



HP Operations Manager i

Software Version: 10.00

RTSM Developer Reference Guide

Document Release Date: January 2015
Software Release Date: January 2015

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 Notice

© Copyright 2015 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.

Citrix® and XenDesktop® are registered trademarks of Citrix Systems, Inc. and/or one more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries.

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

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

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

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.

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

UNIX® is a registered trademark of The Open Group.

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: <https://softwaresupport.hp.com/group/softwaresupport/search-result?keyword=>.

This site requires an HP Passport account. If you do not have one, click the **Create an account** button on the HP Passport Sign in page.

Support

Visit the HP Software Support web site at: <https://softwaresupport.hp.com>

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

HP Software 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 <https://softwaresupport.hp.com> and click **Register**.

To find more information about access levels, go to: <https://softwaresupport.hp.com/web/softwaresupport/access-levels>

HP Software Solutions & Integrations and Best Practices

Visit HP Software Solutions Now at <https://h20230.www2.hp.com/sc/solutions/index.jsp> to explore how the products in the HP Software catalog work together, exchange information, and solve business needs.

Visit the Cross Portfolio Best Practices Library at <https://hpln.hp.com/group/best-practices-hpsw> to access a wide variety of best practice documents and materials.

About this PDF Version of Online Help

This document is a PDF version of the online help. This PDF file is provided so you can easily print multiple topics from the help information or read the online help in PDF format. Because this content was originally created to be viewed as online help in a web browser, some topics may not be formatted properly. Some interactive topics may not be present in this PDF version. Those topics can be successfully printed from within the online help.

Contents

Part 1: Creating Discovery and Integration Adapters	10
Chapter 1: Adapter Development and Writing	11
Adapter Development and Writing Overview	11
Content Creation	11
The Adapter Development Cycle	12
Data Flow Management and Integration	15
Associating Business Value with Discovery Development	16
Researching Integration Requirements	17
Developing Integration Content	19
Developing Discovery Content	21
Discovery Adapters and Related Components	22
Separating Adapters	23
Implement a Discovery Adapter	24
Step 1: Create an Adapter	26
Step 2: Assign a Job to the Adapter	33
Step 3: Create Jython Code	35
Configure Remote Process Execution	35
Chapter 2: Developing Jython Adapters	37
HP Data Flow Management API Reference	37
Create Jython Code	37
Use External Java JAR Files within Jython	38
Execution of the Code	38
Modifying Out-of-the-Box Scripts	38
Structure of the Jython File	39
Imports	40
Main Function – DiscoveryMain	40
Functions Definition	40
Results Generation by the Jython Script	42
The ObjectStateHolder Syntax	42
Sending Large Amounts of Data	43
The Framework Instance	44
Finding the Correct Credentials (for Connection Adapters)	47
Handling Exceptions from Java	49
Troubleshooting Migration from Jython Version 2.1 to 2.5.3	50
Support Localization in Jython Adapters	52

Add Support for a New Language	52
Change the Default Language	53
Determine the Character Set for Encoding	54
Define a New Job to Operate With Localized Data	54
Decode Commands Without a Keyword	55
Work with Resource Bundles	56
API Reference	57
Record DFM Code	59
Jython Libraries and Utilities	60
Chapter 3: Developing Generic Database Adapters	64
Generic Database Adapter Overview	65
TQL Queries for the Generic Database Adapter	65
Reconciliation	66
Hibernate as JPA Provider	67
Prepare for Adapter Creation	69
Prepare the Adapter Package	74
Configure the Adapter – Minimal Method	77
Configure the adapter.conf File	78
Example: Populating a Node and IP Address using the Simplified Method	78
Configure the Adapter – Advanced Method	82
Implement a Plug-in	87
Deploy the Adapter	90
Edit the Adapter	90
Create an Integration Point	90
Create a View	90
Calculate the Results	91
View the Results	91
View Reports	91
Enable Log Files	91
Use Eclipse to Map Between CIT Attributes and Database Tables	92
Adapter Configuration Files	99
The adapter.conf File	101
The simplifiedConfiguration.xml File	102
The orm.xml File	104
The reconciliation_types.txt File	118
The reconciliation_rules.txt File (for backwards compatibility)	118
The transformations.txt File	120
The discriminator.properties File	121
The replication_config.txt File	122

The fixed_values.txt File	122
Persistence.xml File	122
Connect to Database Using NT Authentication	123
Out-of-the-Box Converters	124
Plug-ins	129
Configuration Examples	129
Adapter Log Files	138
External References	140
Troubleshooting and Limitations	140
Chapter 4: Developing Java Adapters	142
Federation Framework Overview	142
Adapter and Mapping Interaction with the Federation Framework	147
Federation Framework for Federated TQL Queries	147
Interactions between the Federation Framework, Server, Adapter, and Mapping Engine ..	149
Federation Framework Flow for Population	158
Adapter Interfaces	159
Debug Adapter Resources	161
Add an Adapter for a New External Data Source	161
Create a Sample Adapter	169
XML Configuration Tags and Properties	170
The DataAdapterEnvironment Interface	172
OutputStream openResourceForWriting(String resourceName) throws FileNotFoundException;	172
InputStream openResourceForReading(String resourceName) throws FileNotFoundException;	173
Properties openResourceAsProperties(String propertiesFile) throws IOException; ...	173
String openResourceAsString(String resourceName) throws IOException;	174
public void saveResourceFromString(String relativeFileName, String value) throws IOException;	174
boolean resourceExists(String resourceName);	175
boolean deleteResource(String resourceName);	175
Collection<String> listResourcesInPath(String path);	175
DataAdapterLogger getLogger();	176
DestinationConfig getDestinationConfig();	176
int getChunkSize();	176
int getPushChunkSize();	176
ClassModel getLocalClassModel();	176
CustomerInformation getLocalCustomerInformation();	176
Object getSettingValue(String name);	177
Map<String, Object> getAllSettings();	177

boolean isMTEnabled();	177
String getUcmdbServerHostName();	177
Chapter 5: Developing Push Adapters	178
Developing and Deploying Push Adapters	178
Build an Adapter Package	179
Troubleshooting	181
TQL Best Practices for Push Adapters	182
Create Mappings	182
Build a Mapping File	183
Prepare the Mapping Files	183
Write Jython Scripts	186
Support Differential Synchronization	190
Generic XML Push Adapter SQL Queries	192
Generic Web Service Push Adapter	192
Mapping File Reference	211
Mapping File Schema	213
Mapping Results Schema	224
Customization	228
Chapter 6: Developing Enhanced Generic Push Adapters	229
Overview	229
The Mapping File	229
The Groovy Traveler	232
Write Groovy Scripts	235
Implement PushConnector Interface	236
Build an Adapter Package	237
Mapping File Schema	238
Chapter 7: Viewing KPIs in External Applications	245
Set Up an Adapter to View KPIs in an External Application	245
Part 2: Using APIs	247
Chapter 8: Introduction to APIs	248
APIs Overview	248
Chapter 9: RTSM API	249
Conventions	249
Using the RTSM API	249
General Structure of an Application	251
Put the API Jar File in the Classpath	253
Create an Integration User	253
UCMDB API Use Cases	255

Examples	256
Chapter 10: RTSM Web Service API	258
Conventions	258
RTSM Web Service API Overview	259
RTSM Web Service API Reference	262
Call the Web Service	262
Query the RTSM	263
Update the RTSM	266
Query the RTSM Class Model	267
getClassAncestors	267
getAllClassesHierarchy	268
getCmdbClassDefinition	268
Query for Impact Analysis	269
RTSM General Parameters	269
RTSM Output Parameters	272
RTSM Query Methods	273
executeTopologyQueryByNameWithParameters	274
executeTopologyQueryWithParameters	274
getChangedCIs	275
getCINeighbours	276
getCIsByID	277
getCIsByType	277
getFilteredCIsByType	278
getQueryNameOfView	281
getTopologyQueryExistingResultByName	282
getTopologyQueryResultCountByName	282
pullTopologyMapChunks	283
releaseChunks	284
UCMDB Update Methods	285
addCIsAndRelations	285
addCustomer	286
deleteCIsAndRelations	286
removeCustomer	287
updateCIsAndRelations	287
RTSM Impact Analysis Methods	288
calculateImpact	288
getImpactPath	289
getImpactRulesByNamePrefix	289
RTSM Web Service API Use Cases	290

