCenter in BSM.

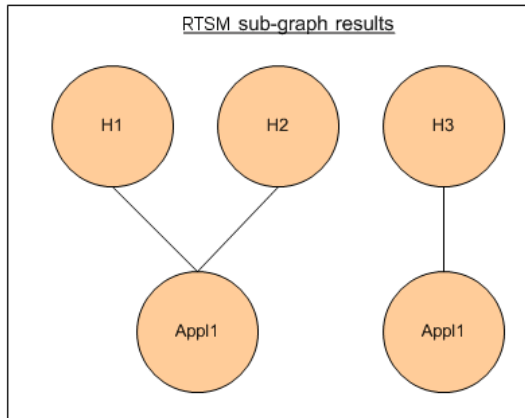


Note: For attribute federation, the adapter's **getTopology** method is called. The reconciliation data is adapted in the user TQL (in this case, the CI element).


- 1 After analyzing the adapter, the Federation Framework recognizes the virtual relationship between **Node** and **Incident** and splits the federated TQL query into two subgraphs:




- 2 The Federation Framework runs the BSM subgraph to request the topology, and receives the following results:




- 3** The Federation Framework requests, from the appropriate Mapping Engine, the reconciliation data for the first data repository (BSM) that contains the information to connect between received data from two data repositories. The reconciliation data in this case is:




- The Federation Framework creates a one-node topology query with the Node and ID conditions on it from the previous result (node in H1, H2, H3), and runs this query with the required reconciliation data on BSM. The result includes Node CIs that are relevant to the ID condition and the appropriate reconciliation object for each CI:




- 5 The reconciliation data for ServiceCenter should contain a condition for node and ip that is derived from the reconciliation objects received from BSM:



- 6 The Federation Framework runs the ServiceCenter subgraph with the reconciliation data to request the topology and appropriate reconciliation objects, and receives the following results:



- 7 The result after connection in Mapping Engine and creating virtual relationships is:



- 8 The Federation Framework requests the original TQL layout for received instances from BSM and ServiceCenter.

Federation Framework Flow for Population

This section includes the following topics:

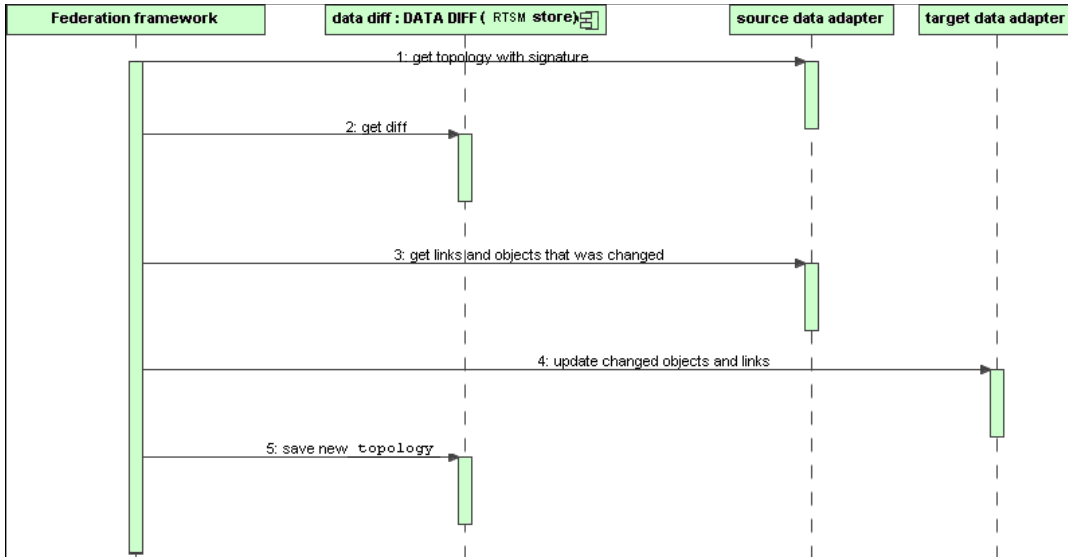
- "Definitions and Terms" on page 231
- "Flow Diagram" on page 232

Definitions and Terms

Signature. Denotes the state of properties in the CI. If changes are made to property values in a CI, the CI signature must also be changed. The CI signature helps to detect whether a CI has changed without retrieving and comparing all CI properties. Both the CI and the CI signature are provided by the appropriate adapter. The adapter is responsible for changing the CI signature when the CI properties are altered.

Flow Diagram

The following sequence diagram illustrates the interaction between the Federation Framework and the source and target adapters in a population flow:



- 1** The Federation Framework receives the topology for the query result from the source adapter. The adapter recognizes the query by its name and runs it on the external data repository. The topology result contains the ID and signature for each CI and relationship in the result. The ID is the logical ID that defines the CI as unique in the external data repository. The signature should be modified if the CI or relationship is modified.
- 2** The Federation Framework uses signatures to compare the newly received topology query results with the saved ones, and to determine which CIs have changed.
- 3** After the Federation Framework finds the CIs and relationships that have changed, it calls the source adapter with the IDs of the changed CIs and relationships as a parameter to retrieve their full layout.
- 4** The Federation Framework sends the update to the target adapter. The target adapter updates the external data source with the received data.
- 5** After the update, the Federation Framework saves the last query result.

Adapter Interfaces

This section includes the following topics:

- "Definitions and Terms" on page 233
- "Adapter Interfaces for Federated TQL Queries" on page 233

Definitions and Terms

The external relation. The relation between two external CI types that are supported by the same adapter.

Adapter Interfaces for Federated TQL Queries

Use the appropriate adapter interface for each adapter, as follows.

A **one Node topology interface** is used when the adapter does not support any external relations; that is, the adapter is never meant to receive a request with more than one external CI. All OneNode interfaces are created to simplify the workflow; for those cases where you need to use a more extensive query, use the `DataAdapter` interface.

Deprecated as of UCMDB 9.00: Pattern topology interface

A **DataAdapter interface** is used to define adapters that support complex federated queries. The reconciliation request in these adapters is part of the single `QueryDefinition` parameter. These adapters may also be used for Population.

OneNode Interfaces

The following interfaces have different types of reconciliation data:

- **OneNodeTopologyIdReconciliationDataAdapter.** Use if the adapter supports a **single-node TQL** and the reconciliation between data repositories is calculated by the ID.
- **OneNodeTopologyPropertyReconciliationDataAdapter.** Use if the adapter supports a **single-node TQL** and the reconciliation between data repositories is done by the properties of one CI.

- ▶ **OneNodeTopologyDataAdapter.** Use if the adapter supports a **single-node TQL** and the reconciliation between data repositories is done by topology.

Data Adapter Interfaces

- ▶ **DataAdapter.** Use this adapter to support complex federated TQL queries. Allows the most diversity.
- ▶ **PopulateDataAdapter.** Use this adapter to support complex federated TQL queries and population flows. In a population flow, this adapter retrieves the entire data set, and lets the probe filter the difference since the last execution of the job.
- ▶ **PopulateChangesDataAdapter.** Use this adapter to support complex federated TQL queries and population flows. In a population flow, this adapter supports the retrieval of only the changes that occurred since the last execution of the job.

Pattern Topology Interfaces (Deprecated as of UCMDB 9.00)

The following interfaces have different types of reconciliation data:

- ▶ **PatternTopologyIdReconciliationDataAdapter.** Use if the adapter supports a **complex TQL** and the reconciliation between data repositories is done by the ID.
- ▶ **PatternTopologyPropertyReconciliationDataAdapter.** Use if the adapter supports a **complex TQL** and the reconciliation between data repositories is done by single-node properties.
- ▶ **PatternTopologyDataAdapter.** Use if the adapter supports a **complex TQL** and the reconciliation between data repositories is done by topology.

Additional Interfaces

- ▶ **SortResultDataAdapter.** Use if you can sort the resulting CIs in the external data repository.
- ▶ **FunctionalLayoutDataAdapter.** Use if you can calculate the functional layout in the external data repository.

Adapter Interfaces for Synchronization

- **SourceDataAdapter.** Use for source adapters in population flows.
- **TargetDataAdapter.** Use for target adapters in data push flows.

Tasks

How to Debug Adapter Resources

This task describes how to use the JMX console to create, view, and delete adapter state resources (any resources created using the resource manipulation methods in the `DataAdapterEnvironment` interface, which are saved in the RTSM database or the Probe database) for debugging and development purposes.

1 Launch the Web browser and enter the server address, as follows:

- For the RTSM server: `http://localhost:8080/jmx-console`
- For the Probe: `http://localhost:1977`

You may have to log in with a user name and password (the defaults are `sysadmin/sysadmin`).

2 To open the JMX MBEAN View page, do one of the following:

- On the RTSM server: click **UCMDB:service=FCMDB Adapter State Resource Services**
- On the Probe: click **type=AdapterStateResources**

3 Enter values in the operations that you want to use, and click **Invoke**.

How to Add an Adapter for a New External Data Source

This task explains how to define an adapter to support a new external data source.

This task includes the following steps:

- "Prerequisites" on page 237
- "Define Valid Relationships for Virtual Relationships" on page 238
- "Define an Adapter Configuration" on page 238
- "Define Supported Classes" on page 242

- "Implement the Adapter" on page 242
- "Define Reconciliation Rules or Implement the Mapping Engine" on page 243
- "Add Jars Required for Implementation to the Class Path" on page 243
- "Deploy the Adapter" on page 243
- "Update the Adapter" on page 244

1 Prerequisites

Model-supported adapter classes for CIs and relationships in the BSM Data Model. As an adapter developer, you should:

- have knowledge of the hierarchy of the BSM CI types to understand how external CITs are related to the BSM CITs
- model the external CITs in the BSM class model
- add the definitions for new CI types and their relationships
- define valid relationships in the BSM class model for the valid relationships between adapter inner classes. (The CITs can be placed at any level of the BSM class model tree.)

Modeling should be the same regardless of federation type (on the fly or replication). For details on adding new CIT definitions to the BSM class model, see "Working with the CI Selector" in the *RTSM Modeling Guide*.

For the adapter to support federated attributes on CITs, add this CIT to the supported classes with supported attributes and the reconciliation rule for this CIT.

2 Define Valid Relationships for Virtual Relationships

Note: This section is relevant only for federation.

To retrieve federated CITs that are connected to local RTSM CITs, a valid link definition must exist between the two CITs in the RTSM.


- a** Create a valid links XML file that contains these links (if they do not already exist).
- b** Add the links XML file to the adapter package in the `\validlinks` folder. For details, see "Package Manager" in the *RTSM Administration Guide*.

Example of Valid Relationship Definition:

In the following example, the relation of type containment between instances of type `node` to instances of type `myclass1` is a valid relationship definition.

```
<Valid-Links>
  <Valid-Link>
    <Class-Ref class-name="containment"/>
    <End1 class-name="node"/>
    <End2 class-name="myclass1"/>
    <Valid-Link-Qualifiers/>
  </Valid-Link>
</Valid-Links>
```

3 Define an Adapter Configuration

- a** Navigate to **Admin > RTSM Administration > Adapter Management**.
-  **b** Click the **Create new resource** button.
- c** In the New adapter dialog box, select **Integration** and **Java Adapter**.
- d** Right-click on the adapter that you created and select **Edit Adapter Source** from the shortcut menu.

- e Edit the following XML tags:

```
<pattern xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" id="newAdapterIdName"
xsi:noNamespaceSchemaLocation="../../../Patterns.xsd" description="Adapter Description"
schemaVersion="9.0" displayName="New Adapter Display Name">
  <deletable>true</deletable>
  <discoveredClasses>
    <discoveredClass>link</discoveredClass>
    <discoveredClass>object</discoveredClass>
  </discoveredClasses>
  <taskInfo className="com.hp.ucmdb.discovery.probe.services.dynamic.core.AdapterService">
    <params
className="com.hp.ucmdb.discovery.probe.services.dynamic.core.AdapterServiceParams"
enableAging="true" enableDebugging="false" enableRecording="false" autoDeleteOnErrors="success"
recordResult="false" maxThreads="1" patternType="java_adapter" maxThreadRuntime="25200000">
      <className >com.yourCompany.adapter.MyAdapter.MyAdapterClass</className>
    </params>
```

```

    <destinationInfo className="com.hp.ucmdb.discovery.probe.tasks.BaseDestinationData">
      <!-- check -->
      <destinationData name="adapterId" description="">${ADAPTER.adapter_id}</
destinationData>
      <destinationData name="attributeValues" description="">${SOURCE.attribute_values}</
destinationData>
      <destinationData name="credentialsId" description="">${SOURCE.credentials_id}</
destinationData>
      <destinationData name="destinationId" description="">${SOURCE.destination_id}</
destinationData>
    </destinationInfo>
    <resultMechanism isEnabled="true">
      <autoDeleteCITs isEnabled="true">
        <CIT>link</CIT>
        <CIT>object</CIT>
      </autoDeleteCITs>
    </resultMechanism>
  </taskInfo>
  <adapterInfo>
    <adapter-capabilities>
      <support-federated-query>
        <!--<supported-classes/> <!--see the section about supported classes-->
      </support-federated-query>
      <topology>
        <pattern-topology /> <!--or <one-node-topology> -->
      </topology>
    </support-federated-query>
    <!--<support-replicatioin-data>
    <source>
      <changes-source/>
    </source>
  </target/>
  </adapter-capabilities>
  <default-mapping-engine />
  <queries />
  <removedAttributes />
  <full-population-days-interval>-1</full-population-days-interval>
</adapterInfo>
<inputClass>destination_config</inputClass>
<protocols />

```



```

<parameters>
  <!--The description attribute may be written in simple text or HTML.-->
  <!--The host attribute is treated as a special case by UCMDB-->
  <!--and will automatically select the probe name (if possible)-->
  <!--according to this attribute's value.-->
  <parameter name="credentialsId" description="Special type of property, handled by UCMDB for
credentials menu" type="integer" display-name="Credentials ID" mandatory="true" order-index="12" />
  <parameter name="host" description="The host name or IP address of the remote machine"
type="string" display-name="Hostname/IP" mandatory="false" order-index="10" />
  <parameter name="port" description="The remote machine's connection port" type="integer"
display-name="Port" mandatory="false" order-index="11" />
</parameters>
<parameter name="myatt" description="is my att true?" type="string" display-name="My Att"
mandatory="false" order-index="15" valid-values="True;False"/>True</parameters>
<collectDiscoveredByInfo>true</collectDiscoveredByInfo>
<integration isEnabled="true">
  <category >My Category</category>
</integration>
<overrideDomain>${SOURCE.probe_name}</overrideDomain>
<inputTQL>
  <resource:XmlResourceWrapper xmlns:resource="http://www.hp.com/ucmdb/1-0-0/
ResourceDefinition" xmlns:ns4="http://www.hp.com/ucmdb/1-0-0/ViewDefinition" xmlns:tql="http://
www.hp.com/ucmdb/1-0-0/TopologyQueryLanguage">
    <resource xsi:type="tql:Query" group-id="2" priority="low" is-live="true" owner="Input TQL"
name="Input TQL">
      <tql:node class="adapter_config" id="-11" name="ADAPTER" />
      <tql:node class="destination_config" id="-10" name="SOURCE" />
      <tql:link to="ADAPTER" from="SOURCE" class="fcmdb_conf_aggregation" id="-12"
name="fcmdb_conf_aggregation" />
    </resource>
  </resource:XmlResourceWrapper>
</inputTQL>
<permissions />
</pattern>

```

For details about the XML tags, see "XML Configuration Tags and Properties" on page 249.

4 Define Supported Classes

Define supported classed either the adapter code by implementing the *getSupportedClasses()* method, or by using the pattern XML file.

```
<supported-classes>
  <supported-class name="HistoryChange" is-derived="false"
  is-reconciliation-supported="false" federation-not-supported="false"
  is-id-reconciliation-supported="false">
    <supported-conditions>
      <attribute-operators attribute-name="change_create_time">
        <operator>GREATER</operator>
        <operator>LESS</operator>
        <operator>GREATER_OR_EQUAL</operator>
        <operator>LESS_OR_EQUAL</operator>
        <operator>CHANGED_DURING</operator>
      </attribute-operators>
    </supported-conditions>
  </supported-class>
```

name	The name of the CI type
is-derived	Specifies whether this definition includes all inheriting children
is-reconciliation-supported	Specifies whether this class is used for reconciliation
is-id-reconciliation-supported	Specifies whether this class is used for id-reconciliation
federation-not-supported	Specifies whether this CIT should not be allowed for federation (blocking certain CITs, for example, a CIT defined solely for federation)
<supported-conditions>	Specifies the supported conditions on each attribute

5 Implement the Adapter

Select the correct adapter implementation class according to its defined capabilities. The adapter implementation class implements the appropriate interfaces according to defined capabilities.

6 Define Reconciliation Rules or Implement the Mapping Engine

If your adapter supports federated TQL queries, you have three options for defining your Mapping Engine:

- Use the default RTSM 9.x default mapping engine, which uses the RTSM's internal reconciliation rules for mapping. To use it, leave the `<default-mapping-engine/>` XML tag empty.

For details, see "reconciliation_types.txt file" on page 179.

- Use the RTSM 8.0x mapping engine. To do this, use the following XML Tag:
`<default-mapping-engine>com.hp.ucmdb.federation.mappingEngine.AdapterMappingEngine</default-mapping-engine>`

For details, see "reconciliation_rules.txt File (for backwards compatibility)" on page 179.

- Write your own mapping engine by implementing the mapping engine interface and placing the JAR with the rest of the adapter code. To do this, use the following XML tag:
`<default-mapping-engine>com.yourcompany.map.MyMappingEngine</default-mapping-engine>`

7 Add Jars Required for Implementation to the Class Path

To implement your classes, add the `federation_api.jar` file to your code editor class path.

8 Deploy the Adapter

Deploy the adapter package. For general details on deploying a package, see "Package Manager" in the *RTSM Administration Guide*.

The package should contain the following entities:

- New CIT definition (optional):
- Used only if the adapter supports new CI types that do not yet exist in BSM.
- The new CIT definitions are located in the class folder in the package.

- ▶ New data type definition (optional):
- ▶ Used only if the new CITs require new data types.
- ▶ The new data type definitions are located in the `typedef` folder in the package.
- ▶ New valid relationships definition (optional):
- ▶ Used only if the adapter supports the federated TQL.
- ▶ The new valid relationships definitions are located in the `validlinks` folder in the package.
- ▶ The pattern configuration XML file should be located in the `discoveryPatterns` folder in the package.
- ▶ **Descriptor.** Defines the package definitions.
- ▶ Place your compiled classes (normally a jar file) in the package under the `adapterCode\<adapter id>` folder.

Note: The `adapter id` folder name has the same value as in the adapter configuration.

- ▶ If you create your own configuration file, you should place the file in the package under the `adapterCode\<adapter id>` folder.

9 Update the Adapter

Changes to any of the adapter's non-binary files may be made in the Adapter Management module. Making changes to configuration files in the Adapter management module causes the adapter to reload with the new configurations.

Updates may also be made by editing the files in the package (both binary and non-binary files), and then redeploying the package by using the Package Manager. For details, see "How to Deploy a Package" in the *RTSM Administration Guide*.

How to Implement the Mapping Engine

The configuration of the mapping engine depends on which mapping engine you are using.

This task includes the following steps:

- "Configure the reconciliation_types.txt File (for the BSM 9.x default mapping engine)" on page 245
- "Configure the reconciliation_rules.txt File (for the BSM 8.0x mapping engine)" on page 245

1 Configure the reconciliation_types.txt File (for the BSM 9.x default mapping engine)

The file is used to define which CI types are used for reconciliation in the adapter.

Write each CI type used for reconciliation on a single line, as follows:

```
node
business_application
```

Place the file in the adapter package in the `adapterCode\<AdapterID>\META-INF\` folder.

To support ID reconciliation (reconciliation based on ID mapping between the RTSM ID in the RTSM to a value on the remote database), you should map a special RTSM attribute called `cmdb_id` to a column in the database of either the string (char, varchar) or byte[] (raw/bytes) type.

2 Configure the reconciliation_rules.txt File (for the BSM 8.0x mapping engine)

This file is used to configure the reconciliation rules. Each row in the file represents a rule. For example:

```
reconciliation_type[node] expression[^node.name OR ip_address.name]
end1_type[node] end2_type[ip_address] link_type[containment]
```

The **reconciliation_type** parameter is filled with the type of CI on which the reconciliation is performed (the BSM class name that is connected to the federated class in the TQL).

The **expression** parameter is the logic that decides whether two reconciliation objects are equal (one reconciliation object from the BSM side and the other from the Federated adapter side).

The expression is composed of ORs and ANDs.

The convention regarding attributes names in the expression part is **[className].[attributeName]**. For example, the attribute **ip_address** in the **ip** class is written **ip.ip_address**.

You can define ordered matches. The ordered match checks the first OR sub expression. If two reconciliation objects have the value on the attributes of the sub expression and it returns that false (the reconciliation objects are not equal) then the second OR sub expression is not compared.

For an ordered match, use **ordered expression** instead of **expression**.

The circumflex sign (^) is used to ignore case during comparisons.

The other parameters (**end1_type**, **end2_type**, and **link_type**) are used only if the reconciliation data contains two nodes and not just the node of the reconciliation type (the topological reconciliation data). In this case, the reconciliation data is **end1_type -(link_type)> end2_type**.

There is no need to add the relevant layout as it is retrieved from the expression.

To perform reconciliation by BSM ID, use **cmdb_id** as the attribute name in expression.

Place the file in the adapter package in the **adapterCode\<AdapterID>\META-INF** folder.

Examples:

- You can add a reconciliation rule for a node CIT only. This is because only node CITs have valid relationships with external CITs. For example, a node CI in the RTSM is matched to a node CI in ServiceCenter through the **node.name** attribute or through the **ip_address.name** attribute.

- The reconciliation rule in this case is a topology rule and the expression is ordered. The rule performs the following checks on the CIs under comparison:
 - If the `node.name` attribute is equal, the rule matches the nodes.
 - If the `node.name` attribute is not equal, the rule does not match the nodes.
 - If the `node.name` attribute is null in one of the compared CIs, the rule checks the `ip_address.name` attribute. If the `ip_address.name` attribute is equal, the rule matches the nodes.

How to Create a Sample Adapter

This example illustrates how to create a sample adapter.

This task includes the following steps:

- "Select Adapter Logic" on page 247
- "Load the Project" on page 248

1 Select Adapter Logic

When you implement an adapter, you must choose how to handle the condition logic in the implementation (property conditions, ID conditions, reconciliation conditions, and link conditions).

- a** Retrieve the entire data into the adapter memory and let it select or filter the needed CI Instances.
- b** Convert all the conditions into the data source language and let it filter and select the data. For example:
 - Convert the condition into a SQL query.
 - Convert the condition into a Java API filter object.
- c** The middle road is to filter some of the data on the remote service, and have the adapter select and filter the remainder.

In the MyAdapter example, the logic in step a is used.

2 Load the Project

Copy the files from the <HP BSM root directory\odb\tools\adapter-dev-kit folder and follow the instructions in the readme files.

Note: If you use an adapter with large data sets, you may need to use caching and indexing to improve performance for Federation.

Online javadocs documentation is available at:

C:\hp\UCMDB\UCMDBServer\deploy\ucmdb-docs\docs\eng\doc_lib\DevRef_guide\DBAdapterFramework_JavaAPI\index.html

Reference

XML Configuration Tags and Properties

id="newAdapterIdName"	Defines the adapter's real name. Used for logs and folder lookups
displayName="New Adapter Display Name"	Defines the adapter's display name, as it appears in the UI.
<className>...</className>	Defines the adapter's interface implementing the Java class.
<category >My Category</category>	Defines the adapter's category.
<parameters>	Defines the properties for the configuration that are available in the UI when setting up a new integration point.
name	The name of the property (used mostly by code)
description	The display hint of the property
type	String or integer (use valid values with string for Boolean).
display-name	The name of the property in the UI.
mandatory	Specifies whether this configuration property is mandatory for the user.
order-index	The placing order of the property (small = up)
valid-values	A list of possible valid values separated by ';' characters (for example, valid-values="Oracle;SQLServer;MySQL" or valid-values="True;False").
<adapterInfo>	Contains the definition of the adapter's static settings and capabilities.
<support-federated-query>	Defines this adapter as capable of federation.
<one-node-topology>	The ability to federated queries with one federated query node.
<pattern-topology>	The ability to federate complex queries.
<support-replicatioin-data>	Defines the capability to run data push and population flows.

<source>	This adapter may be used for population flows.
<changes-source/>	This adapter may be used for population changes flows.
<target>	This adapter may be used for data push flows.
<default-mapping-engine>	Allows definition of a mapping engine for the adapter (by default, the adapter uses the default mapping engine). For any other mapping engine, enter the implementing class name of the mapping engine (for the BSM 8.0x mapping engine use: com.hp.ucmdb.federation.mappingEngine.AdapterMappingEngine)
<removedAttributes>	Forces the removal of specific attributes from the result.
<full-population-days-interval>	Specifies when to execute a full population job instead of a differential job (every 'x' days). Uses the aging mechanism together with the changes flow.

6

Developing Push Adapters

This chapter includes:

Concepts

- ▶ Developing Push Adapters Overview on page 252
- ▶ Differential Synchronization on page 252

Tasks

- ▶ How to Prepare the Mapping Files on page 253
- ▶ How to Write Jython Scripts on page 255
- ▶ How to Support Differential Synchronization on page 259
- ▶ How to Build an Adapter Package on page 261

Reference

- ▶ Mapping File Schema on page 263
- ▶ Mapping Results Schema on page 278

Concepts

Developing Push Adapters Overview

The Generic Push Adapter provides a platform that enables rapid development of integrations that push BSM 9.x data to external data repositories (databases and third-party applications). Developing a custom integration based on Generic Push Adapter requires:

- ▶ an XML mapping file between the BSM CI link types and the external data items.
- ▶ a Jython script to push the data items into the external data repository.

Differential Synchronization

For the Push adapter to support differential synchronization, the **DiscoveryMain** function must return an object implementing the **DataPushResults** interface, which contains the mappings between the IDs that the Jython script receives from the XML and the IDs that the Jython script creates on the remote machine. The latter IDs are of the type **ExternalId**.

The **ExternalIdUtil.restoreExternal** command, which receives the ID of the CI in the CMDB as a parameter, restores the external ID from the ID of the CI in the CMDB. This command can be used, for example, while performing differential synchronization, and a link is received where one of its ends is not in the bulk (it was already synchronized).

If the **DiscoveryMain** method in the Jython script on which the Push adapter is based returns an empty **ObjectStateHolderVector** instance, the adapter will not support differential synchronization. This means that even when a differential synchronization job is run, in actuality, a full synchronization is being performed. Therefore, no data can be updated or removed on the remote system, since all data is added to the RTSM during each synchronization.

Tasks

How to Prepare the Mapping Files

There are two different ways to prepare mapping files:

- ▶ You can prepare a single, global mapping file.

All mappings are placed in a single file named **mappings.xml**.

- ▶ You can prepare a separate file for each push query.

Each each mapping file is called **<query name>.xml**.

For details, see "Mapping File Schema" on page 263.

This task includes the following steps:

- ▶ "Create the Mapping File" on page 253
- ▶ "Map CIs" on page 254
- ▶ "Map Links" on page 255

1 Create the Mapping File

The mapping file structure is as follows

```
<?xml version="1.0" encoding="UTF-8"?>
<integration>
  <info>
    <source name="UCMDB" versions="9.x" vendor="HP" />
    <!-- for example: -->
    <target name="Oracle" versions="11g" vendor="Oracle" />
  </info>
  <targetcis>
    <!-- CI Mappings --->
  </targetcis>
  <targetrelations>
    <!-- Link Mappings --->
  </targetrelations>
</integration>
```

2 Map CIs

There are two ways to map RTSM CI types:

- Map a CI type so that CIs of that type and all inherited types are mapped in the same way:

```
<source_ci_type_tree name="node" mode="update_else_insert">
  <apioutputseq>1</apioutputseq>
  <target_ci_type name="host">
    <targetprimarykey>
      <pkey>name</pkey>
    </targetprimarykey>
    <target_attribute name=" name" datatype="STRING">
      <map type="direct" source_attribute="name" />
    </target_attribute>
    <!-- more target attributes --->
  </target_ci_type>
</source_ci_type_tree>
```

- Map a CI type so that only CIs of that type will be processed. CIs of inherited types will not be processed unless their type is also mapped (in one of the two ways):

```
<source_ci_type name="node" mode="update_else_insert">
  <apioutputseq>1</apioutputseq>
  <target_ci_type name="host">
    <targetprimarykey>
      <pkey>name</pkey>
    </targetprimarykey>
    <target_attribute name=" name" datatype="STRING">
      <map type="direct" source_attribute="name" />
    </target_attribute>
    <!-- more target attributes --->
  </target_ci_type>
</source_ci_type>
```

A CI type which is mapped indirectly (one of its ancestors is mapped using **source_ci_type_tree**), can also override its parent's map by having it appear in its own **source_ci_type_tree** or **source_ci_type**.

It is recommended to use **source_ci_type_tree** wherever possible. Otherwise, resulting CIs of a CI type that do not appear in the mapping files will not be transferred to the Jython script.

3 Map Links

There are two ways to map links:

- Map a link which will also map all of the link types that inherit from that specific link:

```
<source_link_type_tree name="dependency" target_link_type="dependency"
mode="update_else_insert" source_ci_type_end1="webservice"
source_ci_type_end2="sap_gateway">
  <target_ci_type_end1 name="webservice" />
  <target_ci_type_end2 name="sap_gateway" />
  <target_attribute name="name" datatype="STRING">
    <map type="direct" source_attribute="name" />
  </target_attribute>
</source_link_type_tree>
```

- Map a link which will also map only that specific link type and not the link types which inherit from it:

```
<link source_link_type="dependency" target_link_type="dependency"
mode="update_else_insert" source_ci_type_end1="webservice"
source_ci_type_end2="sap_gateway">
  <target_ci_type_end1 name="webservice" />
  <target_ci_type_end2 name="sap_gateway" />
  <target_attribute name="name" datatype="STRING">
    <map type="direct" source_attribute="name" />
  </target_attribute>
</link>
```

How to Write Jython Scripts

The mapping script is a regular Jython script, and should follow the rules for Jython scripts. For details, see "Developing Jython Adapters" on page 63.

The script should contain the **DiscoveryMain** function, which may return either an empty **OSHVResult** or a **DataPushResults** instance upon success.

To report any failure, the script should raise an exception, for example:

```
raise Exception('Failed to insert to remote UCMDB using TopologyUpdateService. See log of the remote UCMDB')
```

In the `DiscoveryMain` function, the data items to be pushed to or deleted from the external application can be obtained as follows:

```
# get add/update/delete result objects (in XML format) from the Framework
addResult = Framework.getTriggerCIData('addResult')
updateResult = Framework.getTriggerCIData('updateResult')
deleteResult = Framework.getTriggerCIData('deleteResult')
```

The client object to the external application can be obtained as follows:

```
oracleClient = Framework.createClient()
```

This client object automatically uses the credentials ID, host name and port number passed by the adapter through the Framework.

If you need to use the connection parameters that you defined for the adapter (for details, see step 2 in "How to Build an Adapter Package" on page 261), use the following code:

```
propValue = str(Framework.getDestinationAttribute('<Connection Property Name'))
```

For example:

```
serverName = Framework.getDestinationAttribute('ip_address')
```

This section also includes:

- ▶ "Working with the Mapping's Results" on page 257
- ▶ "Handling Test Connection in the Script" on page 258

Working with the Mapping's Results

The Generic Push Adapter creates XML strings that describe the data to be added, updated, or deleted from the target system. The Jython script needs to analyze this XML, and then performs the add, update, or delete operation on the target.

In the XML of the add operation that the Jython script receives, the `mamId` attribute for the objects and links is always the BSM identifier of the original object or link before its type, attribute or other information was changed to the schema of the remote system.

In the XML of the update or remove operations, the `mamId` attribute of each object or link contains the string representation of the same `ExternalId` that was returned from the Jython script from the previous synchronization.

Example of the XML result

```

<data>
  <objects>
    <Object mode="update_else_insert" name="UCMDB_NODE"
operation="add" mamId="6ebafa5b84ab54c0e3155b762aec3e71">
      <field name="ID" key="true" datatype="char" length="">nt2</field>
      <<field name="NAME" key="false" datatype="char"
length="255">NT2</field>>
    </Object>
    <Object mode="update_else_insert" name="UCMDB_NODE"
operation="add" mamId="f682342fa70aebc73ad525ce10721462">
      <field name="ID" key="true" datatype="char" length="">node_nt2</
field>
      <field name="NAME" key="false" datatype="char"
length="255">NODE_NT2</field>
    </Object>
  </objects>
  </links>
  <link targetRelationshipClass="dependency" targetParent="node"
targetChild="node" operation="add" mode="update_else_insert"
mamId="end1id%5BUCMDB%0Ant%0A1%0Ainternal_id%3DSTRING%3D6ebafa5
b84ab54c0e3155b762aec3e71%0A%5D%0Aend2id%5BUCMDB%0Anode%0A1%
0Ainternal_id%3DSTRING%3Df682342fa70aebc73ad525ce10721462%0A%5D%0
AUCMDB%0Adependency%0A0%0A">
    <field
name="DiscoveryID1">6ebafa5b84ab54c0e3155b762aec3e71</field>
    <field
name="DiscoveryID2">f682342fa70aebc73ad525ce10721462</field>
    <field name="ID" key="true" datatype="char"
length="255">DEPENDENCY_NT2</field>
    <field name="NAME" key="false" datatype="char"
length="255">DEPENDENCY_NT2</field>
  </link>
</links>
</data>

```

Handling Test Connection in the Script

A Jython script can be invoked to test the connection with an external application. In this case, the `testConnection` destination attribute will be `true`. This attribute can be obtained from the Framework as follows:

```
testConnection = Framework.getTriggerCIData('testConnection')
```

When run in test connection mode, a script should raise an exception if a connection to the external application cannot be established. Otherwise, if the connection is successful, the **DiscoveryMain** function should return an empty **OSHVResult**.

How to Support Differential Synchronization

Important: If you are implementing differential synchronization on an existing adapter that was created in version 9.00 or 9.01, you must use the push-adapter.zip file from version 9.02 or later to recreate your adapter package. For details, see "How to Build an Adapter Package" on page 261.

This task enables the Push adapter to perform differential synchronization. For details, see "Differential Synchronization" on page 252.

The Jython script returns the **DataPushResults** object which contains two Java maps - one for object ID mappings (keys and values are ExternalCid type objects) and one for link IDs (keys and values are ExternalRelationId type objects).

- Add the following **from** statements to your Jython script:

```
from com.hp.ucmdb.federationspi.data.query.types import ExternalIdFactory
from com.hp.ucmdb.adapters.push import DataPushResults
from com.hp.ucmdb.adapters.push import DataPushResultsFactory
from com.mercury.topaz.cmdb.server.fcldb.spi.data.query.types import
ExternalIdUtil
```

- Use the **DataPushResultsFactory** factory class to obtain the **DataPushResults** object from the **DiscoveryMain** function.

```
# Create the UpdateResult object
updateResult = DataPushResultsFactory.createDataPushResults(objectMappings,
linkMappings);
```

- ▶ Use the following commands to create Java maps for the **DataPushResults** object:

```
# Prepare the maps to store the mappings if IDs
objectMappings = HashMap()
linkMappings = HashMap()
```

- ▶ Use the **ExternalIdFactory** class to create the following ExternalId IDs:
 - ▶ ExternalId for objects or links originating in a CMDB (for example, all of the CIs in an add operation are from the CMDB):

```
externaCIId = ExternalIdFactory.createExternalCmdbCild(ciType, ciIDAsString)
externalRelationId = ExternalIdFactory.createExternalCmdbRelationId(linkType,
end1ExternalCIId, end2ExternalCIId, linkIDAsString)
```

- ▶ ExternalId for objects or links not originating in a CMDB (usually, every update and remove operation contains such objects):

```
myIDField = TypesFactory.createProperty("systemID", "1")
myExternalId = ExternalIdFactory.createExternalCild(type, myIDField)
```

Note: If the Jython script updated existing information and the ID of the object (or link) changes, you must return a mapping between the previous external ID and the new one.

- ▶ Use the **restoreCmdbCiIDString** or **restoreCmdbRelationIDString** methods from the **ExternalIdFactory** class to retrieve the UCMDB ID string from an External ID of an object or link that originated in the UCMDB.
- ▶ Use the **restoreExternalCild** and **restoreExternalRelationId** methods from the **ExternalIdUtil** class to restore the **ExternalId** object from the **mamId** attribute value of the XML of the update or remove operations.

Note: ExternalId objects are actually an array of properties. This means that you can use an ExternalId object to store any information you may need that will identify the data on the remote system.

How to Build an Adapter Package

- 1 Extract the content of `<HP BSM root directory>\odb\content\adapters\push-adapter.zip` into a temporary folder. In the adapter package, the `sql_queries` file located in `adapterCode > PushAdapter > sqlTablesCreation`, contains the queries needed to create tables in a new schema in Oracle for testing the adapter. The tables correspond to the `adapterCode\<adapter ID>\mappings\mappings.xml` file.

Note: The `sql_queries` file is not needed for the adapter. It is only an example.

- 2 Edit the `discoveryPatterns\push_adapter.xml` file.
 - a Modify the `<pattern>` tag with a new id and display label. Replace:

```
<pattern id="PushAdapter" xsi:noNamespaceSchemaLocation="../../Patterns.xsd" description="Discovery Pattern Description" schemaVersion="9.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
```

with

```
<pattern id="MyPushAdapter" displayLabel="My Push Adapter" xsi:noNamespaceSchemaLocation="../../Patterns.xsd" description="Discovery Pattern Description" schemaVersion="9.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
```

- b** Update the parameter list, so that the list of parameters reflects the required connection attributes. Do not remove the **probeName** attribute.
- 3** Rename the **adapterCode\PushAdapter** folder with the adapter ID used in step 2 (for example, **adapterCode\MyPushAdapter**).
- 4** In the **discoveryScript** file, there is a **script pushScript.py** script which inserts the CIs and links to an external Oracle database. Replace **discoveryScripts\pushScript.py** with the script you wrote (for details, see "How to Write Jython Scripts" on page 255). If you rename the script, the **jythonScript.name** property in **adapterCode\<adapter ID>\push.properties** should be updated accordingly.
- 5** The **adapterCode\<adapter ID>\mappings\mappings.xml** file, located in **adapterCode\<adapter ID>\mappings**, is a sample mapping file which contains a mapping of the:
 - ▶ Node CI type with all the CI types that inherit from it
 - ▶ UNIX CI type without the CI types that inherit from it
 - ▶ Dependency link with all of the link types that inherit from it
 - ▶ Talk link type without the inherited link types that inherit from it

This mapping example corresponds to the example of the tables created in ORACLE in the **sql_queries** file (see step 1).

Replace the **adapterCode\<adapter ID>\mappings\mappings.xml** file with the mapping files you prepared (for details, see "How to Prepare the Mapping Files" on page 253).

If you want to use a mapping file for each TQL method, assign the name of the corresponding TQL to each XML file, followed by **.xml**. In this case, the **mappings.xml** file will be used as a default, if no specific mapping file is found for the current TQL name. The name of the default mapping file can be modified by changing the **mappingFile.default** property in **adapterCode\<adapter ID>\push.properties**.

Reference

Mapping File Schema

Element		Attributes
Name and Path	Description	
integration	Defines the mapping contents of the file. Must be the outermost block in the file except for the beginning line and any comments.	
info (integration)	Defines information about the data repositories being integrated	
source (integration > info)	Defines information about the source data repository	Name. type Description. Name of the source data repository. Is required. Required Type. String
		Name. versions Description. Version(s) of the source data repositories. Is required. Required Type. String
		Name. vendor Description. Vendor of the source data repository. Is required. Required Type. String

Element		Attributes
Name and Path	Description	
target (integration > info)	Defines information about the target data repository	Name. type Description. Name of the source data repository. Is required. Required Type. String
		Name. versions Description. Version(s) of the source data repository. Is required. Required Type. String
		Name. vendor Description. Vendor of the source data repository. Is required. Required Type. String
targetcis (integration)	Container element for all CIT mappings	

Element		Attributes
Name and Path	Description	
source_ci_type_tree (integration > targetcis)	Defines a source CIT and all of the CI types which inherit from it.	<p>Name. name</p> <p>Description. Name of the source CIT.</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. mode</p> <p>Description. The type of update required for the current CI type.</p> <p>Is required. Required</p> <p>Type. One of the following strings:</p> <ul style="list-style-type: none"> ▶ insert – Use this only if the CI does not already exist. ▶ update – Use this only if the CI is known to exist. ▶ update_else_insert – If the CI exists, update it; otherwise, create a new CI. ▶ ignore – Do nothing with this CI type.

Element		Attributes
Name and Path	Description	
source_ci_type (integration > targetcis)	Defines a source CIT without the CI types which inherit from it.	<p>Name. name</p> <p>Description. Name of the source CIT.</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. mode</p> <p>Description. The type of update required for the current CI type.</p> <p>Is required. Required</p> <p>Type. One of the following strings:</p> <ul style="list-style-type: none"> ▶ insert – Use this only if the CI does not already exist. ▶ update – Use this only if the CI is known to exist. ▶ update_else_insert – If the CI exists, update it; otherwise, create a new CI. ▶ ignore – Do nothing with this CI type.

Element		Attributes
Name and Path	Description	
target_ci_type (integration > targetcis > source_ci_type -OR- integration > targetcis > source_ci_type_tree)	Defines a target CIT	Name. name Description. Target CI type name. Is required. Required Type. String
		Name. schema Description. The name of the schema that will be used to store this CI type at the target. Is required. Not Required Type. String
		Name. namespace Description. Indicates the namespace of this CI type on the target Is required. Not Required Type. String

Element		Attributes
Name and Path	Description	
targetprimarykey (integration > targetcis > source_ci_type) -OR- (integration > targetcis > source_ci_type_tree) -OR- (integration > targetrelations > link) -OR- (integration > targetrelations > source_link_type_ tree)	Identifies target CIT primary key attributes	

Element		Attributes
Name and Path	Description	
<p>pkey (integration > targetcis > source_ci_type > targetprimarykey -OR- integration > targetcis > source_ci_type_tree > targetprimarykey -OR- (integration > targetrelations > link > targetprimarykey -OR- integration > targetrelations > source_link_type_tree > targetprimarykey)</p>	<p>Identifies one primary key attribute</p> <p>Required only if mode is update or insert_else_update</p>	

Element		Attributes
Name and Path	Description	
target_attribute (integration > targetcis > source_ci_type -OR- integration > targetcis > source_ci_type_tree -OR- integration > targetrelations > link -OR- integration > targetrelations > source_link_type_tree)	Defines the target CIT's attribute	Name. name Description. Name of the target CIT's attribute. Is required. Required Type. String
		Name. datatype Description. Data type of the target CIT's attribute. Is required. Required Type. String
		Name. length Description. For string/char data types, integer size of target attribute. Is required. Not Required Type. Integer
		Name. option Description. The conversion function to be applied to the value. Is required. False Type. One of the following strings: <ul style="list-style-type: none"> ▶ uppercase – Convert to uppercase ▶ lowercase – Convert to lowercase ▶ If this attribute is empty, no conversion function will be applied.