Examples	291
Chapter 11: Internal-Global ID Conversion API	292
Using the Internal-Global ID Conversion API	292
How to Convert RTSM Internal IDs to Global IDs	292
Parameters	292
Example	292
How to Convert Global IDs to RTSM Internal IDs	294
Parameters	294
Example	294
Chapter 12: Data Flow Management Java API	296
Using the Data Flow Management Java API	296
Chapter 13: Data Flow Management Web Service API	298
Data Flow Management Web Service API Overview	298
Conventions	299
Calling the Web Service	299
Data Flow Management Methods and Data Structures	299
Data Structures	300
Managing Discovery Job Methods	301
Managing Trigger Methods	302
Domain and Probe Data Methods	304
Credentials Data Methods	307
Data Refresh Methods	309
Code Sample	310
Adding Credentials Example	313
Send Documentation Feedback	317

Part 1: Creating Discovery and Integration Adapters

Chapter 1: Adapter Development and Writing

This chapter includes:

Adapter Development and Writing Overview	11
Content Creation	11
Developing Integration Content	19
Developing Discovery Content	21
Implement a Discovery Adapter	24
Step 1: Create an Adapter	26
Step 2: Assign a Job to the Adapter	33
Step 3: Create Jython Code	35
Configure Remote Process Execution	35

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.)

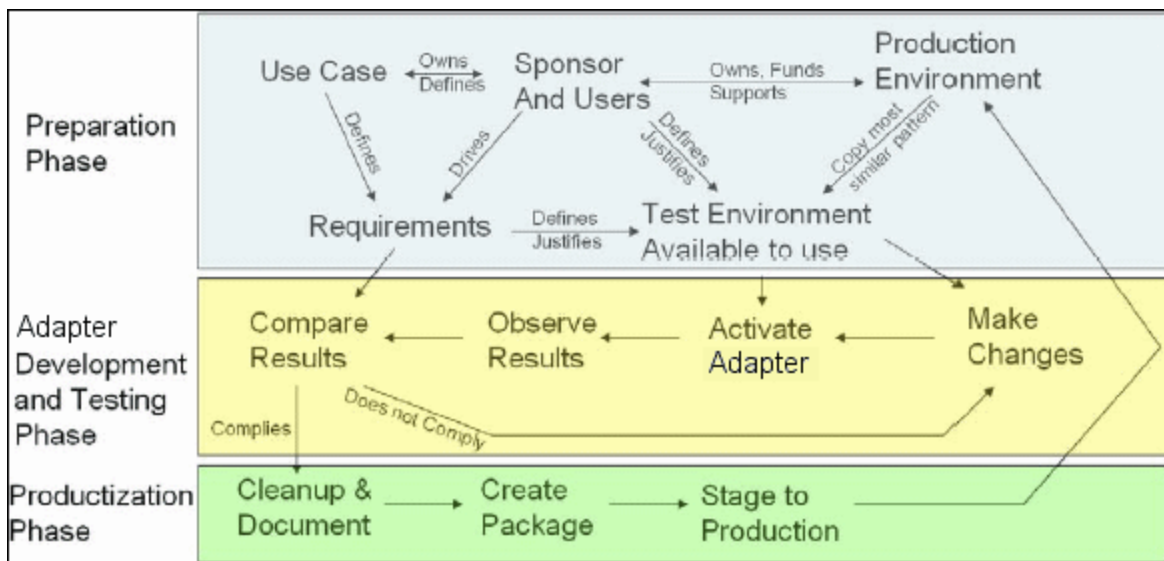
Content Creation

This section includes:

- ["The Adapter Development Cycle" below](#)
- ["Data Flow Management and Integration" on page 15](#)
- ["Associating Business Value with Discovery Development" on page 16](#)
- ["Researching Integration Requirements" on page 17](#)

The 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.



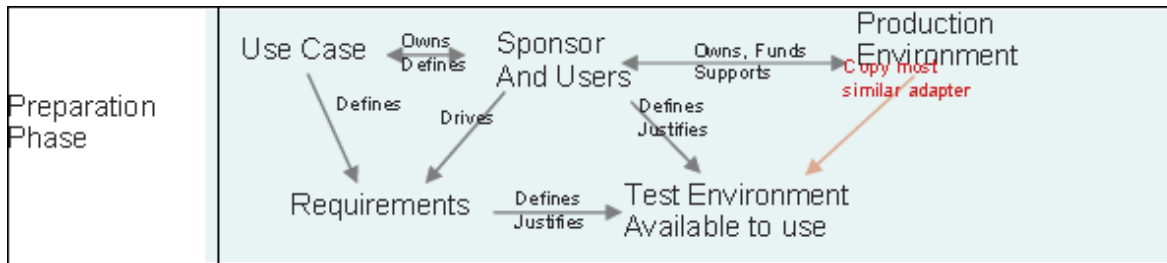
Each phase of adapter development builds on the last one.

Once you are satisfied with the way the adapter looks and works, you are ready to package it. Using either the RTSM Package Manager in the OMi UI 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" on the next page](#)
- ["Adapter Development and Testing" on the next page](#)
- ["Adapter Packaging and Productization " on page 14](#)

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 *Data Flow Management Guide*.
2. Research the method which the adapter should use to 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 the entire package.
5. Unzip the package contents into the 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 use of the content, your plan should answer the following questions:

- 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 to 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 11](#).

This section includes the following topics:

- ["Modifying an Existing Adapter" below](#)
- ["Writing a New Adapter" below](#)
- ["Model Research" on the next page](#)
- ["Technology Research" on the next page](#)
- ["Guidelines for Choosing Ways to Access Data" on the next page](#)
- ["Summary" on page 19](#)

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 RTSM 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 in the **Data Flow Management > Data Flow Probe Setup > Domains and Probes pane**. For details about each protocol, see the section describing the supported protocols in the *HP Universal CMDB Discovery and Integration Content 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?

If you answer yes to these questions, then **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?

If you answer yes to these questions, then **Federation** is needed.

- Should the integration push data to remote data sources?

If you answer yes to these questions, then **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 *Data Flow Management Guide*.

Five different options are available for creating integration adapters:

1. Jython Adapter:

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

For details, see ["Developing Jython Adapters" on page 37](#).

2. 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 connect to a single data source or target.

For details, see ["Developing Java Adapters" on page 142](#).

3. Generic DB Adapter:

- An abstract adapter based on the Java Adapter that uses the Federation SDK Framework.
- Allows creation of adapters that connect to external data repositories.
- Supports both Federation and Population (with a Java plug-in 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 connect to a single data source.

For details, see ["Developing Generic Database Adapters" on page 64](#).

4. 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 connect to a single data target.
- Used for Data Push.

For details, see ["Developing Push Adapters" on page 178](#).

5. Enhanced Generic Push Adapter:

- All the above features of the Generic Push Adapter
- A root-element-based adapter
- Maps a UCMDB tree data structure to a target tree data structure

For details, see ["Developing Enhanced Generic Push Adapters" on page 229](#).