Element		Attributes
Name and Path	Description	
<p>map (integration > targetcis > source_ci_type > target_attribute -OR- integration > targetcis > source_ci_type_tree > target_attribute) -OR- (integration > targetrelations > link > target_attribute -OR- integration > targetrelations > source_link_type_tree > target_attribute)</p>	<p>Specifies how to obtain the source CIT's attribute value</p>	<p>Name. type Description. The type of mapping between the source and target values. Is required. Required Type. One of the following strings:</p> <ul style="list-style-type: none"> ▶ direct – Specifies a 1-to-1 mapping from source attribute's value to target attribute's value ▶ compoundstring – Sub-elements are joined into a single string and the target attribute value is set ▶ childattr – Sub-elements are one or more child CIT's attributes. Child CITs are defined as those with container_f or contained relationship. ▶ constant – Static string
		<p>Name. value Description. Constant string for type=constant Is required. Only required when type=constant Type. String</p>
		<p>Name. attr Description. Source attribute name for type=direct Is required. Only required when type=direct Type. String</p>

Element		Attributes
Name and Path	Description	
aggregation (integration > targetcis > source_ci_type > target_attribute > map -OR- integration > targetcis > source_ci_type_tree > target_attribute > map -OR- (integration > targetrelations > link > target_attribute > map -OR- integration > targetrelations > source_link_type_tree > target_attribute > map) Only valid when the map's type is childattr	Specifies how the source CI's child CI attribute values are combined into a single value to map to the target CI attribute. Optional.	Name. type Description. The type of aggregation function Is required. Required Type. One of the following strings: <ul style="list-style-type: none"> ▶ csv – Concatenates all included values into a comma-separated list (numeric or string/character). ▶ count – Returns a numeric count of all included values. ▶ sum – Returns a numeric count of all included values. ▶ average – Returns a numeric average of all included values. ▶ min – Returns the lowest numeric/ character included value. ▶ max – Returns the highest numeric/ character included value.

Element		Attributes
Name and Path	Description	
validation (integration > targetcis > source_ci_type > target_attribute > map -OR- integration > targetcis > source_ci_type_tree > target_attribute > map -OR- (integration > targetrelations > link > target_attribute > map -OR- integration > targetrelations > source_link_type_tree > target_attribute > map) Only valid when the map's type is childatt	Allows exclusion filtering of the source CI's child CIs based on attribute values. Used with the aggregation sub-element to achieve granularity of exactly which children attributes are mapped to the target CIT's attribute value. Optional.	Name. minlength Description. Excludes strings shorter than the given value. Required. Not required Type. Integer
		Name. maxlength Description. Excludes strings longer than the given value. Required. Not required Type. Integer
		Name. minvalue Description. Excludes numbers smaller than the specified value. Required. Not required Type. Numeric
		Name. maxvalue Description. Excludes numbers greater than the specified value. Required. Not required Type. Numeric
targetrelations (integration)	Container element for all relationship mappings. Optional.	

Element		Attributes
Name and Path	Description	
source_link_type_tree (integration > targetrelations)	Maps a source Relationship type without the types which inherit from it to a target Relationship. Mandatory only if targetrelation is present.	<p>Name. name</p> <p>Description. Source relationship name.</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. target_link_type</p> <p>Description. Target relationship name.</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. nameSpace</p> <p>Description. The namespace for the link that will be created on the target.</p> <p>Is required. Not required</p> <p>Type. String</p>
		<p>Name. mode</p> <p>Description. The type of update required for the current link.</p> <p>Is required. Required</p> <p>Type. On the following strings:</p> <ul style="list-style-type: none"> ➤ insert – Use this only if the CI does not already exist. ➤ update – Use this only if the CI is known to exist. ➤ update_else_insert – If the CI exists, update it; otherwise, create a new CI. ➤ ignore – Do nothing with this CI type.

Element		Attributes
Name and Path	Description	
source_link_type_tree (continued)		Name. source_ci_type_end1 Description. Source relationship's End1 CI type Is required. Required Type. String
		Name. source_ci_type_end2 Description. Source relationship's End2 CI type Is required. Required Type. String

Element		Attributes
Name and Path	Description	
link (integration > targetrelations)	Maps a source Relationship to a target Relationship. Mandatory only if targetrelation is present.	<p>Name. source_link_type Description. Source relationship name. Is required. Required Type. String</p>
		<p>Name. target_link_type Description. Target relationship name. Is required. Required Type. String</p>
		<p>Name. nameSpace Description. The namespace for the link that will be created on the target. Is required. Not required Type. String</p>
		<p>Name. mode Description. The type of update required for the current link. Is required. Required Type. On the following strings:</p> <ul style="list-style-type: none"> ➤ insert – Use this only if the CI does not already exist. ➤ update – Use this only if the CI is known to exist. ➤ update_else_insert – If the CI exists, update it; otherwise, create a new CI. ➤ ignore – Do nothing with this CI type.

Element		Attributes
Name and Path	Description	
link (continued)		Name. source_ci_type_end1 Description. Source relationship's End1 CI type Is required. Required Type. String
		Name. source_ci_type_end2 Description. Source relationship's End2 CI type Is required. Required Type. String
target_ci_type_end1 (integration > targetrelations > link -OR- integration > targetrelations > source_link_type_tree)	Target relationship's End1 CI type	Name. name Description. Name of the target relationship's End1 CI type. Is required. Required Type. String
		Name. superclass Description. Name of the End1 CI type's super-class. Is required. Not required Type. String

Element		Attributes
Name and Path	Description	
target_ci_type_end2 (integration > targetrelations > link -OR- integration > targetrelations > source_link_type_tree)	Target relationship's End2 CI type	<p>Name. name Description. Name of the target relationship's End2 CI type. Is required. Required Type. String</p> <p>Name. superclass Description. Name of the End2 CI type's super-class. Is required. Not required Type. String</p>

Mapping Results Schema

Element		Attributes
Name and Path	Description	
root	The root of the result document	
data (root)	The root of the data itself	
objects (root > data)	The root element for the objects to update	

Element		Attributes
Name and Path	Description	
Object (root > data > objects)	Describes the update operation for a single object and all of its attributes	<p>Name. name</p> <p>Description. Name of the CI type</p> <p>Is required. Required</p> <p>Type. String</p>
		<p>Name. mode</p> <p>Description. The type of update required for the current CI type.</p> <p>Is required. Required</p> <p>Type. One of the following strings:</p> <ul style="list-style-type: none"> ➤ insert – Use this only if the CI does not already exist. ➤ update – Use this only if the CI is known to exist. ➤ update_else_insert – If the CI exists, update it; otherwise, create a new CI. ➤ ignore – Do nothing with this CI type.

Element		Attributes
Name and Path	Description	
Object (continued)		<p>Name. operation</p> <p>Description. The operation to perform with this CI.</p> <p>Is required. Required</p> <p>Type. One of the following strings:</p> <ul style="list-style-type: none"> ➤ add – The CI should be added ➤ update – The CI should be updated ➤ delete – The CI should be deleted <p>If no value is set, then the default value of add is used.</p>
		<p>Name. mamId</p> <p>Description. The ID of the object on the source RTSM.</p> <p>Is required. Required</p> <p>Type. String</p>

Element		Attributes
Name and Path	Description	
field (root > data > objects > Object -OR- root > data > links > link)	Describes the value of a single field for an object. The field's text is the new value in the field, and if the field contains a link, the value is the ID of one of the ends. Each end ID appears as an object (under <objects>).	Name. name Description. Name of the field. Is required. Required Type. String
		Name. key Description. Specifies whether this field is a key for the object. Is required. Required Type. Boolean
		Name. datatype Description. The type of the field. Is required. Required Type. String
		Name. length Description. For string/character data types, this is the integer size of the target attribute. Is required. Not Required Type. Integer

Element		Attributes
Name and Path	Description	
links (root > data)	The root element for the links to update	<p>Name. targetRelationshipClass Description. The name of the relationship (link) in the target system. Is required. Required Type. String</p>
		<p>Name. targetParent Description. The type of first end of the link (parent). Is required. Required Type. String</p>
		<p>Name. targetChild Description. The type of the second end of the link (child). Is required. Required Type. String</p>

Element		Attributes
Name and Path	Description	
links (continued)		<p>Name. mode</p> <p>Description. The type of update required for the current CI type.</p> <p>Is required. Required</p> <p>Type. One of the following strings:</p> <ul style="list-style-type: none"> ▶ insert – Use this only if the CI does not already exist. ▶ update – Use this only if the CI is known to exist. ▶ update_else_insert – If the CI exists, update it; otherwise, create a new CI. ▶ ignore – Do nothing with this CI type.
		<p>Name. operation</p> <p>Description. The operation to perform with this CI.</p> <p>Is required. Required</p> <p>Type. One of the following strings:</p> <ul style="list-style-type: none"> ▶ add – The CI should be added ▶ update – The CI should be updated ▶ delete – The CI should be deleted <p>If no value is set, then the default value of add is used.</p>
		<p>Name. mamId</p> <p>Description. The ID of the object on the source RTSM.</p> <p>Is required. Required</p> <p>Type. String</p>

7

Viewing KPIs in External Applications

This chapter includes:

Tasks

- ▶ How to Set Up an Adapter to View KPIs in an External Application on page 286

Tasks

How to Set Up an Adapter to View KPIs in an External Application

This chapter explains how to retrieve the status and value of Business Service Management KPIs when working in an external application. You can view the results in your application, or you can use the results in calculations, statistics, and so on.

This task describes how to set up the **BACKPIsAdapter**.

This task includes the following steps:

- "Create an Adapter Instance" on page 286
- "View the Attributes in the CIT Manager" on page 287
- "Define a TQL" on page 288
- "Consume KPIs in External Application" on page 288

1 Create an Adapter Instance

In Business Service Management, in the integration point, create an adapter instance named **BACKPIsAdapter**. Use the following definitions:

- **Adapter.** Enter **BACKPIsAdapter**.
- **Is Integration Activated.** This checkbox should be selected.
- **Host name and IP.** Enter details of the BSM Gateway Server.
- **Port.** Enter 80.
- **Customer ID.** Enter 1.
- **Is CMS External.** Select **true** if you are working with an external CMS; select **false** if you are working with the RTSM.

In the second stage, choose the **Kpi** and **KpiObjective** CITs.

For details, see "Integration Point Pane" in the *RTSM Data Flow Management Guide*.

2 View the Attributes in the CIT Manager

The following attributes are used by the **Kpi** CIT:

- **KpiName.** The name of the KPI, for example, Availability, Performance.
- **Origin.** The origin of the KPIs, for example, HP Business Service Management.

The following attributes are used by the **KpiObjective** CIT:

- **Origin.** The origin of the KPIs, for example, HP Business Service Management.
- **KpiValue.** The numeric value of the KPI as calculated by the KPI business logic rule, for example, 0, 100, 3.2. (This value is optional.)
- **KpiStatus.** The status of the KPI as calculated by the business logic rule (can be one of the following values: ok, warning, minor, major, critical, none).
- **KpiUnitOfMeasure.** The measurement unit of the NumericValue field, for example, %, \$. (This value is optional.)
- **LastModifiedTime.** The date when the status last changed.
- **KpiComparisonOperator.** This attribute contains enum operators: na, =, >, <, <=, >=.
- **KpiThresholdOk.** The range of values that defines the OK KPI status.
- **KpiThresholdWarning.** The range of values that defines the WARNING KPI status.
- **KpiThresholdMinor.** The range of values that defines the MINOR KPI status.
- **KpiThresholdMajor.** The range of values that defines MAJOR KPI status.
- **KpiThresholdCritical.** The range of values that defines the CRITICAL KPI status.

For details on working with the CIT Manager, see "CI Type Manager" and "CI Type Attributes" in the *RTSM Modeling Guide*.

For details on KPI rules, see "Business Rule Repository" in *Using Service Health*.

3 Define a TQL

Define a TQL using the **Kpi** and **KpiObjective** CITs. For details, see "How to Define a TQL Query" in the *RTSM Modeling Guide*.

4 Consume KPIs in External Application

Using the API of your external application, send a query to run the relevant TQL in Business Service Management, to retrieve the status and value of the federated KPIs.

Part II

Using APIs

8

Introduction to APIs

This chapter includes:

Concepts

- ▶ APIs Overview on page 292

Concepts

APIs Overview

The following APIs are included with Business Service Management:

- ▶ **RTSM (HP Universal CMDB) API.** Explains how third-party or custom tools can use the Java API to extract data and calculations and to write data to the BSM Run-time Service Model (RTSM). For details, see "RTSM (HP Universal CMDB) API" on page 293.
- ▶ **RTSM (HP Universal CMDB) Web Service API.** Enables writing configuration item definitions and topological relations to the BSM Run-time Service Model (RTSM), and querying the information with TQL and ad hoc queries. For details, see "RTSM (HP Universal CMDB) Web Service API" on page 305.
- ▶ **Data Flow Management (Discovery and Dependency Mapping) Web Service API.** Enables managing probes, jobs, triggers and credentials for Data Flow Management. For details, see "Data Flow Management API" on page 385.
- ▶ **Internal-Global ID Conversion API.** Enables converting RTSM CI IDs to Global CI IDs, and Global CI IDs to RTSM CI IDs. For details, see "Internal-Global ID Conversion API" on page 407.

Note: To gain the full value of the API documentation, it is recommended to access the online documentation. The PDF version does not have the links into the API documentation that is generated in html format.

9

RTSM (HP Universal CMDB) API

This chapter includes:

Concepts

- ▶ Conventions on page 294
- ▶ Using the HP Universal CMDB API on page 294
- ▶ General Structure of an Application on page 296

Tasks

- ▶ How to Put the API Jar File in the Classpath on page 299
- ▶ How to Create an Integration User on page 299

Reference

- ▶ HP Universal CMDB API Reference on page 302
- ▶ Use Cases on page 302
- ▶ Examples on page 303

Concepts

Conventions

Note: The HP Universal CMDB Web Service API can be used with the Business Service Management Run-time Service Model (RTSM). References in this section to HP Universal CMDB (UCMDB) refer to the RTSM in the BSM context.

This chapter uses the following conventions:

- ▶ **UCMDB** refers to the Universal Configuration Management database itself. In the BSM context, it refers to the Run-time Service Model (RTSM). **HP Universal CMDB** refers to the application.
- ▶ UCMDB elements and method arguments are spelled in the case in which they are specified in the interfaces.

For full documentation on the available APIs, refer to the HP UCMDB API Reference.

These files are located in the following folder:

```
\\<BSM Gateway Server root directory>\AppServer\  
webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\UCMDB_JavaAPI  
\index.html
```

Using the HP Universal CMDB API

Use this chapter in conjunction with the API Javadoc, available in the online Documentation Library.

The HP Universal CMDB API is used to integrate applications with the Run-time Service Model (RTSM). The API provides methods to:

- add, remove, and update CIs and relations in the RTSM
- retrieve information about the class model
- retrieve information from the UCMDB history
- run what-if scenarios
- retrieve information about configuration items and relationships

Methods for retrieving information about configuration items and relationships generally use the Topology Query Language (TQL). For details, see "Topology Query Language" in the *RTSM Modeling Guide*.

Users of the HP Universal CMDB API should be familiar with:

- The Java programming language
- HP Universal CMDB

This section includes the following topics:

- "Uses of the API" on page 295
- "Permissions" on page 295

Uses of the API

The API is used to fulfill a number of business requirements. For example, a third-party system can query the class model for information about available configuration items (CIs). For more use cases, see "Use Cases" on page 302.

Permissions

The administrator provides login credentials for connecting with the API. The API client needs the user name and password of an integration user defined in the RTSM. These users do not represent human users of the RTSM, but rather applications that connect to the RTSM.

Caution: The API client can also work with regular users as long as they have API authentication permission. However, this option is not recommended.

For details, see "How to Create an Integration User" on page 299.

General Structure of an Application

There is only one static factory, the `UcmdbServiceFactory`. This factory is the entry point for an application. The `UcmdbServiceFactory` exposes `getServiceProvider` methods. These methods return an instance of the `UcmdbServiceProvider` interface.

The client creates other objects using interface methods. For example, to create a new query definition, the client:

- 1 gets the query service from the main RTSM service object
- 2 gets a query factory object from the service object
- 3 gets a new query definition from the factory

```
UcmdbServiceProvider provider =
    UcmdbServiceFactory.getServiceProvider(HOST_NAME, PORT);
UcmdbService ucmdbService =
    provider.connect(provider.createCredentials(USERNAME,
    PASSWORD), provider.createClientContext("Test"));
TopologyQueryService queryService = ucmdbService.getTopologyQueryService();
TopologyQueryFactory factory = queryService.getFactory();
QueryDefinition queryDefinition = factory.createQueryDefinition("Test Query");
queryDefinition.addNode("Node").ofType("host");
Topology topology = queryService.executeQuery(queryDefinition);
System.out.println("There are " + topology.getAllCIs().size() + " hosts in uCMDB");
```


The services available from **UcmdbService** are:

Service Methods	Use
getClassModelService	Information about types of CIs and relations
getConfigurationService	Infrastructure settings management, for server configuration
getDDMConfigurationService	Configure the Discovery and Dependency Management system
getDDMManagementService	Analyze and view the progress, results, and errors of the Discovery and Dependency Management system
getHistoryService	Information about history of monitored CIs (changes, removals, and so on)
getImpactAnalysisService	Run impact analysis scenario (also known as correlation).
getQueryManagementService	Manage access to queries - save, delete, list existing. Also provides query validation and queries dependencies discovery.
getResourceBundleManagement Service	Resource tagging ("bundling" services. Allows explicit creation of new tags and removal of tags from all tagged resources.
getStateService	Provide services for managing states (list, add, remove, and so on)
getSoftwareSignatureService	Define software items to be discovered by the Discovery and Dependency Management system
getSnapshotService	Provide services for managing snapshots (get, save, compare, and so on)
getTopologyQueryService	Get information about the IT universe
getTopologyUpdateService	Change information in the IT universe

Service Methods	Use
getViewService	View execution service (execute definition, execute saved) and management service (save, delete, list existing). Also provides view validation and dependencies discovery.
getViewArchiveService	View result archiving services. Allows saving the current view result and retrieving previously saved results.
SystemHealthService	Provide system health services (basic system performance indicators, capacity and availability metrics)

The client communicates with the server over HTTP.

Tasks

How to Put the API Jar File in the Classpath

The use of this API set requires the file **ucmdb-api.jar**. Get the jar from the <HP BSM root directory>\lib\ folder.

Put the jar file in the classpath before compiling or running your application.

Note: Usage of the UCMDB Java API Jar requires you to have JRE version 6 or later installed.

How to Create an Integration User

You can create a dedicated user for integrations between other products and BSM. This user enables a product that uses the BSM client SDK to be authenticated in the server SDK and execute the APIs. Applications written with this API set must log on with integration user credentials.

Caution: It is also possible to connect with a regular UCMDB user, (for instance, admin); however, this option is not recommended. To connect with a UCMDB user, you must grant it API authentication permission.

To create an integration user:

1 Launch a Web browser and enter the DPS server address, as follows:

`http://<FQDN of the Data Processing Server>:21212/jmx-console`

You may have to log in with a user name and password.

- 2 Under UCMDB, click **service=UCMDB Security Services**.
- 3 Locate the **CreateIntegrationUser** operation. This method accepts the following parameters:
 - **customerId**. The customer ID.
 - **username**. The integration user's name.
 - **password**. The integration user's password.
 - **dataStoreOrigin**. The name of the product that is going to use this integration user.

The following operations are useful for integration user management:

- **DeleteIntegrationUser**. Deletes the given integration user.
- **ExportIntegrationUser**. Exports the integration user to an XML file in the given path (on the server machine).
- **getIntegrationUser**. Displays the integration user information.
- **changeIntegrationUserPassword**. Changes the integration user's password.
- **canUserAuthenticate**. **isIntegrationUser** is **true**: can the integration user authenticate with the given credentials?

- 4 Click **Invoke**.

Either create more users, or close the JMX console.

- 5 Log on to BSM as an administrator.
- 6 From the **Administration** tab, run **Package Manager**.
- 7 Click the **New** icon.
- 8 Enter a name for the new package, and click **Next**.
- 9 In the Resource Selection tab, under **Administration**, click **Integration Users**.
- 10 Select a user or users that you created using the JMX console.
- 11 Click **Next** and then **Finish**. Your new package appears in the Package Name list in Package Manager.

12 Deploy the package to the users who will run the API applications.

For details, see "How to Deploy a Package" in the *RTSM Administration Guide*.

Note:

The integration user is per customer. To create a stronger integration user for cross-customer usage, use a **systemUser** with the **isSuperIntegrationUser** flag set to **true**. Use the **systemUser** methods (**createSystemUser**, **removeSystemUser**, **showAllSystemUsers**, **changeSystemUserPassword**, **canSuperIntegrationUserAuthenticate**, and so on).

There are two out-of-the-box system users; it is recommended to change their passwords after installation using the **changeSystemUserPassword** method.

- **sysadmin/sysadmin**
- **UISysadmin/UISysadmin** (This user is also the Super Integration User **SuperIntegrationUser**).

If you change the UISysadmin password using **changeSystemUserPassword**, you must execute the following method: in the JMX Console, locate the **UCMDB-UI:name=UCMDB Integration** service. Run **setCMDBSuperIntegrationUser** with the user name and new password of the integration user.

Reference

HP Universal CMDB API Reference

These files are located in the following folder:

```
\\<BSM Gateway Server root directory>\AppServer\  
webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\UCMDB_JavaAPI  
\index.html
```

Use Cases

The following use cases assume two systems:

- HP Universal CMDB server
- A third-party system that contains a repository of configuration items

This section includes the following topics:

- "Populating the RTSM" on page 302
- "Querying the RTSM" on page 303
- "Querying the Class Model" on page 303
- "Analyzing Change Impact" on page 303

Populating the RTSM

Use cases:

- A third-party asset management updates the RTSM with information available only in asset management
- A number of third-party systems populate the RTSM to create a central CMDB that can track changes and perform impact analysis
- A third-party system creates Configuration Items and Relations according to third-party business logic, to leverage the RTSM query capabilities

Querying the RTSM

Use cases:

- A third-party system gets the Configuration Items and Relations that represent the SAP system by retrieving the results of the SAP TQL
- A third-party system gets the list of Oracle servers that have been added or changed in the last five hours
- A third-party system gets the list of servers whose host name contains the lab substring
- A third-party system finds the elements related to a given CI by getting its neighbors

Querying the Class Model

Use cases:

- A third-party system enables users to specify the set of data to be retrieved from the RTSM. A user interface can be built over the class model to show users the possible properties and prompt them for required data. The user can then choose the information to be retrieved.
- A third-party system explores the class model when the user cannot access the BSM user interface.

Analyzing Change Impact

Use case:

A third-party system outputs a list of the business services that could be impacted by a change on a specified host.

Examples

See the following code samples:

- Create a Connection
- Create and Execute an Ad-Hoc Query

- ▶ Create and Execute a View
- ▶ Add and Delete Data
- ▶ Execute an Impact Analysis
- ▶ Query the Class Model
- ▶ Query a History Sample

These files are located in the following directory:

```
\\<BSM Gateway Server root directory>\AppServer\  
webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\JavaSDK_Samples\
```


10

RTSM (HP Universal CMDB) Web Service API

Note: The HP Universal CMDB Web Service API is the same API that is used with RTSM. References in this section to HP Universal CMDB (UCMDB) refer to the RTSM in the BSM context.

This chapter includes:

Concepts

- ▶ Conventions on page 307
- ▶ HP Universal CMDB Web Service API Overview on page 308
- ▶ HP Universal CMDB Web Service API Reference on page 310

Tasks

- ▶ How to Call the Web Service on page 311
- ▶ How to Query the RTSM on page 311
- ▶ How to Update the RTSM on page 316
- ▶ How to Query the BSM Class Model on page 318
- ▶ How to Query for Impact Analysis on page 320

Reference

- ▶ UCMDB General Parameters on page 321
- ▶ UCMDB Output Parameters on page 324
- ▶ UCMDB Query Methods on page 326

- ▶ UCMDB Update Methods on page 340
- ▶ UCMDB Impact Analysis Methods on page 343
- ▶ Use Cases on page 346
- ▶ Examples on page 347

Concepts

Conventions

Note: The HP Universal CMDB Web Service API can be used with the Business Service Management Run-time Service Model (RTSM). References in this section to HP Universal CMDB (UCMDB) refer to the RTSM in the BSM context.

This chapter uses the following conventions:

- ▶ **UCMDB** refers to the Universal Configuration Management database itself. In the BSM context, it refers to the Run-time Service Model (RTSM).
- ▶ BSM elements and method arguments are spelled in the case in which they are specified in the schema. An element or argument to a method is not capitalized. For example, a relation is an element of type `Relation` passed to a method.

For full documentation on the request and response structures, refer to the HP UCMDB Web Service API Reference. These files are located in the following folder:

```
\\<Business Service Management root directory>\AppServer\  
webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\CMDB_Schema  
\webframe.html
```

HP Universal CMDB Web Service API Overview

Note to HP Software-as-a-Service customers: For details on how to use the UCMDB Web Service API in an HP Software-as-a-Service environment, contact HP Software-as-a-Service Support.

Use this chapter in conjunction with the UCMDB schema documentation, available in the online Documentation Library.

The HP Universal CMDB Web Service API is used to integrate applications with the Business Service Management Run-time Service Model (RTSM). The API provides methods to:

- ▶ add, remove, and update CIs and relations in the RTSM
- ▶ retrieve information about the class model
- ▶ retrieve impact analyses
- ▶ retrieve information about configuration items and relationships
- ▶ manage credentials: view, add, update, and remove
- ▶ manage jobs: view status, activate, and deactivate
- ▶ manage Probe ranges: view, add, and update
- ▶ manage triggers: add or remove a trigger CI, and add, remove, or disable a trigger TQL
- ▶ view general data on domains and Probes

Methods for retrieving information about configuration items and relationships generally use the Topology Query Language (TQL). For details, see "Topology Query Language" in the *RTSM Modeling Guide*.

Users of the HP Universal CMDB Web Service API should be familiar with:

- ▶ The SOAP specification
- ▶ An object-oriented programming language such as C++, C# or Java
- ▶ HP Universal CMDB

► Data Flow Management

This section includes the following topics:

- "Uses of the API" on page 309
- "Permissions" on page 309

Uses of the API

The API is used to fulfill a number of business requirements. For example:

- A third-party system can query the class model for information about available configuration items (CIs).
- A third-party asset management tool can update the RTSM with information available only to that tool, thereby unifying its data with data collected by HP applications.
- A number of third-party systems can populate the RTSM to create a central RTSM that can track changes and perform impact analysis.
- A third-party system can create entities and relations according to its business logic, and then write the data to the RTSM to take advantage of the RTSM query capabilities.
- Other systems, such as the Release Control (CCM) system, can use the Impact Analysis methods for change analysis.

Permissions

The administrator provides login credentials for connecting with the Web Service. The required credentials depend on whether you are using HP Universal CMDB as a standalone application or from within Business Service Management:

- **HP Universal CMDB standalone.** Log in using the credentials of a BSM user who has been granted permissions on the discovery and integration resources.
- **HP Universal CMDB embedded in Business Service Management.** Log in using the credentials of a Business Service Management user. The user must have been granted the relevant permissions on the HP Universal CMDB resource in Business Service Management.

When permissions are assigned through HP Universal CMDB, the permission levels are View, Update, and Execute. When they are assigned using Business Service Management, the levels are View and Update, where Update also includes Execution. To view the permissions required for each operation, see each operation's request documentation.

HP Universal CMDB Web Service API Reference

For full documentation on the request and response structures, refer to the HP UCMDB Web Service API Reference. These files are located in the following folder:

```
\\<Business Service Management root directory>\AppServer\  
webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\CMDB_Schema  
\webframe.html
```

Tasks

How to Call the Web Service

You use standard SOAP programming techniques in the HP Universal CMDB Web Service to enable calling server-side methods. If the statement cannot be parsed or if there is a problem invoking the method, the API methods throw a `SoapFault` exception. When a `SoapFault` exception is thrown, BSM populates one or more of the error message, error code, and exception message fields. If there is no error, the results of the invocation are returned.

SOAP programmers can access the WSDL at:

[http://<server>\[:port\]/axis2/services/UcmdbService?wsdl](http://<server>[:port]/axis2/services/UcmdbService?wsdl)

The port specification is only necessary for non-standard installations. Consult your system administrator for the correct port number.

The URL for calling the service is:

[http://<server>\[:port\]/axis2/services/UcmdbService](http://<server>[:port]/axis2/services/UcmdbService)

For examples of connecting to the RTSM, see "Use Cases" on page 346.

How to Query the RTSM

The RTSM is queried using the APIs described in "UCMDB Query Methods" on page 326.

The queries and the returned RTSM elements always contain real BSM IDs.

For examples of the use of the query methods, see "Query Example" on page 351.

This section includes the following topics:

- "Just In Time Response Calculation" on page 312
- "Processing Large Responses" on page 312

- "Specifying Properties to Return" on page 313
- "Concrete Properties" on page 314
- "Derived Properties" on page 314
- "Naming Properties" on page 315
- "Other Property Specification Elements" on page 315

Just In Time Response Calculation

For all query methods, the BSM server calculates the values requested by the query method when the request is received, and returns results based on the latest data. The result is always calculated at the time the request is received, even if the TQL query is active and there exists a previously calculated result. Therefore, the results of running a query returned to the client application may be different to the results of the same query displayed on the user interface.

Tip: If your application uses the results of a given query more than once and the data is not expected to change significantly between uses of the result data, you can improve performance by having the client application store the data rather than repeatedly running the query.

Processing Large Responses

The response to a query always includes the structures for the data requested by the query method, even if no actual data is being transmitted. For many methods where the data is a collection or map, the response also includes the `ChunkInfo` structure, comprised of `chunksKey` and `numberOfChunks`. The `numberOfChunks` field indicates the number of chunks containing data that must be retrieved.

The maximum transmission size of data is set by the system administrator. If the data returned from the query is larger than the maximum size, the data structures in the first response contain no meaningful information, and the value of the `numberOfChunks` field is 2 or greater. If the data is not larger than the maximum, the `numberOfChunks` field is 0 (zero), and the data is transmitted in the first response. Therefore, in processing a response, check the `numberOfChunks` value first. If it is greater than 1, discard the data in the transmission and request the chunks of data. Otherwise, use the data in the response.

For information on handling chunked data, see "pullTopologyMapChunks" on page 338 and "releaseChunks" on page 339.

Specifying Properties to Return

CIs and relations generally have many properties. Some methods that return collections or graphs of these items accept input parameters that specify which property values to return with each item that matches the query. The RTSM does not return empty properties. Therefore, the response to a query may have fewer properties than requested in the query.

This section describes the types of sets used to specify the properties to return.

Properties can be referenced in two ways:

- ▶ By their names
- ▶ By using names of predefined properties rules. Predefined properties rules are used by the RTSM to create a list of real property names.

When an application references properties by name, it passes a `PropertiesList` element.

Tip: Whenever possible, use `PropertiesList` to specify the names of the properties in which you are interested, rather than a rule-based set. The use of predefined properties rules nearly always results in returning more properties than needed, and bears a performance price.

There are two types of predefined properties: qualifier properties and simple properties.

- ▶ **Qualifier properties.** Use when the client application should pass a `QualifierProperties` element (a list of qualifiers that can be applied to properties). The RTSM converts the list of qualifiers passed by the client application to the list of the properties to which at least one of the qualifiers applies. The values of these properties are returned with the `CI` or `Relation` elements.
- ▶ **Simple properties.** To use simple rule-based properties, the client application passes a `SimplePredefinedProperty` or `SimpleTypedPredefinedProperty` element. These elements contain the name of the rule by which the RTSM generates the list of properties to return. The rules that can be specified in a `SimplePredefinedProperty` or `SimpleTypedPredefinedProperty` element are `CONCRETE`, `DERIVED`, and `NAMING`.

Concrete Properties

Concrete properties are the set of properties defined for the specified CIT. The properties added by derived classes are not returned for instances of those derived classes.

A collection of instances returned by a method may consist of instances of a CIT specified in the method invocation and instances of CITs that inherit from that CIT. The derived CITs inherit the properties of the specified CIT. In addition, the derived CITs extend the parent CIT by adding properties.

Example of Concrete Properties:

CIT T1 has properties P1 and P2. CIT T11 inherits from T1 and extends T1 with properties P21 and P22.

The collection of CIs of type T1 includes the instances of T1 and T11. The concrete properties of all instances in this collection are P1 and P2.

Derived Properties

Derived properties are the set of properties defined for the specified CIT and, for each derived CIT, the properties added by the derived CIT.

Example of Derived Properties:

Continuing the example from concrete properties, the derived properties of instances of T1 are P1 and P2. The derived properties of instances of T11 are P1, P2, P21, and P22.

Naming Properties

The naming properties are `display_label` and `data_name`.

Other Property Specification Elements**► PredefinedProperties**

`PredefinedProperties` can contain a `QualifierProperties` element and a `SimplePredefinedProperty` element for each of the other possible rules. A `PredefinedProperties` set does not necessarily contain all types of lists.

► PredefinedTypedProperties

`PredefinedTypedProperties` is used to apply a different set of properties to each CIT. `PredefinedTypedProperties` can contain a `QualifierProperties` element and a `SimpleTypedPredefinedProperty` element for each of the other applicable rules. Because `PredefinedTypedProperties` is applied to each CIT individually, derived properties are not relevant. A `PredefinedProperties` set does not necessarily contain all applicable types of lists.

► CustomProperties

`CustomProperties` can contain any combination of the basic `PropertiesList` and the rule-based property lists. The properties filter is the union of all the properties returned by all the lists.

► CustomTypedProperties

`CustomTypedProperties` can contain any combination of the basic `PropertiesList` and the applicable rule-based property lists. The properties filter is the union of all the properties returned by all the lists.

► **TypedProperties**

TypedProperties is used to pass a different set of properties for each CIT. **TypedProperties** is a collection of pairs composed of type names and properties sets of all types. Each properties set is applied only to the corresponding type.

How to Update the RTSM

You update the RTSM with the update APIs.

For details of the API methods, see "UCMDB Update Methods" on page 340.

For examples of the use of the update methods, see "Update Example" on page 368.

This task includes the following steps:

- "UCMDB Update Parameters" on page 316
- "Use of ID Types with Update Methods" on page 317

UCMDB Update Parameters

This topic describes the parameters used only by the service's update methods. For details, see the schema documentation.

CIsAndRelationsUpdates

The **CIsAndRelationsUpdates** type consists of **CIsForUpdate**, **relationsForUpdate**, **referencedRelations**, and **referencedCIs**. A **CIsAndRelationsUpdates** instance does not necessarily include all three elements.

CIsForUpdate is a **CIs** collection. **relationsForUpdate** is a **Relations** collection. The CI and relation elements in the collections have a **props** element. When creating a CI or relation, properties that have either the **required** attribute or the **key** attribute in the CI Type definition must be populated with values. The items in these collections are updated or created by the method.

`referencedCIs` and `referencedRelations` are collections of CIs that are already defined in the RTSM. The elements in the collection are identified with a temporary ID in conjunction with all the key properties. These items are used to resolve the identities of CIs and relations for update. They are never created or updated by the method.

Each of the CI and relation elements in these collections has a `properties` collection. New items are created with the property values in these collections.

Use of ID Types with Update Methods

The following describes ID CIs, and CIs and relations. When the ID is not a real RTSM ID, the type and key attributes are required.

Deleting or Updating Configuration Items

A temporary or empty ID may be used by the client when calling a method to delete or update an item. In this case, the CI type and the key attributes that identify the CI must be set.

Deleting or Updating Relations

When deleting or updating relations, the relation ID can be empty, temporary, or real.

If a CI's ID is temporary, the CI must be passed in the `referencedCIs` collection and its key attributes must be specified. For details, see `CIsAndrelationsUpdates` in "UCMDB Update Parameters" on page 316.

Inserting New Configuration Items into the RTSM

It is possible to use either an empty ID or a temporary ID to insert a new CI. However, if the ID is empty, the server cannot return the real RTSM ID in the structure `createIDsMap` because there is no `clientID`. For details, see "addCIsAndRelations" on page 340 and "UCMDB Query Methods" on page 326.

Inserting New Relations into the RTSM

The relation ID can be either temporary or empty. However, if the relation is new but the configuration items on either end of the relation are already defined in the RTSM, then those CIs that already exist must be identified by a real RTSM ID or be specified in a `referencedCIs` collection. For details, see `CIsAndrelationsUpdates` in "UCMDB Update Parameters" on page 316.

How to Query the BSM Class Model

The class model methods return information about CITs and relations. The class model is configured using the CI Type Manager. For details, see "CI Type Manager" in the *RTSM Modeling Guide*.

For examples of the use of the class model methods, see "Class Model Example" on page 372.

This section provides information on the following methods that return information about CITs and relations:

- "getClassAncestors" on page 318
- "getAllClassesHierarchy" on page 319
- "getCmdbClassDefinition" on page 319

getClassAncestors

The `getClassAncestors` method retrieves the path between the given CIT and its root, including the root.

Input

Parameter	Comment
<code>cmdbContext</code>	For details, see "CmdbContext" on page 321.
<code>className</code>	The type name. For details, see "Type Name" on page 323.

Output

Parameter	Comment
classHierarchy	A collection of pairs of class names and parent class name.
comments	For internal use only.

getAllClassesHierarchy

The `getAllClassesHierarchy` method retrieves the entire class model tree.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.

Output

Parameter	Comment
classesHierarchy	A collection of pairs of class name and parent class name.
comments	For internal use only.

getCmdbClassDefinition

The `getCmdbClassDefinition` method retrieves information about the specified class.

If you use `getCmdbClassDefinition` to retrieve the key attributes, you must also query the parent classes up to the base class. `getCmdbClassDefinition` identifies as key attributes only those attributes with the `ID_ATTRIBUTE` set in the class definition specified by `className`. Inherited key attributes are not recognized as key attributes of the specified class. Therefore, the complete list of key attributes for the specified class is the union of all the keys of the class and of all its parents, up to the root.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
className	The type name. For details, see "Type Name" on page 323.

Output

Parameter	Comment
cmdbClass	The class definition, consisting of name, classType, displayLabel, description, parentName, qualifiers, and attributes.
comments	For internal use only.

 **How to Query for Impact Analysis**

The Identifier in the impact analysis methods points to the service's response data. It is unique for the current response and is discarded from the server's memory cache after 10 minutes of non-use.

For examples of the use of the impact analysis methods, see "Impact Analysis Example" on page 374.

Reference

UCMDB General Parameters

This section describes the most common parameters of the service's methods. For details, refer to the schema documentation.

This section includes the following topics:

- "CmdbContext" on page 321
- "ID" on page 321
- "Key Attributes" on page 321
- "ID Types" on page 322
- "CIProperties" on page 322
- "Type Name" on page 323
- "Configuration Item (CI)" on page 323
- "Relation" on page 323

CmdbContext

All UCMDB Web Service API service invocations require a `CmdbContext` argument. `CmdbContext` is a `callerApplication` string that identifies the application that invokes the service. `CmdbContext` is used for logging and troubleshooting.

ID

Every CI and Relation has an ID field. It consists of a case-sensitive ID string and an optional temp flag, indicating whether the ID is temporary.

Key Attributes

For identifying a CI or Relation in some contexts, key attributes can be used in place of a RTSM ID. Key attributes are those attributes with the `ID_ATTRIBUTE` set in the class definition.

In the user interface, the key attributes have a key icon next to them in the list of Configuration Item Type attributes in the user interface. For details, see "Add/Edit Attribute Dialog Box" in the *RTSM Modeling Guide*. For information about identifying the key attributes from within the API client application, see "getCmdbClassDefinition" on page 319.

ID Types

An ID element can contain a real ID, or a temporary ID.

A real ID is a string assigned by the RTSM that identifies an entity in the database. A temporary ID can be any string that is unique in the current request.

A temporary ID can be assigned by the client and often represents the ID of the CI as stored by the client. It does not necessarily represent an entity already created in the RTSM. When a temporary ID is passed by the client, if the RTSM can identify an existing data configuration item using the CI key properties, that CI is used as appropriate for the context as though it had been identified with a real ID.

CIProperties

A CIProperties element is composed of collections, each containing a sequence of name-value elements that specify properties of the type indicated by the collection name. None of the collections are required, so the CIProperties element can contain any combination of collections.

CIProperties are used by CI and Relation elements. For details, see "Configuration Item (CI)" on page 323 and "Relation" on page 323.

The properties collections are:

- ▶ dateProps - collection of DateProp elements
- ▶ doubleProps - collection of DoubleProp elements
- ▶ floatProps - collection of FloatProp elements
- ▶ intListProps - collection of intListProp elements
- ▶ intProps - collection of IntProp elements
- ▶ strProps - collection of StrProp elements

- strListProps - collection of StrListProp elements
- longProps - collection of LongProp elements
- bytesProps - collection of BytesProp elements
- xmlProps - collection of XmlProp elements

Type Name

The type name is the class name of a configuration item type or relation type. The type name is used in code to refer to the class. It should not be confused with the display name, which is seen on the user interface where the class is mentioned, but which is meaningless in code.

Configuration Item (CI)

A CI element is composed of an ID, a type, and a props collection.

When using UCMDB Update Methods to update a CI, the ID element can contain a real RTSM ID or a client-assigned temporary ID. If a temporary ID is used, set the temp flag to true. When deleting an item, the ID can be empty. UCMDB Query Methods take real IDs as input parameters and return real IDs in the query results.

The type can be any type name defined in the CI Type Manager. For details, see "CI Type Manager" in the *RTSM Modeling Guide*.

The props element is a CIProperties collection. For details, see "CIProperties" on page 322.

Relation

A Relation is an entity that links two configuration items. A Relation element is composed of an ID, a type, the identifiers of the two items being linked (end1ID and end2ID), and a props collection.

When using UCMDB Update Methods to update a Relation, the value of the Relation's ID can be a real RTSM ID or a temporary ID. When deleting an item, the ID can be empty. UCMDB Query Methods take real IDs as input parameters and return real IDs in the query results.

The relation type is the **Type Name** of the BSM class from which the relation is instantiated. The type can be any of the relation types defined in the RTSM. For further information on classes or types, see "How to Query the BSM Class Model" on page 318.

For details, see "CI Type Manager" in the *RTSM Modeling Guide*.

The two relation end IDs must not be empty IDs because they are used to create the ID of the current relation. However, they both can have temporary IDs assigned to them by the client.

The props element is a CIProperties collection. For details, see "CIProperties" on page 322.

UCMDB Output Parameters

This section describes the most common output parameters of the service methods. For details, refer to the schema documentation.