The following table displays the capabilities of each adapter:

Flow/Adapter	Jython Adapter	Java Adapter	Generic DB Adapter	Generic Push Adapter	Enhanced Generic Push Adapter
Population	✓	✓	✓	✗	✗
Federation	✗	✓	✓	✗	✗
Data Push	✗	✓	✗	✓	✓

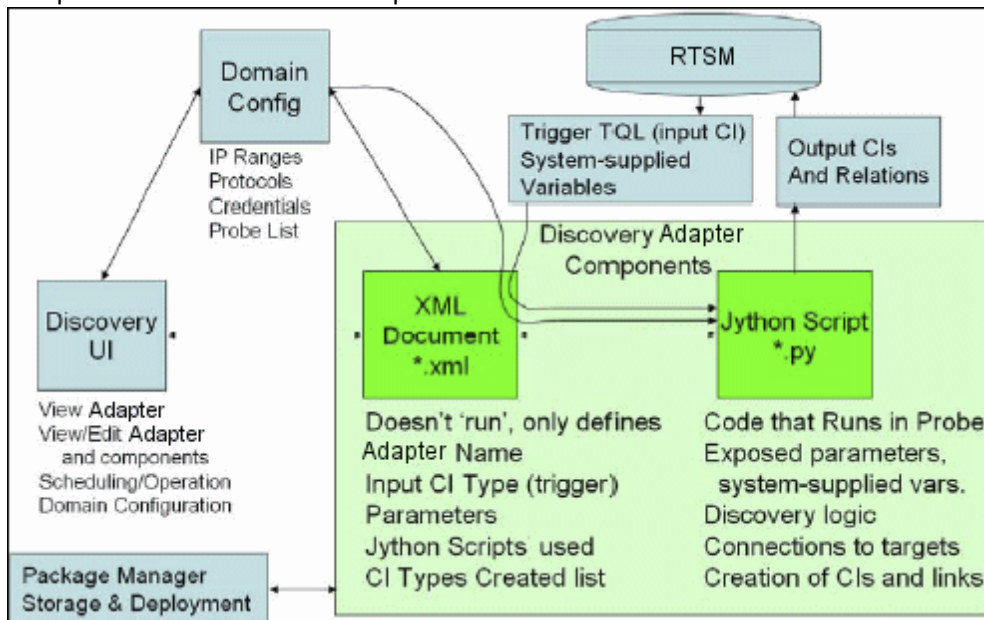
Developing Discovery Content

This section includes:

- ["Discovery Adapters and Related Components " below](#)
- ["Separating Adapters" on the next page](#)

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 Universal Discovery, 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

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.

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 an 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 ["Step 1: Create an Adapter" on page 26](#).

2. **Create a Discovery job.**

You configure a job with scheduling information and a trigger query. For details, see ["Step 2: Assign a Job to the Adapter" on page 33](#).

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 35](#).

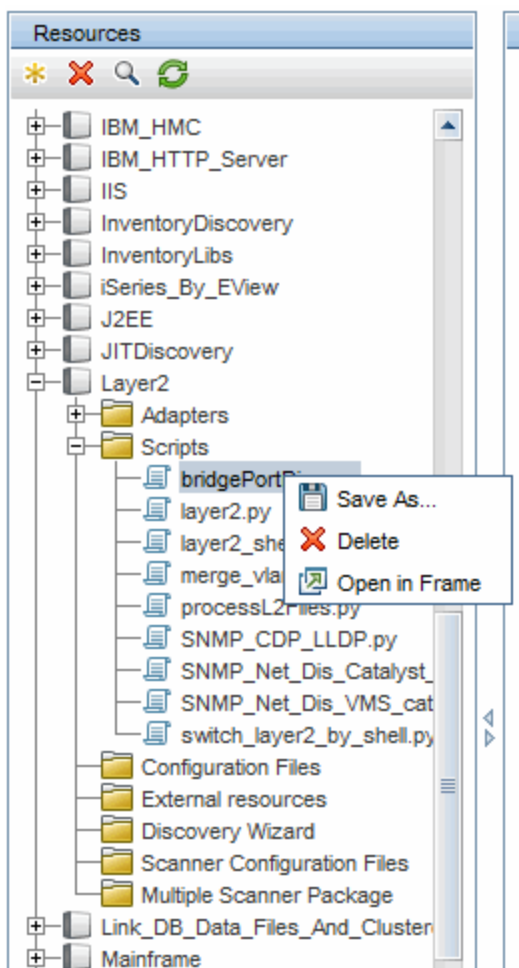
To write new adapters, you create each of the above components, each one of which is automatically bound to the component in the previous step. For example, once you create a job and select the relevant adapter, the adapter file binds to the job.

Adapter Code

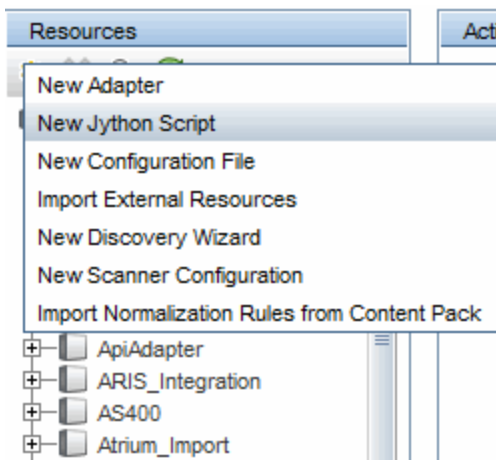
The actual implementation of connecting to the remote system, querying its data, and mapping it as RTSM data is performed by the Jython code. For example, the code contains the logic for connecting to a database and extracting data from it. In this case, the code expects to receive a JDBC URL, a user 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 35](#). 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 *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 [Data Flow Management Concepts](#) in the *Data Flow Management Guide*.

This section includes the following topics:

- ["Define Adapter Input \(Trigger CIT and Input Query\)" below](#)
- ["Define Adapter Output" on page 29](#)
- ["Override Adapter Parameters" on page 31](#)
- ["Override Probe Selection - Optional" on page 32](#)
- ["Configure a classpath for a remote process - Optional" on page 33](#)

1. 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 [How to Add Query Nodes and Relationships to a TQL Query](#) in the 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.
- If the value of the destination data is a list, you can define the number of items from the list to be sent to the probe. To define it, add a colon after the default value followed by the number of items. If there is no default value for the destination data, enter two colons.

For example, if the following destination data is entered: `name=portId, value=${PHYSICALPORT.root_id:NA:1}` or `name=portId, value=${PHYSICALPORT.root_id:1}`, only the first port from the port list is sent to the probe.

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.

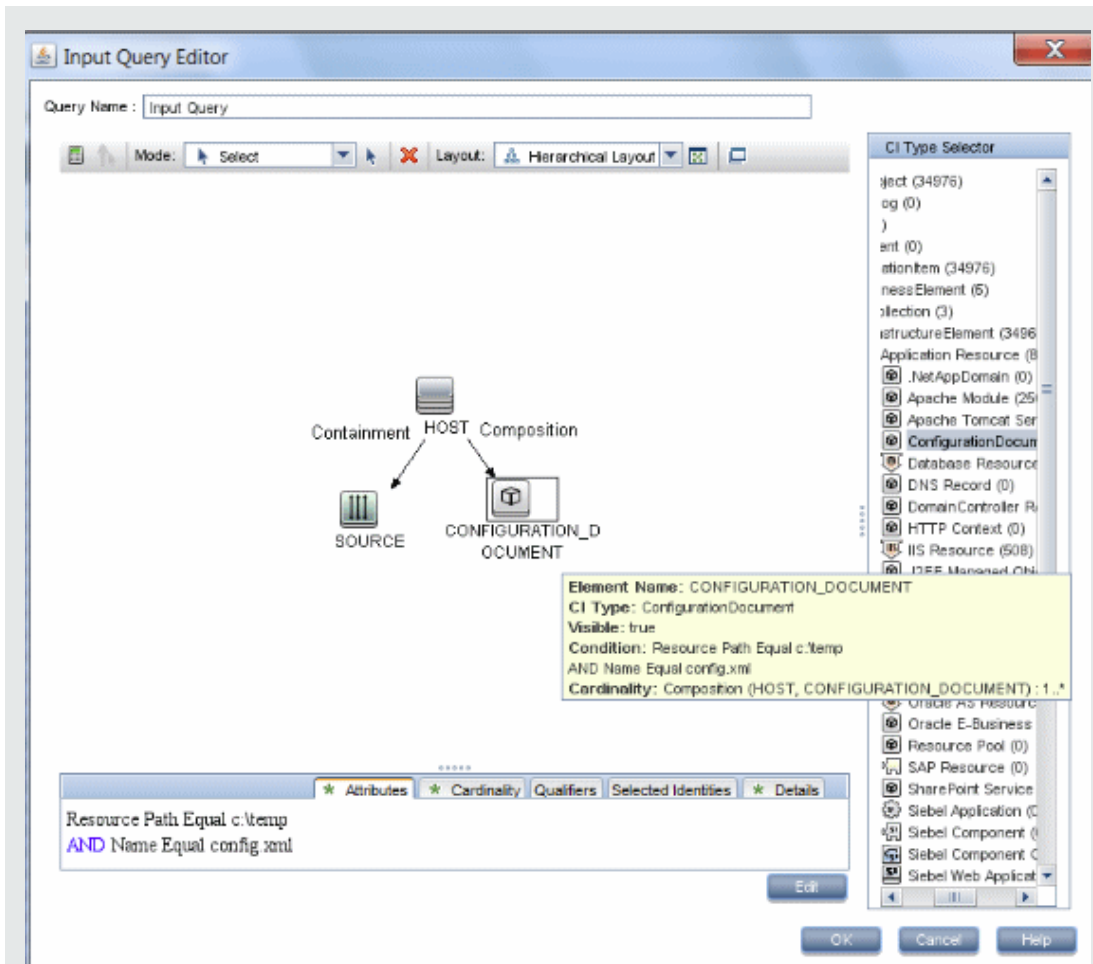
The Triggered CI data for the **IpAddress** CI includes a `fileName` variable. This variable enables the replacement of the **CONFIGURATION_DOCUMENT** node in the Input TQL with the actual values of the configuration file located on a host:

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).



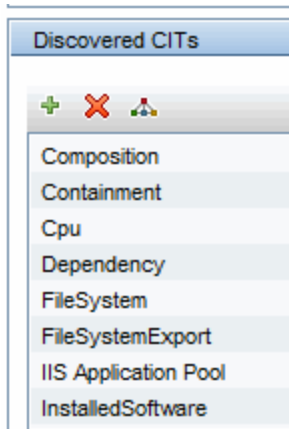
Click thumbnail to view full size image.

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	

2. Define Adapter Output

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



You can also view the CIs as a topology map, that is, the components and the way in which they are linked together (click the **View Discovered CIs 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 42](#).

Example of Adapter Output:

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

- a. Access **Administration > RTSM Administration > Data Flow Management > Adapter Management**.
- b. In the Resources pane, select **Network > Adapters > NSLOOKUP_on_Probe**.
- c. In the Adapter Definition tab, locate the Discovered CITs pane.
- d. The CITs that are to be part of the adapter output are listed. Add CITs to, or remove from, the list. For details, see [Adapter Definition Tab](#) in the *Data Flow Management Guide*.


3. 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.

- a. Access **Administration > RTSM Administration > Data Flow Management > Adapter Management**.
- b. In the Resources pane, select **Database_Basic > Adapters > SQL_NET_Dis_Connection**.
- c. In the Adapter Definition tab, locate the **Adapter Parameters** pane. The `protocolType` parameter has a value of **all**:



Name	Value
protocolType	all

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

Parameters		
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 Definition Tab](#) in the *Data Flow Management Guide*.

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

4. Override Probe Selection - Optional

On the UCMDDB server there is a dispatching mechanism that takes the trigger CIs received by the UCMDDB and automatically chooses which probe should run the job for each trigger CI according to one of the following options:

- **For the IP address CI type:** Take the probe that is defined for this IP.
- **For the running software CI type:** Use the attributes **application_ip** and **application_ip_domain** and choose the probe that is defined for the IP in the relevant domain.
- **For other CI types:** Take the node's IP according to the CI's related node (if it exists).

The automatic probe selection is done according to the CI's related node. After obtaining the CI's related node, the dispatching mechanism chooses one of the node's IPs and chooses the probe according to the probe's network scope definitions.

In the following cases, you need to specify the probe manually and not use the automatic dispatching mechanism:


- You already know which probe should be run for the adapter and you do not need the automatic dispatching mechanism to select the probe (for example if the trigger CI is the probe gateway).
- The automatic probe selection might fail. This can happen in the following situations:
 - A trigger CI does not have a related node (such as the network CIT)
 - A trigger CI's node has multiple IPs, each belonging to a different probe.

To manually specify which probe to use with the adapter:

- Select the adapter and click the **Adapter Configuration** tab.
- Under **Trigger Dispatch Options**, select **Override default probe selection**.
- In the box, enter the Probe in one of the following formats:

Probe name	The name of the Probe
IP address	The Probe's IP address—can be defined in either IPv4 or IPv6 format
IP,Domain	IPv4 format: 16.59.63.86,DefaultDomain IPv6 format: 2001:0:9d46:953c:34a9:1e6b:f2ff:ffe,CustomDomain
Domain name	The domain from which the Probe should be selected.

For example:



The screenshot shows the 'Adapter Configuration' tab in a software interface. Under the 'Trigger Dispatch Options' section, the 'Override default probe selection' checkbox is checked, and the text box next to it contains the value '\${SOURCE.name}'. The 'Supports IPv6' checkbox is unchecked. Below this is the 'Execution Options' section, which includes a 'Create communication log' dropdown set to 'On failure', an 'Include results in communication log' section with 'Yes' and 'No' radio buttons (where 'No' is selected), and a 'Max. threads' text box.

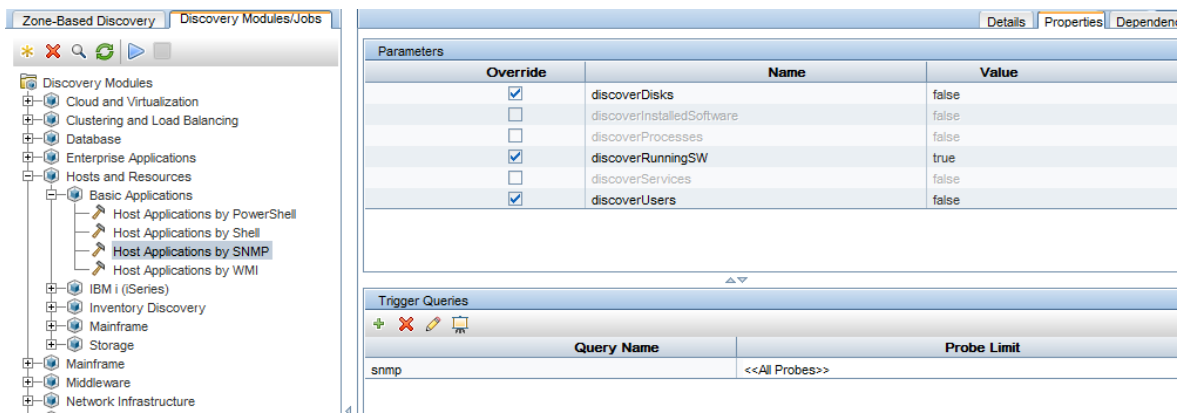
5. Configure a classpath for a remote process - Optional

For details, see ["Configure Remote Process Execution" on page 35](#).

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 sets 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, as shown in the picture below.



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 *Data Flow Management Guide*.

Override Adapter Parameters

When configuring a job you can override the adapter parameters. For details, see ["Override Adapter Parameters" on page 31](#).

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 "[Create Jython Code](#)" on page 37.

Configure Remote Process Execution

You can run discovery for a discovery job in a process separate from the Data Flow Probe's process.

For example, you can run the job in a separate remote process if the job uses **.jar** libraries that are a different version than the Probe's libraries or that are incompatible with the Probe's libraries.

You can also run the job in a separate remote process if the job potentially consumes a lot of memory (brings a lot of data) and you want to isolate the Probe from potential **OutOMemory** problems.