This section includes the following topics:

- "CIs" on page 324
- "ShallowRelation" on page 325
- "Topology" on page 325
- "CINode" on page 325
- "RelationNode" on page 325
- "TopologyMap" on page 325
- "ChunkInfo" on page 326

CIs

CIs is a collection of CI elements.

ShallowRelation

A **ShallowRelation** is an entity that links two configuration items, composed of an ID, a type, and the identifiers of the two items being linked (end1ID and end2ID). The relation type is the Type Name of the RTSM class from which the relation is instantiated. The type can be any of the relation types defined in the RTSM.

Topology

Topology is a graph of CI elements and relations. A **Topology** consists of a **CIs** collection and a **Relations** collection containing one or more **Relation** elements.

CINode

CINode is composed of a **CIs** collection with a **label**. The **label** in the **CINode** is the **label** defined in the node of the TQL used in the query.

RelationNode

RelationNode is a set of **Relations** collections with a **label**. The **label** in the **RelationNode** is the **label** defined in the node of the TQL used in the query.

TopologyMap

TopologyMap is the output of a query calculation that matches a TQL query. The **labels** in the **TopologyMap** are the node **labels** defined in the TQL used in the query.

The data of **TopologyMap** is returned in the following form:

- **CINodes**. This is one or more **CINode** (see "CINode" on page 325).
- **relationNodes**. This is one or more **RelationNode** (see "RelationNode" on page 325).

The **labels** in these two structures order the lists of configuration items and relations.

ChunkInfo

When a query returns a large amount of data, the server stores the data, divided into segments called chunks. The information the client uses to retrieve the chunked data is located in the `ChunkInfo` structure returned by the query. `ChunkInfo` is composed of the `numberOfChunks` that must be retrieved and the `chunksKey`. The `chunksKey` is a unique identifier of the data on the server for this specific query invocation.

For more information, see "Processing Large Responses" on page 312.

UCMDB Query Methods

This section provides information on the following methods:

- "executeTopologyQueryByNameWithParameters" on page 326
- "executeTopologyQueryWithParameters" on page 327
- "getChangedCIs" on page 328
- "getCINeighbours" on page 329
- "getCIsByID" on page 330
- "getCIsByType" on page 331
- "getFilteredCIsByType" on page 331
- "getQueryNameOfView" on page 336
- "getTopologyQueryExistingResultByName" on page 336
- "getTopologyQueryResultCountByName" on page 337
- "pullTopologyMapChunks" on page 338
- "releaseChunks" on page 339

executeTopologyQueryByNameWithParameters

The `executeTopologyQueryByNameWithParameters` method retrieves a `topologyMap` element that matches the specified parameterized query.

The values for the query parameters are passed in the `parameterizedNodes` argument. The specified TQL must have unique labels defined for each `CINode` and each `relationNode` or the method invocation fails.

Input

Parameter	Comment
<code>cmdbContext</code>	For details, see "CmdbContext" on page 321.
<code>queryName</code>	The name of the parameterized TQL in the RTSM for which to get the map.
<code>parameterizedNodes</code>	The conditions each node must meet to be included in the query results.
<code>queryTypedProperties</code>	A collection of sets of properties to retrieve to items of a specific Configuration Item Type.

Output

Parameter	Comment
<code>topologyMap</code>	For details, see "TopologyMap" on page 325.
<code>chunkInfo</code>	For details, see: "ChunkInfo" on page 326, "Processing Large Responses" on page 312.

executeTopologyQueryWithParameters

The `executeTopologyQueryWithParameters` method retrieves a `topologyMap` element that matches the parameterized query.

The query is passed in the `queryXML` argument. The values for the query parameters are passed in the `parameterizedNodes` argument. The TQL must have unique labels defined for each `CINode` and each `relationNode`.

The `executeTopologyQueryWithParameters` method is used to pass ad-hoc queries, rather than accessing a query defined in the RTSM. You can use this method when you do not have access to the BSM user interface to define a query, or when you do not want to save the query to the database.

To use an exported TQL as the input to this method, do the following:

- 1 Launch the Web browser and enter the following address:
http://localhost:8080/jmx-console.

You may have to log in with a user name and password. The default is **sysadmin/sysadmin**
- 2 Click **UCMDB:service=TQL Services.**
- 3 Locate the **exportTql** operation.
 - In the **customerId** parameter box, enter **1** (the default).
 - In the **patternName** parameter box, enter a valid TQL name.
- 4 Click **Invoke.**

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
queryXML	An XML string representing a TQL without resource tags.
parameterizedNodes	The conditions each node must meet to be included in the query results.

Output

Parameter	Comment
topologyMap	For details, see "TopologyMap" on page 325.
chunkInfo	For details, see "ChunkInfo" on page 326 and "Processing Large Responses" on page 312.

getChangedCIs

The `getChangedCIs` method returns the change data for all CIs related to the specified CIs.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
ids	The list of the IDs of the root CIs whose related CIs are checked for changes. Only real RTSM IDs are valid in this collection.
fromDate	The beginning of the period in which to check if CIs changed.
toDate	The end of the period in which to check if CIs changed.

Output

Parameter	Comment
changeDataInfo	Zero or more collections of ChangedDataInfo elements.

getCIneighbours

The `getCIneighbours` method returns the immediate neighbors of the specified CI.

For example, if the query is on the neighbors of CI A, and CI A contains CI B which uses CI C, CI B is returned, but CI C is not. That is, only neighbors of the specified type are returned.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
ID	The ID of the CI with which to retrieve the neighbors. This must be a real RTSM ID.

Parameter	Comment
neighbourType	The CIT name of the neighbors to retrieve. Neighbors of the specified type and of types derived from that type are returned. For details, see "Type Name" on page 323.
CIProperties	The data to be returned on each configuration item, called the Query Layout in the user interface. For details, see "TypedProperties" on page 316.
relationProperties	The data to be returned on each relation (called the Query Layout in the user interface). For details, see "TypedProperties" on page 316.

Output

Parameter	Comment
topology	For details, see "Topology" on page 325.
comments	For internal use only.

getCIsByID

The getCIsByID method retrieves configuration items by their RTSM IDs.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
CIsTypedProperties	A typed properties collection. For details, see "Other Property Specification Elements" on page 315.
IDs	Only real RTSM IDs are valid in this collection.

Output

Parameter	Comment
CIs	Collection of CI elements.
chunkInfo	For details, see: "ChunkInfo" on page 326 and "Processing Large Responses" on page 312.

getCIsByType

The `getCIsByType` method returns the collection of configuration items of the specified type and of all types that inherit from the specified type.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
type	The class name. For details, see "Type Name" on page 323.
properties	The data to be returned on each configuration item. For details, see "CustomProperties" on page 315.

Output

Parameter	Comment
CIs	Collection of CI elements.
chunkInfo	For details, see: "ChunkInfo" on page 326 and "Processing Large Responses" on page 312.

getFilteredCIsByType

The `getFilteredCIsByType` method retrieves the CIs of the specified type that meet the conditions used by the method. A condition is comprised of:

- a name field containing the name of a property
- an operator field containing a comparison operator

- an optional value field containing a value or list of values

Together, they form a Boolean expression:

```
<item>.property.value [operator] <condition>.value
```

For example, if the condition name is `root_actualdeletionperiod`, the condition value is `40` and the operator is `Equal`, the Boolean statement is:

```
<item>.root_actualdeletionperiod.value == 40
```

The query returns all items whose `root_actualdeletionperiod` is `40`, assuming there are no other conditions.

If the `conditionsLogicalOperator` argument is `AND`, the query returns the items that meet all conditions in the `conditions` collection. If `conditionsLogicalOperator` is `OR`, the query returns the items that meet at least one of the conditions in the `conditions` collection.

The following table lists the comparison operators:

Operator	Type of Condition/Comments
ChangedDuring	<p>Date</p> <p>This is a range check. The condition value is specified in hours. If the value of the date property lies in the range of the time the method is invoked plus or minus the condition value, the condition is true.</p> <p>For example, if the condition value is <code>24</code>, the condition is true if the value of the date property is between yesterday at this time and tomorrow at this time.</p> <p>Note: The name <code>ChangedDuring</code> is kept to preserve backward compatibility. In previous versions, the operator was used only with create and modify time properties.</p>
Equal	String and numerical
EqualIgnoreCase	String

Operator	Type of Condition/Comments
Greater	Numerical
GreaterEqual	Numerical
In	String, numerical, and list The condition's value is a list. The condition is true if the value of the property is one of the values in the list.
InList	List The condition's value and the property's value are lists. The condition is true if all the values in the condition's list also appear in the item's property list. There can be more property values than specified in the condition without affecting the truth of the condition.
IsNull	String, numerical, and list The item's property has no value. When operator IsNull is used, the value of the condition is ignored, and in some cases can be nil.
Less	Numerical
LessEqual	Numerical
Like	String The condition's value is a substring of the value of the property's value. The condition's value must be bracketed with percentage signs (%). For example, %Bi% matches Bismark and Bay of Biscay, but not biscuit.
LikeIgnoreCase	String Use the LikeIgnoreCase operator as you use the Like operator. The match, however is not case-sensitive. Therefore, %Bi% matches biscuit.

Operator	Type of Condition/Comments
NotEqual	String and numerical
UnchangedDuring	<p>Date</p> <p>This is a range check. The condition value is specified in hours. If the value of the date property is in the range of the time the method is invoked plus or minus the condition value, the condition is false. If it lies outside that range, the condition is true.</p> <p>For example, if the condition value is 24, the condition is true if the value of the date property is before yesterday at this time or after tomorrow at this time.</p> <p>Note: The name UnchangedDuring is kept to preserve backward compatibility. In previous versions, the operator was used only with create and modify time properties.</p>

Example of Setting Up a Condition:

```
FloatCondition fc = new FloatCondition();
FloatProp fp = new FloatProp();
fp.setName("attr_name");
fp.setValue(11);
fc.setCondition(fp);
fc.setFloatOperator(FloatCondition.floatOperatorEnum.Equal);
```

Example of Querying for Inherited Properties:

The target CI is `sample` which has two attributes, `name` and `size`. `samplell` extends the CI with two attributes, `level` and `grade`. This example sets up a query for the properties of `samplell` that were inherited from `sample` by specifying them by name.

```
GetFilteredCIsByType request = new GetFilteredCIsByType()
request.setCmdbContext(cmdbContext)
request.setType("samplell")
CustomProperties customProperties = new CustomProperties();
PropertiesList propertiesList = new PropertiesList();
propertiesList.addPropertyName("name");
propertiesList.addPropertyName("size");
customProperties.setPropertiesList(propertiesList);
request.setProperties(customProperties)
```

Input

Parameter	Comment
<code>cmdbContext</code>	For details, see "CmdbContext" on page 321.
<code>type</code>	The class name. For details, see "Type Name" on page 323. The type can be any of the types defined using the CI Type Manager. For details, see "CI Type Manager" in the <i>RTSM Modeling Guide</i> .
<code>properties</code>	The data to be returned on each CI (called the Query Layout in the user interface). For details, see "CustomProperties" on page 315.
<code>conditions</code>	A collection of name-value pairs and the operators that relate one to the other. For example, <code>host_hostname</code> like QA.
<code>conditionsLogicalOperator</code>	<ul style="list-style-type: none"> ➤ AND. All the conditions must be met. ➤ OR. At least one of the conditions must be met.

Output

Parameter	Comment
CIs	Collection of CI elements.
chunkInfo	For details, see "ChunkInfo" on page 326 and "Processing Large Responses" on page 312.

getQueryNameOfView

The `getQueryNameOfView` method retrieves the name of the TQL on which the specified view is based.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
viewName	The name of a view, that is, a sub-set of the class model in the RTSM.

Output

Parameter	Comment
queryName	The name of the TQL in the RTSM on which the view is based.

getTopologyQueryExistingResultByName

The `getTopologyQueryExistingResultByName` method retrieves the most recent result of running the specified TQL. The call does not run the TQL. If there are no results from a previous run, nothing is returned.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
queryName	The name of a TQL.
queryTypedProperties	A collection of sets of properties to retrieve for items of a specific Configuration Item Type.

Output

Parameter	Comment
queryName	The name of the TQL in the RTSM on which the view is based.



getTopologyQueryResultCountByName

The `getTopologyQueryResultCountByName` method retrieves the number of instances of each node that matches the specified query.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
queryName	The name of a TQL.
countInvisible	If true, the output includes CIs defined as invisible in the query.

Output

Parameter	Comment
queryName	The name of the TQL in the RTSM on which the view is based.

pullTopologyMapChunks

The pullTopologyMapChunks method retrieves one of the chunks that contain the response to a method.

Each chunk contains a topologyMap element that is part of the response. The first chunk is numbered 1, so the retrieval loop counter iterates from 1 to `<response object>.getChunkInfo().getNumberOfChunks()`.

For details, see "ChunkInfo" on page 326 and "How to Query the RTSM" on page 311.

The client application must be able to handle the partial maps. See the following example of handling a CI collection and the example of merging chunks to a map in "Query Example" on page 351.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
ChunkRequest	The number of the chunk to retrieve and the ChunkInfo that is returned by the query method.

Output

Parameter	Comment
topologyMap	For details, see "TopologyMap" on page 325.
comments	For internal use only.

Example of Handing Chunks:

```

GetClsByType request =
    new GetClsByType(cmdbContext, typeName, customProperties);
GetClsByTypeResponse response =
    ucmdbService.getClsByType(request);
ChunkRequest chunkRequest = new ChunkRequest();
chunkRequest.setChunkInfo(response.getChunkInfo());
for(int j=1 ; j < response.getChunkInfo().getNumberOfChunks() ; j++) {
    chunkRequest.setChunkNumber(j);
    PullTopologyMapChunks req = new PullTopologyMapChunks(cmdbContext,
    chunkRequest);
    PullTopologyMapChunksResponse res =
        ucmdbService.pullTopologyMapChunks(req);
    for(int m=0 ;
        m < res.getTopologyMap().getCINodes().sizeCINodeList() ;
        m++) {
        Cls cis =
            res.getTopologyMap().getCINodes().getCINode(m).getCls();
        for(int i=0 ; i < cis.sizeCICollection() ; i++) {
            // your code to process the Cls
        }
    }
}
}

```

 **releaseChunks**

The `releaseChunks` method frees the memory of the chunks that contain the data from the query.

Tip: The server discards the data after ten minutes. Calling this method to discard the data as soon as it has been read conserves server resources.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
chunksKey	The identifier of the data on the server that was chunked. The key is an element of ChunkInfo.

UCMDB Update Methods

This section provides information on the following methods:

- "addCIsAndRelations" on page 340
- "addCustomer" on page 341
- "deleteCIsAndRelations" on page 342
- "removeCustomer" on page 342
- "updateCIsAndRelations" on page 342

addCIsAndRelations

The `addCIsAndRelations` method adds or updates CIs and relations.

If the CIs or relations do not exist in the RTSM, they are added and their properties are set according to the contents of the `CIsAndRelationsUpdates` argument.

If the CIs or relations do exist in the RTSM, they are updated with the new data, if `updateExisting` is **true**.

If `updateExisting` is **false**, `CIsAndRelationsUpdates` cannot reference existing configuration items or relations. Any attempt to reference existing items when `updateExisting` is **false** results in an exception.

If `updateExisting` is **true**, the add or update operation is performed without validating the CIs, regardless of the value of `ignoreValidation`.

If `updateExisting` is **false** and `ignoreValidation` is **true**, the add operation is performed without validating the CIs.

If `updateExisting` is **false** and `ignoreValidation` is **false**, the CIs are validated before the add operation.

Relations are never validated.

`CreatedIDsMap` is a map or dictionary of type `ClientIDToCmdbID` that connects the client's temporary IDs with the corresponding real RTSM IDs.

Input

Parameter	Comment
<code>cmdbContext</code>	For details, see "CmdbContext" on page 321.
<code>updateExisting</code>	Set to <i>true</i> to update items that already exist in the RTSM. Set to <i>false</i> to throw an exception if any item already exists.
<code>CIsAndRelationsUpdates</code>	The items to update or create. For details, see "CIsAndRelationsUpdates" on page 316.
<code>ignoreValidation</code>	If true, no check is performed before updating the RTSM.

Output

Parameter	Comment
<code>CreatedIDsMap</code>	The map of client IDs to RTSM IDs.
<code>comments</code>	For internal use only.

addCustomer

The `addCustomer` method adds a customer.

Input

Parameter	Comment
<code>CustomerID</code>	The numeric ID of the customer.

deleteCIsAndRelations

The `deleteCIsAndRelations` method removes the specified configuration items and relations from the RTSM.

When a CI is deleted and the CI is one end of one or more `Relation` items, those `Relation` items are also deleted.

Input

Parameter	Comment
<code>cmdbContext</code>	For details, see "CmdbContext" on page 321.
<code>CIsAndRelationsUpdates</code>	The items to delete. For details, see "CIsAndRelationsUpdates" on page 316

removeCustomer

The `removeCustomer` method deletes a customer record.

Input

Parameter	Comment
<code>CustomerID</code>	The numeric ID of the customer.

updateCIsAndRelations

The `updateCIsAndRelations` method updates the specified CIs and relations.

Update uses the property values from the `CIsAndRelationsUpdates` argument. If any of the CIs or relations do not exist in the RTSM, an exception is thrown.

`CreatedIDsMap` is a map or dictionary of type `ClientIDToCmdbID` that connects the client's temporary IDs with the corresponding real RTSM IDs.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
CIsAndRelationsUpdates	The items to update. For details, see "CIsAndRelationsUpdates" on page 316.
ignoreValidation	If true, no check is performed before updating the RTSM.

Output

Parameter	Comment
CreatedIDsMap	The map of client IDs to RTSM IDs. For details, see "addCIsAndRelations" on page 340.

UCMDB Impact Analysis Methods

This section provides information on the following methods:

- "calculateImpact" on page 343
- "getImpactPath" on page 344
- "getImpactRulesByNamePrefix" on page 345

calculateImpact

The calculateImpact method calculates which CIs are affected by a given CI according to the rules defined in the RTSM.

This shows the effect of an event triggering of the rule. The identifier output of calculateImpact is used as input for getImpactPath.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
impactCategory	The type of event that would trigger the rule being simulated.
IDs	A collection of ID elements.
impactRulesNames	A collection of ImpactRuleName elements.
severity	The severity of the triggering event.

Output

Parameter	Comment
impactTopology	For details, see "Topology" on page 325.
identifier	The key to the server response.

 **getImpactPath**

The `getImpactPath` method retrieves the topology graph of the path between the affected CI and the CI that affects it.

The `identifier` output of `calculateImpact` is used as the `identifier` input argument of `getImpactPath`.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
identifier	The key to the server response that was returned by <code>calculateImpact</code> .
relation	A Relation based on one of the ShallowRelations returned by <code>calculateImpact</code> in the <code>impactTopology</code> element.

Output

Parameter	Comment
impactPathTopology	A CIs collection and an ImpactRelations collection.
comments	For internal use only.

An ImpactRelations element consists of an ID, type, end1ID, end2ID, a rule, and an action.

 **getImpactRulesByNamePrefix**

The getImpactRulesByNamePrefix method retrieves rules using a prefix filter.

This method applies to impact rules that are named with a prefix that indicates the context to which they apply, for example, SAP_myrule, ORA_myrule, and so on. This method filters all impact rule names for those beginning with the prefix specified by the ruleNamePrefixFilter argument.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 321.
ruleNamePrefixFilter	A string containing the first letters of the rule names to match.

Output

Parameter	Comment
impactRules	impactRules is composed of zero or more impactRule. An impactRule, which specifies the effect of a change, is composed of ruleName, description, queryName, and isActive.

Use Cases

The following use cases assume two systems:

- ▶ Business Service Management server
- ▶ A third-party system that contains a repository of configuration items

This section includes the following topics:

- ▶ "Populating the RTSM" on page 346
- ▶ "Querying the RTSM" on page 346
- ▶ "Querying the Class Model" on page 347
- ▶ "Analyzing Change Impact" on page 347

Populating the RTSM

Use cases:

- ▶ A third-party asset management updates the RTSM with information available only in asset management
- ▶ A number of third-party systems populate the RTSM to create a central RTSM that can track changes and perform impact analysis
- ▶ A third-party system creates Configuration Items and Relations according to third-party business logic to leverage the RTSM query capabilities

Querying the RTSM

Use cases:

- ▶ A third-party system gets the Configuration Items and Relations that represent the SAP system by getting the results of the SAP TQL
- ▶ A third-party system gets the list of Oracle servers that have been added or changed in the last five hours
- ▶ A third-party system gets the list of servers whose host name contains the substring *lab*
- ▶ A third-party system finds the elements related to a given CI by getting its neighbors

Querying the Class Model

Use cases:

- A third-party system enables users to specify the set of data to be retrieved from the RTSM. A user interface can be built over the class model to show users the possible properties and prompt them for required data. The user can then choose the information to be retrieved.
- A third-party system explores the class model when the user cannot access the BSM user interface.

Analyzing Change Impact

Use case:

A third-party system outputs a list of the business services that could be impacted by a change on a specified host.