To configure a job to run as a remote process, define the following parameters in its adapter's configuration file:

Parameter	Description
remoteJVMArgs	JVM parameters for the remote Java process.
runInSeparateProcess	When set to true , the discovery job runs in a separate process.

Parameter	Description
remoteJVMClasspath	<p>(Optional) Enables customization of the classpath of the remote process, overriding the default Probe classpath. This is useful if there might be version incompatibility between the Probe's jars and custom jars required for the customer-defined discovery.</p> <p>If the remoteJVMClasspath parameter is not defined, or is left empty, the default Probe classpath is used.</p> <p>If you develop a new discovery job and you want to ensure that the Probe jar library version does not collide with the job's jar libraries, you must use at least the minimal classpath required to execute basic discovery. The minimal classpath is defined in the DataFlowProbe.properties file in the basic_discovery_minimal_classpath parameter.</p> <p>Examples of remoteJVMClasspath customization:</p> <ul style="list-style-type: none">• To prepend or append custom jars to the default Probe classpath. customize the remoteJVMClasspath parameter as follows: <code>custom1.jar;%classpath%;custom2.jar -</code> In this case, custom1.jar is placed before default Probe classpath, and custom2.jar is appended to the Probe classpath.• To use the minimal classpath, customize the remoteJVMClasspath parameter as follows: <code>custom1.jar;%minimal_classpath%;custom2.jar</code>

Chapter 2: Developing Jython Adapters

This chapter includes:

HP Data Flow Management API Reference	37
Create Jython Code	37
Support Localization in Jython Adapters	52
Record DFM Code	59
Jython Libraries and Utilities	60

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:

\\< Gateway Server root directory>\AppServer\webapps\site.war\amdocs\eng\APIs\DDM_JavaDoc\index.html

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 to 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 Universal Discovery should be compatible with Jython version 2.5.3.
- 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" below](#)
- ["Execution of the Code" below](#)
- ["Modifying Out-of-the-Box Scripts" below](#)
- ["Structure of the Jython File" on the next page](#)
- ["Results Generation by the Jython Script" on page 42](#)
- ["The Framework Instance" on page 44](#)
- ["Finding the Correct Credentials \(for Connection Adapters\)" on page 47](#)
- ["Handling Exceptions from Java" on page 49](#)

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 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 Run-time 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 = ipplanet_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 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.createClient(Framework.getTriggerCIData
        (BaseClient.CREDENTIALS_ID))
    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 (**Administration > 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 CI name and a list of attributes and their values. The `ObjectStateHolderVector` is a vector of `ObjectStateHolder` instances.

The ObjectStateHolder Syntax

This section explains 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', sbIsvrName) 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 **sbIsvrName** 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 *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 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 that 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"](#) on page 62.

Sending Large Amounts of Data

Sending large amounts of data (usually more than 20 KB) is difficult to process in the RTSM. Data of this size should be split into smaller chunks before sending to the RTSM. In order for all the chunks to be correctly inserted to the RTSM, each chunk needs to contain required identification information for the CIs in the chunk. This is a common scenario when developing Jython integrations. The [sendObjects](#) method is used to send the results in chunks. If the Jython script sends a large number of results (the default value is 20,000, but this value can be configured in the ["DataFlowProbe.properties File"](#) using the **appilog.agent.local.maxTaskResultSize** key) it should chunk the results according to their topology. This chunking should be performed taking into account identification rules so that the results are entered correctly in the RTSM. If the Jython script does not chunk the results, the probe attempts to chunk them; however, this can lead to poor performance for a large result set.

Note: Chunking should be used for Jython integration adapters and not for regular discovery jobs. This is because discovery jobs usually discover information regarding a specific trigger and do not send large amounts of information. With Jython integrations, large amounts of data are discovered on the single trigger of the integration.

It is also possible to use chunking for a small number of results. In such a case, there is a relationship between CIs in different chunks and the developer of the Jython script has two options:

- Send the entire CI and all of its identification information again in every chunk that contains a link to it.
- Use the RTSM ID of the CI. To do this, the Jython script has to wait for each chunk to be processed in the RTSM server in order to get the RTSM IDs. To enable this mode (called synchronous result sending), add the `SendJythonResultsSynchronously` tag to the adapter. This tag ensures that when you finish sending the chunk, the RTSM IDs of the CIs in the chunk have already been received by the probe. The adapter developer can use the UCMDB IDs for generating the next chunk. To use the RTSM IDs, use the framework API `getIdMapping`.

Example of Using `getIdMapping`

In the first chunk you send nodes. In the second chunk you send processes. The root container of the process is a node. Instead of sending the entire `objectStateHolder` of the node in the process `root_container` attribute, you can get the RTSM ID of the node using the `getIdMapping` API and use only the node ID in the process `root_container` attribute to make the chunk smaller.

The 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 37](#).

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 the next page](#)
- ["Framework.createClient\(credentialsId, props\)" on the next page](#)
- ["Framework.getParameter \(String parameterName\)" on page 46](#)
- ["Framework.reportError\(String message\) and Framework.reportWarning\(String message\)" on page 47](#)

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:

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}

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 the *HP Universal CMDB Discovery and Integration Content 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:

BaseClient
getCredentialId()
getIpAddress()
getProtocol()

*SSH
Telnet
Ntcmd*
ShellClient
isWinOS()
isUnixOS()
executeCmd()

*Wmi
Snmp
Sql*
QueryClient
executeCmd()

For details on the clients and their supported APIs, see [BaseClient](#), [ShellClient](#), and [QueryClient](#) in the [DFM Framework](#). These files are located in the following folder:

\\<HP Operations Manager i root directory>\AppServer\webapps\site.war\amdocs\eng\APIs\DDM_Schema\webframe.html

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

HP Operations Manager i (10.00)

Page 46 of 317

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.createClient(Framework.getTriggerCIData
    (BaseClient.CREDENTIALS_ID))

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.

For details on error messages, see [Error Messages](#).

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 is the credentials ID. The connection script gains access to possible credential sets and tries them one by one using the Framework.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).

Note: Access to sensitive data (passwords, private keys, and so on) is blocked for the following protocol types:

sshprotocol, ntadminprotocol, as400protocol, vmwareprotocol, wmiprotocol, vcloudprotocol, sapjmxprotocol, websphereprotocol, siebelgtwyprotocol, sapprotocol, ldapprotocol, udaprotocol, ntcmdprotocol, snmpprotocol, jbossprotocol, telnetprotocol, powershellprotocol, sqlprotocol, weblogicprotocol

Utilization of these protocol types should be done by using dedicated clients.

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 23](#).

Example

```
import logger
import netutils
import sys
import errorcodes
import errorobject

# Java imports
from java.util import Properties
from com.hp.ucmdb.discovery.common import CollectorsConstants
from appilog.common.system.types.vectors import ObjectStateHolderVector
from com.hp.ucmdb.discovery.library.clients import ClientsConsts
from com.hp.ucmdb.discovery.library.scope import DomainScopeManager

TRUE = 1
FALSE = 0

def mainFunction(Framework, isClient, ip_address = None):
    _vector = ObjectStateHolderVector()
    errStr = ''
    ip_domain = Framework.getDestinationAttribute('ip_domain')
    # Get the Trigger CI data values
    ip_address = Framework.getDestinationAttribute('ip_address')

    if (ip_domain == None):
        ip_domain = DomainScopeManager.getDomainByIp(ip_address, None)

    protocols = netutils.getAvailableProtocols(Framework,
        ClientsConsts.SNMP_PROTOCOL_NAME, ip_address, ip_domain)
    if len(protocols) == 0:
        errStr = 'No credentials defined for the triggered ip'
        logger.debug(errStr)
        errObj = errorobject.createError(errorcodes.NO_CREDENTIALS_FOR_
            TRIGGERED_IP, [ClientsConsts.SNMP_PROTOCOL_NAME], errStr)
        return (_vector, errObj)

    connected = 0
    # Go over all the protocol credentials
    for protocol in protocols:
        client = None
        try:
```



```
        try:
            logger.debug('try to get snmp agent for: %s:%s' % (ip_
address, ip_domain))
            if (isClient == TRUE):
                properties = Properties()
                properties.setProperty
(CollectorsConstants.DESTINATION_DATA_IP_ADDRESS, ip_address)
                properties.setProperty
(CollectorsConstants.DESTINATION_DATA_IP_DOMAIN, ip_domain)
                client = Framework.createClient(protocol, properties)
            else:
                properties = Properties()
                properties.setProperty
(CollectorsConstants.DESTINATION_DATA_IP_ADDRESS, ip_address)
                client = Framework.createClient(protocol, properties)
            logger.debug('Running test connection queries')
            testConnection(client)
            Framework.saveState(protocol)
            logger.debug('got snmp agent for: %s:%s' % (ip_address,
ip_domain))
            isMultiOid = client.supportMultiOid()
            logger.debug('snmp server isMultiOid state=%s'
%isMultiOid)

            client.close()
            client = None

        except:
            if client != None:
                client.close()
                client = None
            logger.debugException('Unexpected SNMP_AGENT Exception:')
            lastExceptionStr = str(sys.exc_info()[1]).strip()
        finally:
            if client != None:
                client.close()
                client = None

    return (_vector, error)
```

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.

Note: It is very important to import the Java base exception class as shown in the following example due to the presence of the base exception class in Python with the same name.

```
from java.lang import Exception as JException
try:
    client = Framework.createClient(Framework.getTriggerCIData(BaseClient.CREDENTIALS_ID))
except JException, ex:
    # process java exceptions only
    Framework.reportError('Connection failed')
    logger.debugException(str(ex))
    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.

Troubleshooting Migration from Jython Version 2.1 to 2.5.3

Universal Discovery now uses Jython version 2.5.3. All out-of-the-box scripts have been properly migrated. If you developed your own Jython scripts prior to this upgrade for use by Discovery, you may run into the following issues and have to make the fixes indicated.

Note: You must be an experienced Jython developer to make these changes.

String Formatting

- **Error message:** `TypeError: int argument required`
- **Possible cause:** Using string formatting to decimal integer from string variable containing integer data.

- **Problematic Jython 2.1 code:**

```
variable = "43"
print "%d" % variable
```

- **Correct Jython 2.5.3 code:**

```
variable = "43"
print "%s" % variable
```

or

```
variable = "43"
print "%d" % int(variable)
```

Checking String Type

The code below may not work correctly if input contains unicode strings:

- **Problematic Jython 2.1 code:** `isinstance(unicodeStringVariable, '')`
- **Correct Jython 2.5.3 code:** `isinstance(unicodeStringVariable, basestring)`

The comparison should be done with `basestring` to test whether an object is an instance of `str` or `unicode`.

Non-ASCII character in file

- **Error Message:**
`SyntaxError: Non-ASCII character in file 'x', , but no encoding declared; see http://www.python.org/peps/pep-0263.html for details`
- **Correct Jython 2.5.3 code:** (add this to the first line in the file)

```
# coding: utf-8
```

Import sub-packages

- **Error message:**
`AttributeError: 'module' object has no attribute 'sub_package_name'`
- **Possible cause:** A sub-package is imported without explicitly specifying the name of sub-package in the import statement.
- **Problematic Jython 2.1 code:**

```
import a
print dir(a.b)
```

The sub-package is not explicitly imported.

- **Correct Jython 2.5.3 code:**

```
import a.b

or

from a import b
```

Iterator Changes

Starting from Jython 2.2, the `__iter__` method is used to loop over a collection in the scope of a **for-in** block. The iterator should implement the `next` method, returning an appropriate element or throw the

StopIteration error if it reached the end of the collection. If the `__iter__` method is not implemented, the `getitem` method is used instead.

Raising Exceptions

- **Jython 2.1 method for raising exceptions is obsolete:**
`raise Exception, 'Failed getting contents of file'`
- **Recommended Jython 2.5.3 method for raising exceptions:**
`raise Exception('Failed getting contents of file')`

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.

This section includes:

- ["Add Support for a New Language" below](#)
- ["Change the Default Language" on the next page](#)
- ["Determine the Character Set for Encoding" on page 54](#)
- ["Define a New Job to Operate With Localized Data" on page 54](#)
- ["Decode Commands Without a Keyword" on page 55](#)
- ["Work with Resource Bundles" on page 56](#)
- ["API Reference" on page 57](#)

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 \(*.properties Files\)" below](#)
- ["Declare and Register the Language Object" on the next page](#)

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
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 56](#).

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 *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 57](#). For details on the Class Locale object, see <http://java.sun.com/j2se/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 52](#).

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, use the API for the **ShellUtils** class. For details, see ["RTSM Web Service API Overview" on page 259](#).

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 the step about retrieving a keyword from the resource bundle in ["Define a New Job to Operate With Localized Data" on the previous page](#).

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:

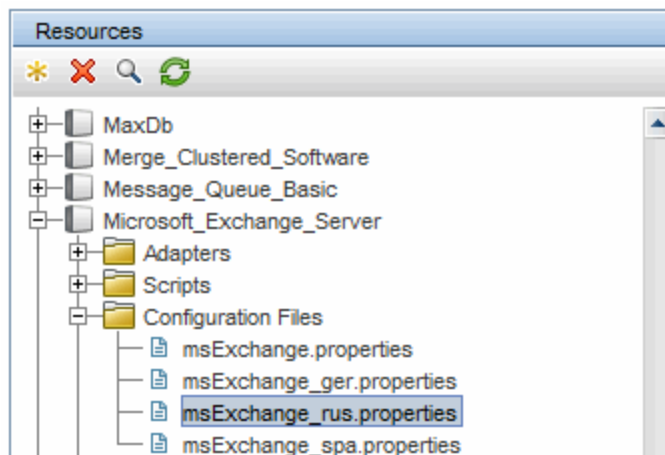
1. Invoke the **executeCommandAndDecode** method once.
2. Obtain the most recently used character set name through the **getCharsetName** method.
3. Make **shellUtils** use this character set by default, by invoking the **useCharset** method on the **ShellUtils** object.
4. Invoke the **execCmd** method of **ShellUtils** one or more times. The output is returned with the character set specified in the previous step. 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 *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 for a resource bundle to contain data for different languages, in DFM each resource bundle contains data for one language only.)

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:

- ["The Language Class" below](#)
- ["The executeCommandAndDecode Method" on the next page](#)
- ["The getCharsetName Method" on the next page](#)
- ["The useCharset Method" on the next page](#)
- ["The getLanguageBundle Method" on page 59](#)
- ["The osLanguage Field" on page 59](#)

The 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.

Name	Description
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).

The 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 <code>sudo</code> should be used (relevant only for UNIX machine clients).
language	Enables specifying the language directly instead of automatically detecting a language.

The getCharsetName Method

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

The 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 ["The getCharsetName Method" above](#).

The 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 <code>ShellUtils.osLanguage</code> should be passed here.
framework	The Framework, common object which is passed to every executable Jython script in DFM.

The osLanguage Field


This field contains an object that represents the language.

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 **Administration > RTSM Administration > Data Flow Management > Universal Discovery**. Right-click the job whose run must be logged and select **Go to Adapter** to open the Adapter Management application.
2. Locate the **Execution Options** pane in the **Adapter Configuration** tab, as shown below.



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

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

	Job name	Trigger CI data
<pre> - <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> </pre>		

The following example shows the message and stacktrace parameters:

Stacktrace	
<pre> - <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"> </pre>	

Jython Libraries and Utilities

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

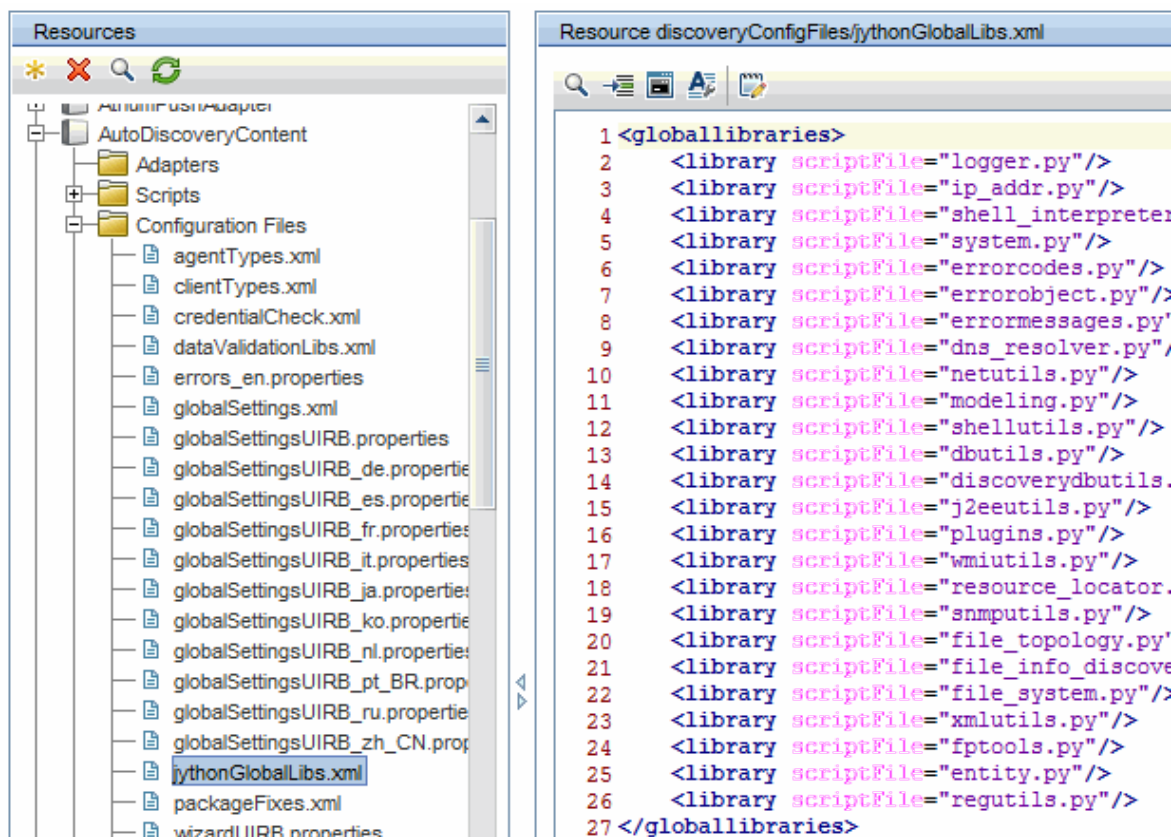
C:\hp\UCMDB\DataFlowProbe\runtime\probeManager\discoveryScripts with the other scripts that are downloaded to the Probe.

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 the next page](#)
- ["modeling.py" on the next page](#)
- ["netutils.py" on page 63](#)
- ["shellutils.py" on page 63](#)

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](#).