Examples

This section includes the following topics:

- "Example Base Class" on page 348
- "Query Example" on page 351
- "Update Example" on page 368
- "Class Model Example" on page 372
- "Impact Analysis Example" on page 374
- "Adding Credentials Example" on page 379

 **Example Base Class**

```
package com.hp.ucmdb.demo;

import com.hp.ucmdb.generated.services.UcmdbService;
import com.hp.ucmdb.generated.services.UcmdbServiceStub;
import com.hp.ucmdb.generated.types.CmdbContext;
import org.apache.axis2.AxisFault;
import org.apache.axis2.transport.http.HTTPConstants;
```

```
import org.apache.axis2.transport.http.HttpTransportProperties;

import java.net.MalformedURLException;
import java.net.URL;
```

```
/**
 * User: hbarkai
 * Date: Jul 12, 2007
 */
abstract class Demo {
```

```
    UcmdbService stub;
    CmdbContext context;
```

```
    public void initDemo() {
        try {
            setStub(createUcmdbService("admin", "admin"));
            setContext();
        } catch (Exception e) {
            //handle exception
        }
    }
}
```

```
    public UcmdbService getStub() {
        return stub;
    }
}
```

```
public void setStub(UcmdbService stub) {
    this.stub = stub;
}
```

```
public CmdbContext getContext() {
    return context;
}
```

```
public void setContext() {
    CmdbContext context = new CmdbContext();
    context.setCallerApplication("demo");
    this.context = context;
}
```

```
//connection to service - for axis2/jibx client
```

```
private static final String PROTOCOL = "http";
private static final String HOST_NAME = "host_name";
private static final int PORT = 21212;
private static final String FILE = "/axis2/services/UcmdbService";
```

```
protected UcmdbService createUcmdbService
(String username, String password) throws Exception{
    URL url;
    UcmdbServiceStub serviceStub;
```

```
try {
    url = new URL
        (Demo.PROTOCOL, Demo.HOST_NAME,
        Demo.PORT, Demo.FILE);
    serviceStub = new UcmdbServiceStub(url.toString());
    HttpTransportProperties.Authenticator auth =
        new HttpTransportProperties.Authenticator();
    auth.setUsername(username);
    auth.setPassword(password);
    serviceStub._getServiceClient().getOptions().setProperty
        (HTTPConstants.AUTHENTICATE,auth);
```

```
    } catch (AxisFault axisFault) {  
        throw new Exception  
            ("Failed to create SOAP adapter for "  
             + Demo.HOST_NAME , axisFault);
```

```
    } catch (MalformedURLException e) {  
  
        throw new Exception  
            ("Failed to create SOAP adapter for "  
             + Demo.HOST_NAME, e);  
    }  
    return serviceStub;  
}  
}
```

 **Query Example**

```

package com.hp.ucmdb.demo;

import com.hp.ucmdb.generated.params.query.*;
import com.hp.ucmdb.generated.services.UcmdbFaultException;
import com.hp.ucmdb.generated.services.UcmdbService;
import com.hp.ucmdb.generated.types.*;
import com.hp.ucmdb.generated.types.props.*;

import java.rmi.RemoteException;

public class QueryDemo extends Demo{

    UcmdbService stub;
    CmdbContext context;

    public void getClsByTypeDemo() {
        GetClsByType request = new GetClsByType();
        //set cmdbcontext
        CmdbContext cmdbContext = getContext();
        request.setCmdbContext(cmdbContext);
        //set Cls type
        request.setType("anyType");
        //set Cls propeties to be retrieved
        CustomProperties customProperties = new CustomProperties();
        PredefinedProperties predefinedProperties =
            new PredefinedProperties();
        SimplePredefinedProperty simplePredefinedProperty =
            new SimplePredefinedProperty();
        simplePredefinedProperty.setName
            (SimplePredefinedProperty.nameEnum.DERIVED);
        SimplePredefinedPropertyCollection
            simplePredefinedPropertyCollection =
            new SimplePredefinedPropertyCollection();
    }
}

```

```

simplePredefinedPropertyCollection.addSimplePredefinedProperty
    (simplePredefinedProperty);
predefinedProperties.setSimplePredefinedProperties
    (simplePredefinedPropertyCollection);
customProperties.setPredefinedProperties(predefinedProperties);
request.setProperties(customProperties);
try {
    GetCIsByTypeResponse response =
        getStub().getCIsByType(request);
    TopologyMap map =
        getTopologyMapResultFromCIs
            (response.getCIs(), response.getChunkInfo());
} catch (RemoteException e) {
    //handle exception
} catch (UcmdbFaultException e) {
    //handle exception
}
}
}

```

```

public void getCIsByIdDemo() {
    GetCIsById request = new GetCIsById();
    CmdbContext cmdbContext = getContext();
    //set cmdbcontext
    request.setCmdbContext(cmdbContext);
    //set ids
    ID id1 = new ID();
    id1.setBase("cmdbobjectidCIT1");
    ID id2 = new ID();
    id2.setBase("cmdbobjectidCIT2");
    IDs ids = new IDs();
    ids.addID(id1);
    ids.addID(id2);
    request.setIDs(ids);
    //set CIs properties to be retrieved
    TypedPropertiesCollection properties =
        new TypedPropertiesCollection();

```

```

TypedProperties typedProperties1 =
    new TypedProperties();
typedProperties1.setType("CIT1");

```



```

CustomTypedProperties customProperties1 =
    new CustomTypedProperties();
PredefinedTypedProperties predefinedProperties1 =
    new PredefinedTypedProperties();
SimpleTypedPredefinedProperty simplePredefinedProperty1 =
    new SimpleTypedPredefinedProperty();
simplePredefinedProperty1.setName
    (SimpleTypedPredefinedProperty.nameEnum.CONCRETE);
SimpleTypedPredefinedPropertyCollection
    simplePredefinedPropertyCollection1 =
        new SimpleTypedPredefinedPropertyCollection();
simplePredefinedPropertyCollection1
    .addSimpleTypedPredefinedProperty
        (simplePredefinedProperty1);

```

```

predefinedProperties1.
    setSimpleTypedPredefinedProperties
        (simplePredefinedPropertyCollection1);
customProperties1.
    setPredefinedTypedProperties
        (predefinedProperties1);
typedProperties1.setProperties(customProperties1);
properties.addTypedProperties(typedProperties1);

```

```

TypedProperties typedProperties2 =
    new TypedProperties();
typedProperties2.setType("CIT2");
CustomTypedProperties customProperties2 =
    new CustomTypedProperties();
PredefinedTypedProperties predefinedProperties2 =
    new PredefinedTypedProperties();
SimpleTypedPredefinedProperty simplePredefinedProperty2 =
    new SimpleTypedPredefinedProperty();
simplePredefinedProperty2.setName
    (SimpleTypedPredefinedProperty.nameEnum.NAMING);
SimpleTypedPredefinedPropertyCollection
    simplePredefinedPropertyCollection2 =
        new SimpleTypedPredefinedPropertyCollection();

```

```
simplePredefinedPropertyCollection2.  
    addSimpleTypedPredefinedProperty  
        (simplePredefinedProperty2);
```

```
predefinedProperties2.setSimpleTypedPredefinedProperties  
    (simplePredefinedPropertyCollection2);  
customProperties2.setPredefinedTypedProperties  
    (predefinedProperties2);  
typedProperties2.setProperties(customProperties2);  
properties.addTypedProperties(typedProperties2);
```

```
request.setClsTypedProperties(properties);  
try {  
    GetClsByIdResponse response =  
        getStub().getClsById(request);  
    Cls cis = response.getCls();  
} catch (RemoteException e) {  
    //handle exception  
} catch (UcmdbFaultException e) {  
    //handle exception  
}  
  
}
```

```
public void getFilteredClsByTypeDemo() {  
    GetFilteredClsByType request = new GetFilteredClsByType();  
    CmdbContext cmdbContext = getContext();  
    //set cmdbcontext  
    request.setCmdbContext(cmdbContext);  
    //set Cls type  
    request.setType("anyType");  
    //sets Filter conditions  
    Conditions conditions = new Conditions();  
    IntConditions intConditions = new IntConditions();  
    IntCondition intCondition = new IntCondition();  
    IntProp intProp = new IntProp();  
    intProp.setName("int_attr1");
```

```

intProp.setValue(100);
intCondition.setCondition(intProp);
intCondition.setIntOperator
    (IntCondition.intOperatorEnum.Greater);
intConditions.addIntCondition(intCondition);

```

```

conditions.setIntConditions(intConditions);
request.setConditions(conditions);
//set logical operator for conditions
request.setConditionsLogicalOperator
    (GetFilteredCIsByType.conditionsLogicalOperatorEnum.AND);
//set CIs properties to be retrieved
CustomProperties customProperties =
    new CustomProperties();
PredefinedProperties predefinedProperties =
    new PredefinedProperties();
SimplePredefinedProperty simplePredefinedProperty =
    new SimplePredefinedProperty();
simplePredefinedProperty.setName
    (SimplePredefinedProperty.nameEnum.NAMING);

```

```

SimplePredefinedPropertyCollection
    simplePredefinedPropertyCollection =
        new SimplePredefinedPropertyCollection();
simplePredefinedPropertyCollection.
    addSimplePredefinedProperty
        (simplePredefinedProperty);
predefinedProperties.setSimplePredefinedProperties
    (simplePredefinedPropertyCollection);
customProperties.setPredefinedProperties
    (predefinedProperties);

```

```

request.setProperties(customProperties);
try {
    GetFilteredCIsByTypeResponse response =
        getStub().getFilteredCIsByType(request);
    TopologyMap map =
        getTopologyMapResultFromCIs
            (response.getCIs(), response.getChunkInfo());
}

```

```
    } catch (RemoteException e) {  
        //handle exception  
    } catch (UcmdbFaultException e) {  
        //handle exception  
    }  
  
}
```

```
public void executeTopologyQueryByNameDemo() {  
    ExecuteTopologyQueryByName request = new  
ExecuteTopologyQueryByName();  
    CmdbContext cmdbContext = getContext();  
    //set cmdbcontext  
    request.setCmdbContext(cmdbContext);  
    //set query name  
    request.setQueryName("queryName");
```

```
    try {  
        ExecuteTopologyQueryByNameResponse response =  
            getStub().executeTopologyQueryByName(request);  
        TopologyMap map =  
            getTopologyMapResult  
                (response.getTopologyMap(), response.getChunkInfo());  
    } catch (RemoteException e) {  
        //handle exception  
    } catch (UcmdbFaultException e) {  
        //handle exception  
    }  
  
}
```

```

// assume the follow query was defined at UCMDB
// Query Name: exampleQuery
// Query sketch:
//           Host
//           / \
//           ip Disk
// Query Parameters:
//   Host-
//     host_os (like)
//   Disk-
//     disk_failures (equal)

```

```

public void executeTopologyQueryByNameWithParametersDemo() {
    ExecuteTopologyQueryByNameWithParameters request =
        new ExecuteTopologyQueryByNameWithParameters();
    CmdbContext cmdbContext = getContext();
    //set cmdbcontext
    request.setCmdbContext(cmdbContext);
    //set query name
    request.setQueryName("queryName");
    //set parameters
    ParameterizedNode hostParametrizedNode =
        new ParameterizedNode();
    hostParametrizedNode.setNodeLabel("Host");
    CIProperties parameters = new CIProperties();
    StrProps strProps = new StrProps();
    StrProp strProp = new StrProp();
    strProp.setName("host_os");
    strProp.setValue("%2000%");
    strProps.addStrProp(strProp);
    parameters.setStrProps(strProps);
    hostParametrizedNode.setParameters(parameters);
    request.addParameterizedNodes(hostParametrizedNode);
    ParameterizedNode diskParametrizedNode =
        new ParameterizedNode();

```

```

    diskParametrizedNode.setNodeLabel("Disk");
    CIProperties parameters1 = new CIProperties();
    IntProps intProps = new IntProps();

```

```

IntProp intProp = new IntProp();
intProp.setName("disk_failures");
intProp.setValue(30);
intProps.addIntProp(intProp);
parameters1.setIntProps(intProps);
diskParametrizedNode.setParameters(parameters1);

```

```

request.addParameterizedNodes(diskParametrizedNode);
try {
    ExecuteTopologyQueryByNameWithParametersResponse
        response =
            getStub().executeTopologyQueryByNameWithParameters
                (request);
    TopologyMap map =
        getTopologyMapResult
            (response.getTopologyMap(), response.getChunkInfo());
} catch (RemoteException e) {
    //handle exception
} catch (UcmdbFaultException e) {
    //handle exception
}
}

```

```

/ // assume the follow query was defined at UCMDB
// Query Name: exampleQuery
// Query sketch:
//           Host
//           / \
//           ip Disk
// Query Parameters:
//   Host-
//     host_os (like)
//   Disk-
//     disk_failures (equal)

```

```

public void executeTopologyQueryWithParametersDemo() {
    ExecuteTopologyQueryWithParameters request =
        new ExecuteTopologyQueryWithParameters();
    CmdbContext cmdbContext = getContext();
    //set cmdbcontext
    request.setCmdbContext(cmdbContext);
    //set query definition
    String queryXml = "<xml that represents the query above>";
    request.setQueryXml(queryXml);
    //set parameters
    ParameterizedNode hostParametrizedNode =
        new ParameterizedNode();

```

```

    hostParametrizedNode.setNodeLabel("Host");
    CIProperties parameters = new CIProperties();
    StrProps strProps = new StrProps();
    StrProp strProp = new StrProp();
    strProp.setName("host_os");
    strProp.setValue("%2000%");
    strProps.addStrProp(strProp);
    parameters.setStrProps(strProps);
    hostParametrizedNode.setParameters(parameters);
    request.addParameterizedNodes(hostParametrizedNode);
    ParameterizedNode diskParametrizedNode =
        new ParameterizedNode();
    diskParametrizedNode.setNodeLabel("Disk");
    CIProperties parameters1 = new CIProperties();
    IntProps intProps = new IntProps();
    IntProp intProp = new IntProp();
    intProp.setName("disk_failures");
    intProp.setValue(30);
    intProps.addIntProp(intProp);
    parameters1.setIntProps(intProps);
    diskParametrizedNode.setParameters(parameters1);
    request.addParameterizedNodes(diskParametrizedNode);

```

```

try {
    ExecuteTopologyQueryWithParametersResponse
    response = getStub().executeTopologyQueryWithParameters
        (request);
    TopologyMap map =
        getTopologyMapResult
            (response.getTopologyMap(), response.getChunkInfo());

```

```

    } catch (RemoteException e) {
        //handle exception
    } catch (UcmdbFaultException e) {
        //handle exception
    }
}

```

```

public void getCI NeighboursDemo() {
    GetCI Neighbours request = new GetCI Neighbours();
    //set cmdbcontext
    CmdbContext cmdbContext = getContext();
    request.setCmdbContext(cmdbContext);
    // set CI id
    ID id = new ID();
    id.setBase("cmdbobjectidCIT1");
    request.setID(id);
    //set neighbour type
    request.setNeighbourType("neighbourType");
    //set Neighbours CIs properties to be retrieved
    TypedPropertiesCollection properties =
        new TypedPropertiesCollection();
    TypedProperties typedProperties1 = new TypedProperties();
    typedProperties1.setType("neighbourType");
    CustomTypedProperties customProperties1 =
        new CustomTypedProperties();
    PredefinedTypedProperties predefinedProperties1 =
        new PredefinedTypedProperties();

```



```

QualifierProperties qualifierProperties =
    new QualifierProperties();
qualifierProperties.addQualifierName("ID_ATTRIBUTE");
predefinedProperties1.setQualifierProperties(qualifierProperties);
customProperties1.setPredefinedTypedProperties
    (predefinedProperties1);
typedProperties1.setProperties(customProperties1);
properties.addTypedProperties(typedProperties1);
request.setCIProperties(properties);

```

```

TypedPropertiesCollection relationsProperties =
    new TypedPropertiesCollection();
TypedProperties typedProperties2 = new TypedProperties();
typedProperties2.setType("relationType");
CustomTypedProperties customProperties2 =
    new CustomTypedProperties();

```

```

PredefinedTypedProperties predefinedProperties2 =
    new PredefinedTypedProperties();
SimpleTypedPredefinedProperty simplePredefinedProperty2 =
    new SimpleTypedPredefinedProperty();
simplePredefinedProperty2.setName

```

```

    (SimpleTypedPredefinedProperty.nameEnum.CONCRETE);
SimpleTypedPredefinedPropertyCollection
    simplePredefinedPropertyCollection2 =
        new SimpleTypedPredefinedPropertyCollection();
simplePredefinedPropertyCollection2.
    addSimpleTypedPredefinedProperty
        (simplePredefinedProperty2);
predefinedProperties2.
    setSimpleTypedPredefinedProperties
        (simplePredefinedPropertyCollection2);
customProperties2.setPredefinedTypedProperties
    (predefinedProperties2);
typedProperties2.setProperties(customProperties2);
relationsProperties.addTypedProperties(typedProperties2);
request.setRelationProperties(relationsProperties);

```

```
try {
    GetCINeighboursResponse response =
        getStub().getCINeighbours(request);
    Topology topology = response.getTopology();
} catch (RemoteException e) {
    //handle exception
} catch (UcmdbFaultException e) {
    //handle exception
}
}
```

```
//get Topology Map for chunked/non-chunked result
```

```
private TopologyMap getTopologyMapResult(TopologyMap topologyMap, ChunkInfo
chunkInfo) {
    if(chunkInfo.getNumberOfChunks() == 0) {
        return topologyMap;
    } else {
```

```
        topologyMap = new TopologyMap();
        for(int i=1 ; i <= chunkInfo.getNumberOfChunks() ; i++) {
            ChunkRequest chunkRequest = new ChunkRequest();
            chunkRequest.setChunkInfo(chunkInfo);
            chunkRequest.setChunkNumber(i);
            PullTopologyMapChunks req =
                new PullTopologyMapChunks();
            req.setChunkRequest(chunkRequest);
            req.setCmdbContext(getContext());
            PullTopologyMapChunksResponse res = null;
```

```

        try {
            res = getStub().pullTopologyMapChunks(req);
            TopologyMap map = res.getTopologyMap();
            topologyMap = mergeMaps(topologyMap, map);
        } catch (RemoteException e) {
            //handle exception
        } catch (UcmdbFaultException e) {
            //handle exception
        }
    }
}
return topologyMap;
}

```

```

private TopologyMap getTopologyMapResultFromCIs(CIs cis, ChunkInfo chunkInfo)
{
    TopologyMap topologyMap = new TopologyMap();
    if(chunkInfo.getNumberOfChunks() == 0) {
        CINode ciNode = new CINode();
        ciNode.setLabel("");
        ciNode.setCIs(cis);
        CINodes ciNodes = new CINodes();
        ciNodes.addCINode(ciNode);
        topologyMap.setCINodes(ciNodes);
    } else {

```

```

        for(int i=1 ; i <= chunkInfo.getNumberOfChunks() ; i++) {
            ChunkRequest chunkRequest =
                new ChunkRequest();
            chunkRequest.setChunkInfo(chunkInfo);
            chunkRequest.setChunkNumber(i);
            PullTopologyMapChunks req =
                new PullTopologyMapChunks();
            req.setChunkRequest(chunkRequest);
            req.setCmdbContext(getContext());
            PullTopologyMapChunksResponse res = null;

```

```

    try {
        res = getStub().pullTopologyMapChunks(req);
    } catch (RemoteException e) {
        //handle exception
    } catch (UcmdbFaultException e) {
        //handle exception
    }
    TopologyMap map = res.getTopologyMap();
    topologyMap = mergeMaps(topologyMap, map);
}

```

```

//release chunks
ReleaseChunks req = new ReleaseChunks();
req.setChunksKey(chunkInfo.getChunksKey());
req.setCmdbContext(getContext());

```

```

    try {
        getStub().releaseChunks(req);
    } catch (RemoteException e) {
        //handle exception
    } catch (UcmdbFaultException e) {
        //handle exception
    }
}
return topologyMap;
}

```

```

//=====================================================
/* WARNING merge will be correct only if a each node is given
   a unique name. This applies to both CI and Relation nodes .*/
//=====================================================
private TopologyMap mergeMaps(TopologyMap topologyMap, TopologyMap
newMap) {
    for(int i=0 ; i < newMap.getCINodes().sizeCINodeList() ; i++ ) {
        CINode ciNode = newMap.getCINodes().getCINode(i);
        boolean alreadyExist = false;
        if(topologyMap.getCINodes() == null) {
            topologyMap.setCINodes(new CINodes());
        }
    }
}

```

```

for(int j=0 ; j < topologyMap.getCINodes().sizeCINodeList() ; j++) {
    CINode ciNode2 = topologyMap.getCINodes().getCINode(j);
    if(ciNode2.getLabel().equals(ciNode.getLabel())){

```

```

        Cls cisTOAdd = ciNode.getCIs();
        Cls cis =
            mergeCIsGroups
            (topologyMap.getCINodes().getCINode(j).getCIs(),
             cisTOAdd);
        topologyMap.getCINodes().getCINode(j).setCIs(cis);
        alreadyExist = true;
    }
}
if(!alreadyExist) {
    topologyMap.getCINodes().addCINode(ciNode);
}
}

```

```

for(int i=0 ; i < newMap.getRelationNodes().sizeRelationNodeList() ; i++ ) {
    RelationNode relationNode =
        newMap.getRelationNodes().getRelationNode(i);
    boolean alreadyExist = false;
    if(topologyMap.getRelationNodes() == null) {
        topologyMap.setRelationNodes(new RelationNodes());
    }
}

```

```

for(int j=0 ;
    j < topologyMap.getRelationNodes().sizeRelationNodeList() ;
    j++){
    RelationNode relationNode2 =
        topologyMap.getRelationNodes().getRelationNode(j);
    if(relationNode2.getLabel().equals(relationNode.getLabel())){
        Relations relationsTOAdd = relationNode.getRelations();
        Relations relations =
            mergeRelationsGroups
            (topologyMap.getRelationNodes().
                getRelationNode(j).getRelations(),
                relationsTOAdd);
        topologyMap.getRelationNodes().
            getRelationNode(j).setRelations(relations);
        alreadyExist = true;
    }
}

```

```

    if(!alreadyExist) {
        topologyMap.getRelationNodes().addRelationNode(relationNode);
    }
}

return topologyMap;
}