```
#####  
#####          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\probeMgr-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 = Framework.createClient(Framework.getTriggerCIData
(BaseClient.CREDENTIALS_ID), Props)
clientShUtils = shellutils.ShellUtils(ttyClient)
if (clientShUtils.isWinOs()):
    logger.debug ('discovering Windows..')
```

Chapter 3: Developing Generic Database Adapters

This chapter includes:

Generic Database Adapter Overview	65
TQL Queries for the Generic Database Adapter	65
Reconciliation	66
Hibernate as JPA Provider	67
Prepare for Adapter Creation	69
Prepare the Adapter Package	74
Configure the Adapter – Minimal Method	77
Configure the Adapter – Advanced Method	82
Implement a Plug-in	87
Deploy the Adapter	90
Edit the Adapter	90
Create an Integration Point	90
Create a View	90
Calculate the Results	91
View the Results	91
View Reports	91
Enable Log Files	91
Use Eclipse to Map Between CIT Attributes and Database Tables	92
Adapter Configuration Files	99
Out-of-the-Box Converters	124
Plug-ins	129
Configuration Examples	129
Adapter Log Files	138
External References	140
Troubleshooting and Limitations	140

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 *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 69](#).
- Reconciliation with the RTSM ID is also supported
- Reconciliation with the Global 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 140](#).

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 69](#) below.

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" below](#)
- ["Associations" on page 69](#)
- ["Usability" on page 69](#)

Examples of Object-Relational Mapping

The following examples describe object-relational mapping:

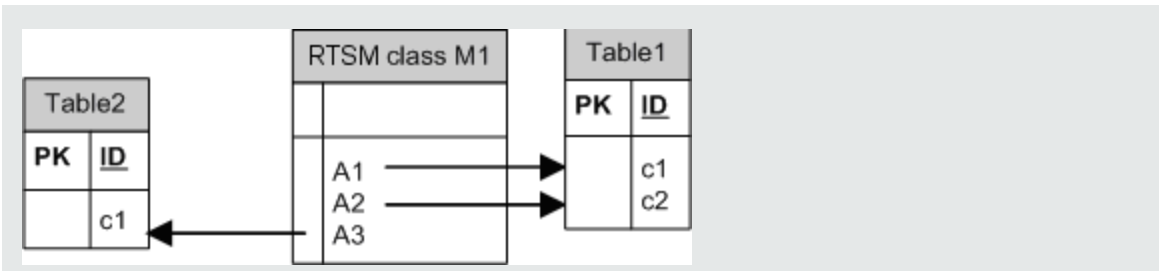
Example of One RTSM Class Mapped to One Database Table:

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.

RTSM class M1	
A1	→
A2	→
A3	→


Table1	
PK	ID
	c1
	c2
	c3

Example of One RTSM Class Mapped to Two Database Tables:



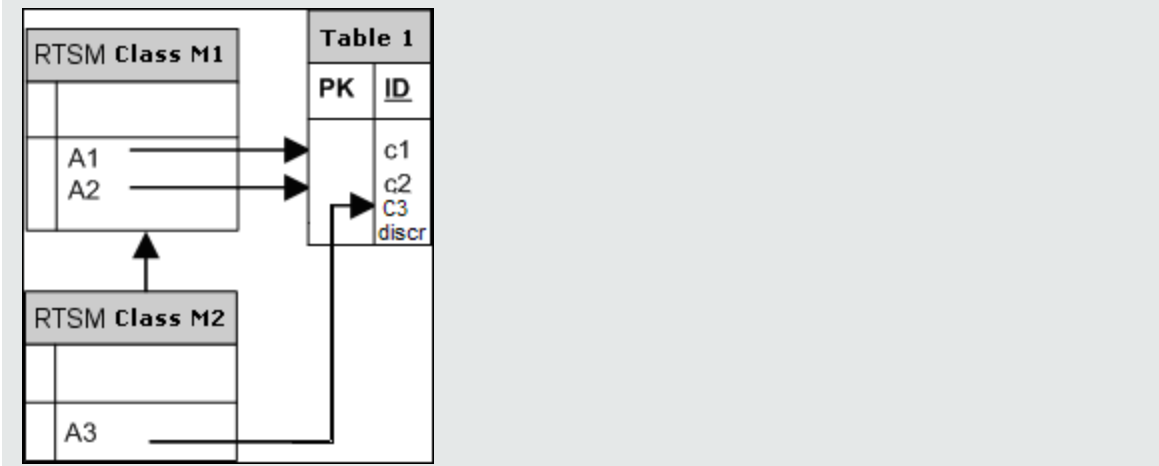
Example of Inheritance:

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.



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 make it easier to define associations.

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.

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 plug-in interfaces.

This task includes the following steps:

- ["Prerequisites" below](#)
- ["Create a CI Type" on page 71](#)
- ["Create a Relationship" on page 71](#)

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 66](#).

ID	NAME	IP_ADDRESS
31	BABA	16.59.33.60

ID	NAME	IP_ADDRESS
33	ext3.devlab.ad	16.59.59.116
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 67](#).

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

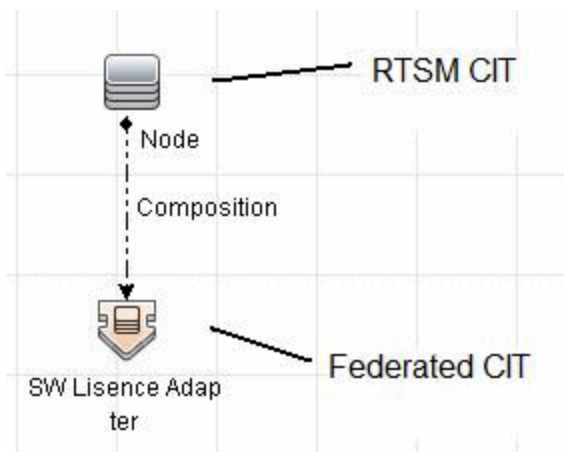
In this step you create a CIT that represents the data in the RDBMS (the external data source).

- In BSM, access the CI Type Manager and create a new CI Type. For details, see [Create a CI Type](#) in the *HP Universal CMDB Modeling Guide*.
- 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.

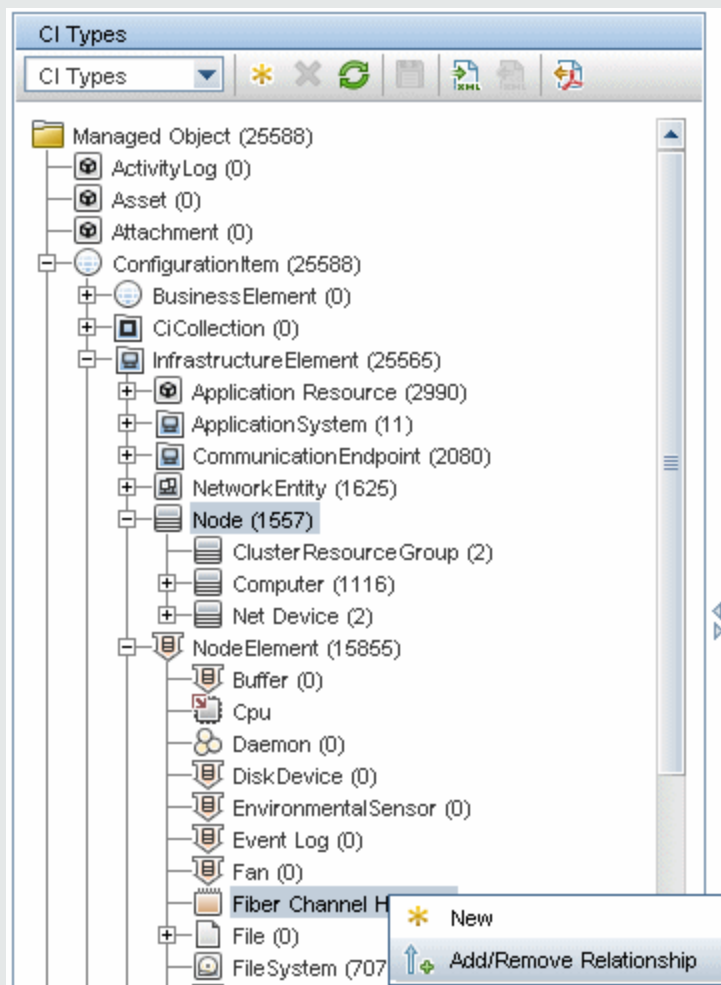
Add appropriate, valid relationships to the new CIT. For details, see [Add/Remove Relationship Dialog Box](#) in the *HP Universal CMDB Modeling Guide*.



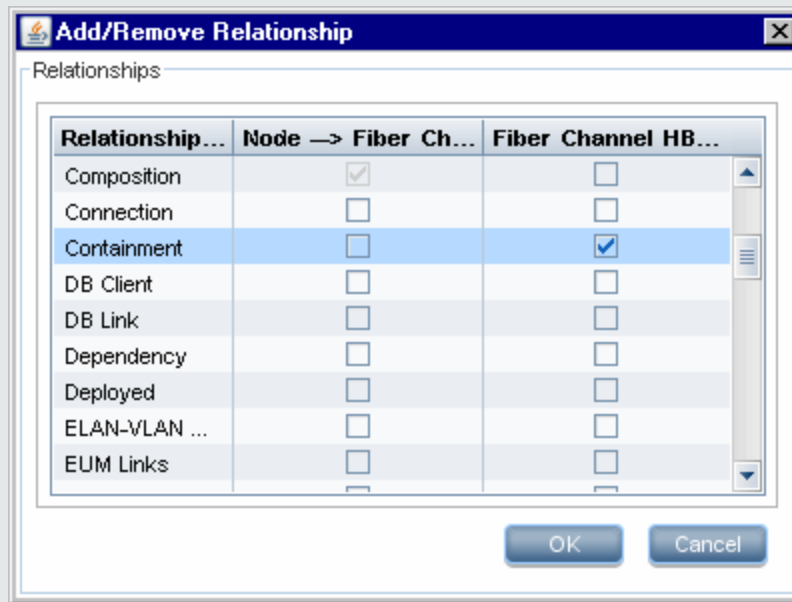