```

```

private Relations mergeRelationsGroups(Relations relations1, Relations relations2)
{
    for(int i=0 ; i < relations2.sizeRelationList() ; i++) {
        relations1.addRelation(relations2.getRelation(i));
    }
    return relations2;
}

```

```
private Cls mergeClsGroups(Cls cis1, Cls cis2) {  
    for(int i=0 ; i < cis2.sizeCIList() ; i++) {  
        cis1.addCI(cis2.getCI(i));  
    }  
    return cis1;  
}  
  
}
```

 **Update Example**

```
package com.hp.ucmdb.demo;

import com.hp.ucmdb.generated.params.update.AddCIsAndRelations;
import com.hp.ucmdb.generated.params.update.AddCIsAndRelationsResponse;
import com.hp.ucmdb.generated.params.update.UpdateCIsAndRelations;
import com.hp.ucmdb.generated.params.update.DeleteCIsAndRelations;
import com.hp.ucmdb.generated.services.UcmdbFaultException;
import com.hp.ucmdb.generated.types.*;
import com.hp.ucmdb.generated.types.update.CIsAndRelationsUpdates;
import com.hp.ucmdb.generated.types.update.ClientIDToCmdbID;

import java.rmi.RemoteException;

public class UpdateDemo extends Demo{
```

```
    public void getAddCIsAndRelationsDemo() {
        AddCIsAndRelations request = new AddCIsAndRelations();
        request.setCmdbContext(getContext());
        request.setUpdateExisting(true);
        CIsAndRelationsUpdates updates = new CIsAndRelationsUpdates();
        CIs cis = new CIs();
        CI ci = new CI();
        ID id = new ID();
        id.setBase("temp1");
        id.setTemp(true);
```

```
        ci.setID(id);
        ci.setType("host");
```

```
        CIProperties props = new CIProperties();
        StrProps strProps = new StrProps();
        StrProp strProp = new StrProp();
        strProp.setName("host_key");
        String value = "blabla";
        strProp.setValue(value);
```



```

strProps.addStrProp(strProp);
props.setStrProps(strProps);
ci.setProps(props);
cis.addCI(ci);
updates.setCIsForUpdate(cis);
request.setCIsAndRelationsUpdates(updates);

```

```

try {
    AddCIsAndRelationsResponse response =
        getStub().addCIsAndRelations(request);
    for(int i = 0 ; i < response.sizeCreatedIDsMapList() ; i++) {
        ClientIDToCmdbID idsMap = response.getCreatedIDsMap(i);
        //do something
    }
} catch (RemoteException e) {
    //handle exception
} catch (UcmdbFaultException e) {
    //handle exception
}
}

```

```

public void getUpdateCIsAndRelationsDemo() {
    UpdateCIsAndRelations request = new UpdateCIsAndRelations();
    request.setCmdbContext(getContext());

```

```

CIsAndRelationsUpdates updates =
    new CIsAndRelationsUpdates();
CIs cis = new CIs();
CI ci = new CI();
ID id = new ID();

```

```

id.setBase("temp1");
id.setTemp(true);
ci.setID(id);
ci.setType("host");
CIProperties props = new CIProperties();
StrProps strProps = new StrProps();

```

```
StrProp hostKeyProp = new StrProp();
hostKeyProp.setName("host_key");
String hostKeyValue = "blabla";
hostKeyProp.setValue(hostKeyValue);
strProps.addStrProp(hostKeyProp);
```

```
StrProp hostOSProp = new StrProp();
hostOSProp.setName("host_os");
String hostOSValue = "winXP";
hostOSProp.setValue(hostOSValue);
strProps.addStrProp(hostOSProp);
```

```
StrProp hostDNSProp = new StrProp();
hostDNSProp.setName("host_dnsname");
String hostDNSValue = "dnsname";
hostDNSProp.setValue(hostDNSValue);
strProps.addStrProp(hostDNSProp);
```

```
props.setStrProps(strProps);
ci.setProps(props);
cis.addCI(ci);
updates.setCIsForUpdate(cis);
request.setCIsAndRelationsUpdates(updates);
```

```
try {
    getStub().updateCIsAndRelations(request);
} catch (RemoteException e) {
    //handle exception
} catch (UcmdbFaultException e) {
    //handle exception
}
}
```

```

public void getDeleteCIsAndRelationsDemo() {
    DeleteCIsAndRelations request =
        new DeleteCIsAndRelations();
    request.setCmdbContext(getContext());
    CIsAndRelationsUpdates updates =
        new CIsAndRelationsUpdates();
    CIs cis = new CIs();
    CI ci = new CI();
    ID id = new ID();
    id.setBase("stam");
    id.setTemp(true);
    ci.setID(id);
    ci.setType("host");

```

```

    CIProperties props = new CIProperties();
    StrProps strProps = new StrProps();
    StrProp strProp1 = new StrProp();
    strProp1.setName("host_key");
    String value1 = "for_delete";
    strProp1.setValue(value1);
    strProps.addStrProp(strProp1);
    props.setStrProps(strProps);
    ci.setProps(props);
    cis.addCI(ci);
    updates.setCIsForUpdate(cis);
    request.setCIsAndRelationsUpdates(updates);

```

```

        try {
            getStub().deleteCIsAndRelations(request);
        } catch (RemoteException e) {
            //handle exception
        } catch (UcmdbFaultException e) {
            //handle exception
        }
    }
}
}

```

 **Class Model Example**

```
package com.hp.ucmdb.demo;

import com.hp.ucmdb.generated.params.classmodel.*;
import com.hp.ucmdb.generated.services.UcmdbFaultException;
import com.hp.ucmdb.generated.types.classmodel.UcmdbClassModelHierarchy;
import com.hp.ucmdb.generated.types.classmodel.UcmdbClass;

import java.rmi.RemoteException;

public class ClassmodelDemo extends Demo{
```

```
    public void getClassAncestorsDemo() {
        GetClassAncestors request =
            new GetClassAncestors();
        request.setCmdbContext(getContext());
        request.setClassName("className");
```

```
        try {
            GetClassAncestorsResponse response =
                getStub().getClassAncestors(request);
            UcmdbClassModelHierarchy hierarchy =
                response.getClassHierarchy();
        } catch (RemoteException e) {
            //handle exception
        } catch (UcmdbFaultException e) {
            //handle exception
        }
    }
}
```

```

public void getAllClassesHierarchyDemo() {
    GetAllClassesHierarchy request =
        new GetAllClassesHierarchy();
    request.setCmdbContext(getContext());
    try {
        GetAllClassesHierarchyResponse response =
            getStub().getAllClassesHierarchy(request);
        UcmdbClassModelHierarchy hierarchy =
            response.getClassesHierarchy();
    } catch (RemoteException e) {
        //handle exception
    } catch (UcmdbFaultException e) {
        //handle exception
    }
}
}

```

```

public void getCmdbClassDefinitionDemo() {
    GetCmdbClassDefinition request =
        new GetCmdbClassDefinition();
    request.setCmdbContext(getContext());
    request.setClassName("className");

```

```

    try {
        GetCmdbClassDefinitionResponse response =
            getStub().getCmdbClassDefinition(request);
        UcmdbClass ucmdbClass = response.getUcmdbClass();
    } catch (RemoteException e) {
        //handle exception
    } catch (UcmdbFaultException e) {
        //handle exception
    }
}
}
}

```

 **Impact Analysis Example**

```

package com.hp.ucmdb.demo;

import com.hp.ucmdb.generated.params.impact.*;
import com.hp.ucmdb.generated.services.UcmdbFaultException;
import com.hp.ucmdb.generated.types.*;
import com.hp.ucmdb.generated.types.impact.*;

import java.rmi.RemoteException;

/**
 * Date: Jul 17, 2007
 */
public class ImpactDemo extends Demo{

//Impact Rule Name : impactExample
//Impact Query:
//      Network
//      |
//      Host
//      |
//      IP
//Impact Action: network affect on ip ;severity 100% ; category: change
//
public void calculateImpactAndGetImpactPathDemo() {
    CalculateImpact request = new CalculateImpact();
    request.setCmdbContext(getContext());
    //set root cause ids
    IDs ids = new IDs();
    ID id = new ID();
    id.setBase("rootCauseCmdbID");
    ids.addID(id);
}
}

```

```

request.setIDs(ids);
//set impact category
request.setImpactCategory("change");
//set rule Names
ImpactRuleNames impactRuleNames = new ImpactRuleNames();
ImpactRuleName impactRuleName = new ImpactRuleName();
impactRuleName.setBase("impactExample");
impactRuleNames.addImpactRuleName(impactRuleName);
request.setImpactRuleNames(impactRuleNames);
//set severity
request.setSeverity(100);
CalculateImpactResponse response =
    new CalculateImpactResponse();

```

```

request.setIDs(ids);
//set impact category
request.setImpactCategory("change");
//set rule Names
ImpactRuleNames impactRuleNames = new ImpactRuleNames();
ImpactRuleName impactRuleName = new ImpactRuleName();
impactRuleName.setBase("impactExample");
impactRuleNames.addImpactRuleName(impactRuleName);
request.setImpactRuleNames(impactRuleNames);
//set severity
request.setSeverity(100);
CalculateImpactResponse response =
    new CalculateImpactResponse();

```

```

try {
    response = getStub().calculateImpact(request);
} catch (RemoteException e) {
    //handle exception
}

```

```

    } catch (UcmdbFaultException e) {
        //handle exception
    }
    Identifier identifier= response.getIdentifier();
    Topology topology = response.getImpactTopology();
    Relation relation = topology.getRelations().getRelation(0);
    GetImpactPath request2 = new GetImpactPath();
    //set cmdb context
    request2.setCmdbContext(getContext());
    //set impact identifier
    request2.setIdentifier(identifier);
    //set shallowRelation
    ShallowRelation shallowRelation = new ShallowRelation();
    shallowRelation.setID(relation.getID());
    shallowRelation.setEnd1ID(relation.getEnd1ID());
    shallowRelation.setEnd2ID(relation.getEnd2ID());
    shallowRelation.setType(relation.getType());
    request2.setRelation(shallowRelation);

```

```

try {
    GetImpactPathResponse response2 =
        getStub().getImpactPath(request2);
    ImpactTopology impactTopology =
        response2.getImpactPathTopology();
} catch (RemoteException e) {
    //To change body of catch statement
    // use File | Settings | File Templates.
    e.printStackTrace();
} catch (UcmdbFaultException e) {
    //To change body of catch statement
    // use File | Settings | File Templates.
    e.printStackTrace();
}
}

```



```

public void getImpactRulesByGroupName() {
    GetImpactRulesByGroupName request =
        new GetImpactRulesByGroupName();
    //set cmdb context
    request.setCmdbContext(getContext());
    //set group names list
    request.addRuleGroupNameFilter("groupName1");
    request.addRuleGroupNameFilter("groupName2");

```

```

    try {
        GetImpactRulesByGroupNameResponse response =
            getStub().getImpactRulesByGroupName(request);
        ImpactRules impactRules = response.getImpactRules();
    } catch (RemoteException e) {
        //handle exception
    } catch (UcmdbFaultException e) {
        //handle exception
    }
}

```

```

public void getImpactRulesByNamePrefix() {
    GetImpactRulesByNamePrefix request =
        new GetImpactRulesByNamePrefix();
    //set cmdb context
    request.setCmdbContext(getContext());
    //set prefixes list
    request.addRuleNamePrefixFilter("prefix1");

```

```
try {
    GetImpactRulesByNamePrefixResponse response =
        getStub().getImpactRulesByNamePrefix(request);
    ImpactRules impactRules = response.getImpactRules();
} catch (RemoteException e) {
    //handle exception
} catch (UcmdbFaultException e) {
    //handle exception
}
}
}
```

 **Adding Credentials Example**

```
import java.net.URL;

import org.apache.axis2.transport.http.HTTPConstants;
import org.apache.axis2.transport.http.HttpTransportProperties;

import com.hp.ucmdb.generated.params.discovery.*;
import com.hp.ucmdb.generated.services.DiscoveryService;
import com.hp.ucmdb.generated.services.DiscoveryServiceStub;
import com.hp.ucmdb.generated.types.BytesProp;
import com.hp.ucmdb.generated.types.BytesProps;
import com.hp.ucmdb.generated.types.CIProperties;
import com.hp.ucmdb.generated.types.CmdbContext;
import com.hp.ucmdb.generated.types.StrList;
import com.hp.ucmdb.generated.types.StrProp;
import com.hp.ucmdb.generated.types.StrProps;

public class test {
    static final String HOST_NAME = "hostname";
    static final int PORT = 808021212;

    private static final String PROTOCOL = "http";
    private static final String FILE = "/axis2/services/DiscoveryService";

    private static final String PASSWORD = "admin";
    private static final String USERNAME = "admin";

    private static CmdbContext cmdbContext = new CmdbContext("ws tests");
```

```
public static void main(String[] args) throws Exception {
    // Get the stub object
    DiscoveryService discoveryService = getDiscoveryService();

    // Activate Job
    discoveryService.activateJob(new ActivateJobRequest("Range IPs by ICMP",
    cmdbContext));

    // Get domain & probes info
    getProbesInfo(discoveryService);

    // Add credentilas entry for ntcmd protocol
    addNTCMDCredentialsEntry();
}
```

```

public static void addNTCMDCredentialsEntry() throws Exception {
    DiscoveryService discoveryService = getDiscoveryService();

    // Get domain name
    StrList domains =
        discoveryService.getDomainsNames(new
GetDomainsNamesRequest(cmdbContext)).getDomainNames();
    if (domains.sizeStrValueList() == 0) {
        System.out.println("No domains were found, can't create credentials");
        return;
    }
    String domainName = domains.getStrValue(0);

    // Create properties with one byte param
    CIProperties newCredsProperties = new CIProperties();

    // Add password property - this is of type bytes
    newCredsProperties.setBytesProps(new BytesProps());
    setPasswordProperty(newCredsProperties);

    // Add user & domain properties - these are of type string
    newCredsProperties.setStrProps(new StrProps());
    setStringProperties("protocol_username", "test user", newCredsProperties);
    setStringProperties("ntadminprotocol_ntdomain", "test domain",
newCredsProperties);

    // Add new credentials entry
    discoveryService.addCredentialsEntry(new
AddCredentialsEntryRequest(domainName, "ntadminprotocol", newCredsProperties,
cmdbContext));

    System.out.println("new credentials created for domain: " + domainName + " in
ntcmd protocol");
}

```

```

private static void setPasswordProperty(CIProperties newCredsProperties) {
    BytesProp bProp = new BytesProp();
    bProp.setName("protocol_password");
    bProp.setValue(new byte[] {101, 103, 102, 104});
    newCredsProperties.getBytesProps().addBytesProp(bProp);
}

```

```

private static void setStringProperties(String propertyName, String value,
CIProperties newCredsProperties) {
    StrProp strProp = new StrProp();
    strProp.setName(propertyName);
    strProp.setValue(value);
    newCredsProperties.getStrProps().addStrProp(strProp);
}

```

```

private static void getProbesInfo(DiscoveryService discoveryService) throws
Exception {
    GetDomainsNamesResponse result =
discoveryService.getDomainsNames(new GetDomainsNamesRequest(cmdbContext
));

    // Go over all the domains
    if (result.getDomainNames().sizeStrValueList() > 0) {
        String domainName = result.getDomainNames().getStrValue(0);
        GetProbesNamesResponse probesResult =
            discoveryService.getProbesNames(new
GetProbesNamesRequest(domainName, cmdbContext));

        // Go over all the probes
        for (int i=0; i<probesResult.getProbesNames().sizeStrValueList(); i++) {
            String probeName = probesResult.getProbesNames().getStrValue(i);

            // Check if connected
            IsProbeConnectedResponse connectedRequest =
                discoveryService.isProbeConnected(new
IsProbeConnectedRequest(domainName, probeName, cmdbContext));
            Boolean isConnected = connectedRequest.getIsConnected();

            // Do something ...
            System.out.println("probe " + probeName + " isconnect=" +
isConnected);
        }
    }
}

```

```
private static DiscoveryService getDiscoveryService() throws Exception {
    DiscoveryService discoveryService = null;
    try {

        // Create service
        URL url = new URL(PROTOCOL,HOST_NAME,PORT, FILE);
        DiscoveryServiceStub serviceStub = new
DiscoveryServiceStub(url.toString());

        // Authenticate info
        HttpTransportProperties.Authenticator auth = new
HttpTransportProperties.Authenticator();
        auth.setUsername(USERNAME);
        auth.setPassword(PASSWORD);

serviceStub._getServiceClient().getOptions().setProperty(HTTPConstants.AUTHENTIC
ATE,auth);

        discoveryService = serviceStub;
    } catch (Exception e) {
        throw new Exception("cannot create a connection to service ", e);
    }

    return discoveryService;
}
} // End class
```


11

Data Flow Management API

This chapter includes:

Concepts

- ▶ Data Flow Management API Overview on page 386
- ▶ Conventions on page 386
- ▶ Discovery and Dependency Mapping Web Service on page 387

Tasks

- ▶ Call the Web Service on page 388

Reference

- ▶ Discovery and Dependency Mapping Methods on page 389
- ▶ Code Sample on page 402

Concepts

Data Flow Management API Overview

This chapter explains how third-party or custom tools can use the HP Discovery and Dependency Mapping Web Service to manage Data Flow Management.

For full documentation on the available operations, see *HP Discovery and Dependency Mapping Schema Reference*. These files are located in the following folder: \\<Business Service Management root directory>\AppServer\webapps\site.war\amdocs\eng\doc_lib\DevRef_guide\DDM_Schema\webframe.html

Conventions

This chapter uses the following conventions:

- ▶ This style Element indicates that an item is an entity in the database or an element defined in the schema, including structures passed to or returned by methods. Plain text indicates that the item is being discussed in a general context.
- ▶ DDM elements and method arguments are spelled in the case in which they are specified in the schema. This usually means that a class name or generic reference to an instance of the class is capitalized. An element or argument to a method is not capitalized. For example, a credential is an element of type Credential passed to a method.

Discovery and Dependency Mapping Web Service

The HP Discovery and Dependency Mapping Web Service is an API used to integrate applications with HP Business Service Management. The API provides methods to:

- **Manage credentials.** View, add, update, and remove.
- **Manage jobs.** View status, activate, and deactivate.
- **Manage probe ranges.** View, add, and update.
- **Manage triggers.** Add or remove a trigger CI, and add, remove, or disable a trigger TQL.
- **View general data.** Data on domains and probes.

Users of the HP Discovery and Dependency Mapping Web Service should be familiar with:

- The SOAP specification
- An object-oriented programming language such as C++, C# or Java
- HP Business Service Management
- Discovery and Dependency Mapping

Permissions

The administrator provides login credentials for connecting with the Web service. The permission levels are View, Update, and Execute. To view the permissions required for each operation, see each operation's request documentation in the *HP Discovery and Dependency Mapping Schema Reference*.

Tasks

Call the Web Service

The HP Discovery and Dependency Mapping Web Service enables calling server-side methods using standard SOAP programming techniques. If the statement cannot be parsed or if there is a problem invoking the method, the API methods throw a `SoapFault` exception. When a `SoapFault` exception is thrown, the service populates one or more of the error message, error code, and exception message fields. If there is no error, the results of the invocation are returned.

To call the service, use:

Protocol: http or https (depending on server configuration)
URL: <BSM Data Processing Server>:21212 /axis2/services/
DiscoveryService
Default password: "admin"
Default username: "admin"

SOAP programmers can access the WSDL at:

`axis2/services/DiscoveryService?wsdl`

Reference

Discovery and Dependency Mapping Methods

This section contains a list of the Web service operations and a brief summary of their use. For full documentation of the request and response for each operation, see *HP Discovery and Dependency Mapping Schema Reference*.

This section includes the following topics:

- ▶ “Data Structures” on page 389
- ▶ “Managing Discovery Job Methods” on page 390
- ▶ “Managing Trigger Methods” on page 392
- ▶ “Domain and Probe Data Methods” on page 394
- ▶ “Credentials Data Methods” on page 397
- ▶ “Data Refresh Methods” on page 399

Data Structures

These are some of the data structures used in the DDM Web Service API.

CIProperties

CIProperties is a collection of collections. Each collection contains properties of a different data type. For example, there can be a `dateProps` collection, a `strListProps` collection, an `xmlProps` collection, and so on.

Each type collection contains individual properties of the given type. The names of these properties elements is the same as the container, but in singular. For example, `dateProps` contains `dateProp` elements. Each property is a name-value pair.

See CIProperties in the *HP Discovery and Dependency Mapping Schema Reference*.

IPList

A list of IP elements, each of which contains an IPv4 Address.

See IPList in the *HP Discovery and Dependency Mapping Schema Reference*.

IPRange

An IPRange has two elements, the **Start** and the **End** elements. Each contains an **Address** element which is an IPv4 Address.

See IPLRange in the *HP Discovery and Dependency Mapping Schema Reference*.

Scope

Two IPRanges. **Exclude** is a collection of IPRanges to exclude from the job. **Include** is a collection of IPRanges to include in the job.

See Scope in the *HP Discovery and Dependency Mapping Schema Reference*.

Managing Discovery Job Methods

activateJob

Activates the specified job.

See example activateJob

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.

deactivateJob

Deactivates the specified job.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.

dispatchAdHocJob

Dispatches a job on the probe ad-hoc. The job must be active and contain the specified trigger CI.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.
CIID	The ID of the trigger CI.
ProbeName	The name of the probe.
Timeout	In milliseconds

getDiscoveryJobsNames

Returns the list of job names.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.

Output

Parameter	Comment
strList	The list of job names.

isJobActive

Checks whether the job is active.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job to check.

Output

Parameter	Comment
JobState	True if the job is active.

Managing Trigger Methods**addTriggerCI**

Adds a new trigger CI to the specified job.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.
CIID	The ID of the trigger CI.

addTriggerTQL

Adds a new trigger TQL to the specified job.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.
TqlName	The name of the TQL to add.

disableTriggerTQL

Prevents the TQL from triggering the job, but does not permanently remove it from the list of queries that trigger the job.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.

removeTriggerCI

Removes the specified CI from the list of CIs that trigger the job.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The job name.
CIID	The ID of the trigger CI.

removeTriggerTQL

Removes the specified TQL from the list of queries that trigger the job.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	Collection of job names to check.
CIID	The ID of the TQL to remove.

setTriggerTQLProbesLimit

Restrict the probes in which the TQL is active in the job to the specified list.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
JobName	The name of the job.
tqlName	The TQL name.
probesLimit	The list of probes for which the TQL is active.

Domain and Probe Data Methods**getDomainType**

Returns the domain type.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The name of the domain.

Output

Parameter	Comment
domainType	The domain type.

getDomainsNames

Returns the names of the current domains.

See example getDomainsNames

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.

Output

Parameter	Comment
domainNames	The list of domain names.

getProbeIPs

Returns the IP addresses of the specified probe.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to check.
probeName	The name of the probe used on that domain.

Output

Parameter	Comment
probeIPs	The IList of the addresses in the probe.

getProbesNames

Returns the names of the probes in the specified domain.

See example `getProbesNames`

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to check.

Output

Parameter	Comment
probesName	The list of probes on the domain.

getProbeScope

Returns the scope definition of the specified probe.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to check.
probeName	The name of the probe.

Output

Parameter	Comment
probeScope	The Scope of the probe.

isProbeConnected

Checks whether the specified probe is connected.

See example isProbeConnected

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to check.
probeName	The probe to check

Output

Parameter	Comment
isConnected	True if the probe is connected.

updateProbeScope

Sets the scope of the specified probe, overriding the existing scope.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain.
probeName	The probe to update.
newScope	The Scope to set for the probe.

Credentials Data Methods**addCredentialsEntry**

Adds a credentials entry to the specified protocol for the specified domain.

See example addCredentialsEntry

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to update.
protocolName	The name of the protocol.
credentialsEntryParameters	The CIProperties collection of the new credentials.

Output

Parameter	Comment
credentialsEntryID	The CI ID of the new credential entry.

getCredentialsEntriesIDs

Returns the IDs of the credentials defined for the specified protocol.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to get the credentials for.
protocolName	The name of a protocol used on that domain.

Output

Parameter	Comment
credentialsEntryIDs	The list of credential IDs for the protocol on the domain.

getCredentialsEntry

Returns the credentials defined for the specified protocol. Encrypted attributes are returned empty.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to get the credentials for.
protocolName	A protocol used on that domain.
credentialsEntryID	The credential ID to get.

Output

Parameter	Comment
credentialsEntryParameters	The CIProperties collection of the credentials.

removeCredentialsEntry

Removes the specified credentials from the protocol.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain.
protocolName	A protocol used on the domain.
credentialsEntryID	The ID of the credential to remove.

updateCredentialsEntry

Sets new values for properties of the specified credentials entry.

The existing properties are deleted and these properties are set. Any property whose value is not set in this call is left undefined.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
domainName	The domain to update credentials in.
protocolName	A protocol used on the domain.
credentialsEntryID	The ID of the credentials to update.
credentialsEntryParameters	The CIProperties collection to set as properties for the credentials.

Data Refresh Methods**rediscoverCIs**

Locates the triggers that discovered the specified CI objects and reruns those triggers.

rediscoverCIs runs asynchronously. Call **checkDiscoveryProgress** to determine when the rediscovery is complete.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
CmdbIDs	Collection of IDs of the objects to rediscover.

Output

Parameter	Comment
isSucceed	True if the CIs rediscovery succeeded.

checkDiscoveryProgress

Returns the progress of the most recent **rediscoverCIs** call on the specified IDs. The response is a value from 0 to 1. When the response is 1, the **rediscoverCIs** call has completed.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
CmdbIDs	Collection of IDs of the objects in the rediscover call to track.

Output

Parameter	Comment
progress	A completed job has a progress of 1. Jobs that have not completed have a fraction less than 1.

rediscoverViewCIs

Locates the triggers that created the data to populate the specified view, and reruns those triggers.

rediscoverViewCIs runs asynchronously. Call **checkViewDiscoveryProgress** to determine when the rediscovery is complete.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
viewName	The views to check.

Output

Parameter	Comment
isSucceed	True if CIs rediscovery succeeded.

checkViewDiscoveryProgress

Returns the progress of the most recent **rediscoverViewCIs** call on the specified view. The response is a value from 0 to 1. When the response is 1, the **rediscoverCIs** call has completed.

Input

Parameter	Comment
cmdbContext	For details, see “CmdbContext” on page 321.
viewName	The collection of views to check.

Output

Parameter	Comment
progress	A completed job has a progress of 1. Jobs that have not completed have a fraction less than 1.



Code Sample

```
import java.net.URL;

import org.apache.axis2.transport.http.HTTPConstants;
import org.apache.axis2.transport.http.HttpTransportProperties;

import com.hp.ucmdb.generated.params.discovery.*;
import com.hp.ucmdb.generated.services.*;
import com.hp.ucmdb.generated.types.*;

public class test {
    static final String HOST_NAME = "<my_hostname>";
    static final int PORT = 21212;

    private static final String PROTOCOL = "http";
    private static final String FILE = "/axis2/services/DiscoveryService";

    private static final String PASSWORD = "<my_password>";
    private static final String USERNAME = "<my_username>";

    private static CmdbContext cmdbContext = new CmdbContext("ws tests");

    public static void main(String[] args) throws Exception {
        // Get the stub object
        DiscoveryService discoveryService = getDiscoveryService();

        // Activate Job
        discoveryService.activateJob(new ActivateJobRequest(
            "Range IPs by ICMP", cmdbContext));

        // Get domain & probes info
        getProbesInfo(discoveryService);

        // Add credentilas entry for ntcmd protocol
        addNTCMDCredentialsEntry();
    }
}
```

```

public static void addNTCMDCredentialsEntry() throws Exception {
    DiscoveryService discoveryService = getDiscoveryService();

    // Get domain name
    StrList domains =
        discoveryService.getDomainsNames(
            new GetDomainsNamesRequest(cmdbContext)).
            getDomainNames();
    if (domains.sizeStrValueList() == 0) {
        System.out.println("No domains were found, can't create credentials");
        return;
    }
    String domainName = domains.getStrValue(0);

    // Create properties with one byte param
    CIProperties newCredsProperties = new CIProperties();

    // Add password property - this is of type bytes
    newCredsProperties.setBytesProps(new BytesProps());
    setPasswordProperty(newCredsProperties);

    // Add user & domain properties - these are of type string
    newCredsProperties.setStrProps(new StrProps());
    setStringProperties("protocol_username", "test user", newCredsProperties);
    setStringProperties("ntadminprotocol_ntdomain",
        "test doamin", newCredsProperties);

    // Add new credentials entry
    discoveryService.addCredentialsEntry(
        new AddCredentialsEntryRequest(domainName,
            "ntadminprotocol", newCredsProperties, cmdbContext));

    System.out.println("new credentials craeted for domain: " + domainName + "
in ntcmd protocol");
}

```

```
private static void setPasswordProperty(CIProperties newCredsProperties) {
    BytesProp bProp = new BytesProp();
    bProp.setName("protocol_password");
    bProp.setValue(new byte[] {101,103,102,104});
    newCredsProperties.getBytesProps().addBytesProp(bProp);
}
```

```
private static void setStringProperties(String propertyName, String value,
CIProperties newCredsProperties) {
    StrProp strProp = new StrProp();
    strProp.setName(propertyName);
    strProp.setValue(value);
    newCredsProperties.getStrProps().addStrProp(strProp);
}
```

```

private static void getProbesInfo(DiscoveryService discoveryService) throws
Exception {
    GetDomainsNamesResponse result =
discoveryService.getDomainsNames(new
GetDomainsNamesRequest(cmdbContext ));

    // Go over all the domains
    if (result.getDomainNames().sizeStrValueList() > 0) {
        String domainName =
result.getDomainNames().getStrValue(0);
        GetProbesNamesResponse probesResult =
discoveryService.getProbesNames(
            new GetProbesNamesRequest(domainName, cmdbContext));

        // Go over all the probes
        for (int i=0; i<probesResult.getProbesNames().sizeStrValueList(); i++) {
            String probeName = probesResult.getProbesNames().getStrValue(i);

            // Check if connected
            IsProbeConnectedResponse connectedRequest =
discoveryService.isProbeConnected(
                new IsProbeConnectedRequest(
                    domainName, probeName, cmdbContext));
            Boolean isConnected = connectedRequest.getIsConnected();

            // Do something ...
            System.out.println("probe " + probeName + " isconnect=" +
isConnected);
        }
    }
}

```

```
private static DiscoveryService getDiscoveryService() throws Exception {
    DiscoveryService discoveryService = null;
    try {

        // Create service
        URL url = new URL(PROTOCOL,HOST_NAME,PORT, FILE);
        DiscoveryServiceStub serviceStub =
            new DiscoveryServiceStub(url.toString());

        // Authenticate info
        HttpTransportProperties.Authenticator auth =
            new HttpTransportProperties.Authenticator();
        auth.setUsername(USERNAME);
        auth.setPassword(PASSWORD);
        serviceStub._getServiceClient().getOptions().setProperty(
            HTTPConstants.AUTHENTICATE,auth);

        discoveryService = serviceStub;
    } catch (Exception e) {
        throw new Exception("cannot create a connection to service ", e);
    }

    return discoveryService;
}
}
```

12

Internal-Global ID Conversion API

This chapter includes:

Concepts

- ▶ Using the Internal-Global ID Conversion API on page 408

Tasks

- ▶ How to Convert RTSM Internal IDs to Global IDs on page 409
- ▶ How to Convert Global IDs to RTSM Internal IDs on page 411

Concepts

Using the Internal-Global ID Conversion API

Important: Use of this API requires the installation of a dedicated patch, available on the HP Software Support website.

BSM APIs work with internal CI IDs, while other tools integrated through a UCMDB work with global CI IDs. To use BSM APIs when you have global CI IDs, you must first convert them to internal CI IDs. Likewise, to send input from BSM to external software that integrates through a UCMDB, the internal CI IDs must first be converted to global CI IDs.

The Internal-Global ID Conversion API contains two methods, to convert RTSM internal CI IDs to global CI IDs, and to convert global CI IDs to RTSM internal CI IDs.

The log file for this API is located on the Gateway Server, at **<HPBSM root directory>\log\EJBContainer\bsm_sdk_utils.log**

Tasks

How to Convert RTSM Internal IDs to Global IDs

Important: Use of this API requires the installation of a dedicated patch, available on the HP Software Support website.

To convert a given internal CI ID to a global CI ID, use the following URL:

`http://<bsm.server.example>/topaz/bsmservices/customers/{customerId}/rtsm/convertRtsmToGlobalIds?cilds={cilds}`

Parameters

- ▶ `customerId` - customer ID (use 1 for non-SaaS deployment)
- ▶ `cilds` - comma separated CI IDs

If the CI ID does not exist, it will not appear in the result map.

If no global ID is defined for a CI, 32 '0' characters are retrieved as the global ID mapped to the internal ID.

Example

HTTP Command

```
http://bsm.server.example/topaz/bsmservices/customers/1/rtsm/
convertRtsmToGlobalIds?cilds=35845d64ba6250d4d45d0e75e3e59001,316e438afc58
135f9e87113bd4427d37,175fe30654e9a9dd16ddab66b2237c92,cb8a27f8ce900aa380
c41ac1c74d8153,8066fe1fcb712926fe3793a75b0e733c,bbbc37f7ce900aa380c41ac1c
74d8153
```

Result map

```

<map>
  <entry>
    <key>8066fe1fcb712926fe3793a75b0e733c</key>
    <value>00000000000000000000000000000000</value>
  </entry>
  <entry>
    <key>cb8a27f8ce900aa380c41ac1c74d8153</key>
    <value>cb8a27f8ce900aa380c41ac1c74d8153</value>
  </entry>
  <entry>
    <key>175fe30654e9a9dd16ddab66b2237c92</key>
    <value>175fe30654e9a9dd16ddab66b2237c92</value>
  </entry>
  <entry>
    <key>316e438afc58135f9e87113bd4427d37</key>
    <value>316e438afc58135f9e87113bd4427d37</value>
  </entry>
  <entry>
    <key>35845d64ba6250d4d45d0e75e3e59001</key>
    <value>00000000000000000000000000000000</value>
  </entry>
</map>

```

In this example, there are two CIs that have no global ID and will be mapped to 00000000000000000000000000000000.

Since there is no CI with ID that matches the last passed ID, it does not appear in the result map.

Error handling

Name	Error Code	Reason
BAD_REQUEST	400	Missing CI IDs or wrong input for ciIds parameter. Number of passed CI IDs reaches the limit, which is 10000.
INTERNAL_SERVER_ERROR	500	General failure to run query against RTSM

How to Convert Global IDs to RTSM Internal IDs

Important: Use of this API requires the installation of a dedicated patch, available on the HP Software Support website.

To convert a given global CI ID to an internal CI ID, use the following URL:

http://<bsm.server.example>/topaz/bsmservices/customers/{customerId}/rtsm/convertGlobalToRtsmIds?cilds={cilds}

Parameters

- customerId - customer ID (use 1 for non-SaaS deployment)
- cilds - comma separated CI IDs

If no CI with the given global ID appears in the RTSM model, no entry will appear in the result map.

Example

HTTP Command

```
http://bsm.server.example/topaz/bsmservices/customers/1/rtsm/
convertGlobalToRtsmIds?cilds=35845d64ba6250d4d45d0e75e3e59001,316e438afc58
135f9e87113bd4427d37,175fe30654e9a9dd16ddab66b2237c92,cb8a27f8ce900aa380
c41ac1c74d8153
```

Result map

```
<map>
  <entry>
    <key>316e438afc58135f9e87113bd4427d37</key>
    <value>316e438afc58135f9e87113bd4427d37</value>
  </entry>
  <entry>
    <key>175fe30654e9a9dd16ddab66b2237c92</key>
    <value>175fe30654e9a9dd16ddab66b2237c92</value>
  </entry>
  <entry>
    <key>cb8a27f8ce900aa380c41ac1c74d8153</key>
```

```
<value>cb8a27f8ce900aa380c41ac1c74d8153</value>  
</entry>  
</map>
```

Since there is no CI with a global ID that matches the first ID, it does not appear in the result map.

Error handling

Name	Error Code	Reason
BAD_REQUEST	400	Missing CI IDs or wrong input for ciIds parameter. Number of passed CI IDs reaches the limit, which is 10000.
INTERNAL_SERVER_ERROR	500	General failure to run query against RTSM

Index

A

- accessing data
 - guidelines 30
- adapter
 - add for new external data source 236
 - deploying 148, 243
 - loading 148
- adapter code 38
- adapter resources
 - debugging 236
- adapter.conf 162
- adapters
 - assigning jobs to 49
 - creating 40
 - defining input (Trigger CIT, Input TQL) 41
 - defining output 46
 - development and testing 23
 - finding correct credentials for connections 77
 - implementing 37
 - interaction with the federation framework 214
 - interfaces 233
 - modifying existing 28
 - overriding parameters 48
 - packaging and productization 24
 - preparations before creating 124
 - preparing package 129
 - prerequisites 124
 - scheduling 50
 - separating 35
 - Trigger TQL 49
 - upgrading from 9.00 and 9.01 133
 - writing new pattern 29

- adapter-writing
 - introduction 20
 - research stage 28
- API 291
 - DDM webservice 385
 - Internal-Global ID Conversion 407
 - ODB Web service 305
- APIs
 - included with HP Business Service Management 292
 - RTSM Java
 - RTSM Java API 293
 - Web service overview 386

B

- BDM
 - accessing documentation 60
- blueprint 27

C

- character set
 - determining encoding 82
- configuration files for the generic database
 - adapter 161
- configuration type
 - UCMDB Web service API 323
- content
 - creating 21
- content development and adapter-writing 19
- converters
 - generic database adapter 186

D

- Data Flow Management
 - Web Service, adding credentials
 - example 379
- data push flow 211
- data source
 - add adapter for new data source 236
- database adapter
 - configuration examples 191
- DDM
 - Web service APIs, overview 386
 - webservice 385
 - webservice, calling 388
 - webservice, exceptions 388
 - webservice, managing query methods 390
 - webservice, mapping methods 389
 - webservice, permissions 387
- debug adapter resources 236
- deploying an adapter 243
- derived properties 314
- DFM
 - development cycle 21
 - discovery adapters and related components 34
 - integration 24
- DFM code
 - recording 109
- differential synchronization 252, 259
- Discovery
 - content migration guidelines 54
 - content migration guidelines, new infrastructure features 54
 - content migration, accessing BDM
 - documentation 60
 - content migration, implementation tips 59
 - cross-data model script development
 - guidelines 59
 - migrating content 53
 - package migration for content
 - migration guidelines 61
- discovery
 - business value 26
- discovery adapters
 - implementing 37

- discovery adapters and related components 34
- Discovery Analyzer
 - running from Eclipse 99
 - working with 90
- discovery content
 - developing 34
- DiscoveryMain function 69
- discriminator.properties 184
- documentation updates 16

E

- Eclipse
 - mapping between CI attributes and database tables 151
 - running Discovery Analyzer 99
- encoding
 - determining for character sets 82
- executeCommandAndDecode
 - method 88

F

- federated database adapter
 - troubleshooting 205
- federation flow 209
- federation framework
 - adapter and mapping interaction 214
 - adapter interfaces 233
 - overview 208
- fixed_values.txt 186
- Framework instance 73

G

- GDBA adapter
 - TQL queries 119
- generic database adapter
 - configuration files 161
 - converters 186
 - overview 119
 - plugins 191
 - reconciliation 120

getCharsetName
 method 88
 getLanguageBundle
 method 89

H

Hibernate mapping tool 121
 HP Data Flow Management API Reference 64
 HP Software Support Web site 15
 HP Software Web site 16

I

integration
 federation framework flow for
 federated TQL queries 215
 federation framework flow for
 population 231
 integration content
 developing 31
 Integration framework SDK 207
 Internal-Global ID Conversion API 407

J

Java
 RTSM API 293
 Java adapter
 create sample 247
 Java adapters
 developing 207
 XML configuration tags 249
 Java exceptions
 handling 79
 Jython
 generating results 71
 libraries and utilities 112
 structure of the file 68
 Jython adapters
 developing 63
 localization 79
 Jython scripts
 writing 255

K

Knowledge Base 15
 KPIs
 view in external application 285

L

log files
 enabling 150
 for federated database 203
 logger.py 113

M

mapping
 interaction with the federation
 framework 214
 mapping file
 schema 263, 278
 mapping files
 preparing 253
 methods
 executeCommandAndDecode 88
 getCharsetName 88
 getLanguageBundle 89
 useCharset 89
 modeling.py 114
 multi-lingual locales
 adding support for new language 80
 API reference 86
 changing default 82
 decoding commands without
 keyword 84
 writing a new job 83

N

netutils.py 114

O

ODB
 querying
 Web service 311
 online resources 15

Index

orm.xml 166
osLanguage 89

P

persistence.xml 182
plugins
 generic database adapter 191
 implementing 144
properties
 derived 314
Push adapter
 building package 261
push adapter mapping file
 schema 263, 278
Push adapters
 developing 252

R

reconciliation_rules.txt 179
reconciliation_types.txt 179
relation
 UCMDB Web service API 323
replication_config.txt 186
reports
 viewing 150
resource bundles 85

S

scripts
 modifying out-of-the-box 66
shellutils.py 115
simplifiedConfiguration.xml 163
synchronization
 supporting differential 259

T

third-party integrations 291
TQL
 TQL queries for Generic Database
 adapter 119
transformations.txt 181
Troubleshooting and Knowledge Base 15

U

UCMDB Java API
 application structure 296
 integration user, creating 299
 jar file 299
 permissions 295
 using 294
UCMDB Web Service API
 errors 311
 exceptions 311
 getCmdbClassDefinition 319
 getQueryNameOfView 336
 parameter format 316
UCMDB Web service API
 using 308
 Web Service, calling 311
UCMDB Web service API
 addCIsAndRelations 340
 addCustomer 341
 calculateImpact 343
 chunkInfo 326
 CIT name 323
 class name 323
 configuration type name 323
 deleteCIsAndRelations 342
 executeTopologyQueryByNameWith
 Parameters 326
 executeTopologyQueryWithParamete
 rs 327
 getAllClassesHierarchy 319
 getChangedCIs 328
 getCIsByID 330
 getCIsByType 331
 getClassAncestors 318
 getFilteredCIsByType 331
 getImpactPath 344
 getImpactRulesByNamePrefix 345
 getTopologyQueryExistingResultByN
 ame 336
 getTopologyQueryResultCountByNa
 me 337
 identifier in impact analysis methods
 320
 inherited properties query 335
 key attributes 321
 parameter format 321, 324

- permissions 309
- query methods 326
- query the UCMDB class model 318
- query, properties returned 313
- relation 323
- removeCustomer 342
- ShallowRelation 325
- update methods 340, 343
- updateCIsAndRelations 342
- updates, documentation 16
- useCharset
 - method 89

V

- views
 - creating 149, 150

W

- Web Service
 - ODB API 305
 - UCMDB Web Service API 311
- Web service APIs
 - overview 386
- webservice
 - in DDM 385

X

- XML configuration tags 249