Note: At this stage, you cannot yet view the federated data or populate the external data, as you have not yet defined the method for bringing in the data.

Example of Creating a Containment Relationship:

- a. In the CIT Manager, select the two CITs:



- b. Create a **Containment** relationship between the two CITs:



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 HP Operations Manager i.

When creating a Generic DB Adapter it is not necessary to edit the **changes-source** tag in the **support-replicatioin-data** tag. If the **FcmdbPluginForSyncGetChangesTopology** plug-in is implemented, the changed topology from the last run will be returned. If the plug-in is not implemented, the full topology will be returned and the auto-delete will be performed according to the returned CIs.

For details about populating the RTSM with data, see ["Integration Studio Page"](#) in the *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;B0" 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 CMDVRTSM-->
  <!--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 the previous step.

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 ["Configure the Adapter – Minimal Method" below](#).
6. Create a *.zip file with the same name as you gave to the **adapter id** attribute, as described in the step ["Edit the adapter XML file:" on page 74](#).

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**.

Configure the Adapter – Minimal Method

The simplified (minimal) method is a method for creating the **simplifiedConfiguration.xml** mapping file that is used by the adapter. This method enables a basic population or federation of a single CIT.

The instructions provided in this section describes a method of mapping the class model for certain CI Types in the RTSM to an RDBMS.

All of the configuration files mentioned in this section are located in the **db-adapter.zip** package in the **<HP BSM root directory>\odb\content\adapters** folder that you extracted in ["Prepare the Adapter Package" on page 74](#).

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

Configure the adapter.conf File


To 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**.

Example: Populating a Node and IP Address using the Simplified Method

This example demonstrates populating a **Node** related by a containment link to **IP Address** into the BSM. The RDBMS has a table named **simpleNode** that contains data on the computer name, computer node, and the IP address of the computer.

The content of the **simpleNode** table is shown below:

	host_id	host_name	note	ip_address
	1	Comp1	Test Computer 1	12.33.211.52
	2	Comp2	Test Computer 2	12.33.211.53
	3	Comp3	Test Computer 3	12.33.211.54
	4	Comp4	Test Computer 4	12.33.211.55

The population is performed in three stages, as follows:

1. ["Create the simplifiedConfiguration.xml" below](#)
2. ["Create the TQL" on page 80](#)
3. ["Create an Integration Point" on page 81](#)

Create the `simplifiedConfiguration.xml`

Create the **simplifiedConfiguration.xml** as follows:

1. Create a **cmdb-class** entity as follows:

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

The CI Type is **node** and the RDBMS table name is **simpleNode**.

2. Set the primary-key of the table as follows:

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

This primary key is equivalent to entity id in the **orm.xml** file.

3. Set the **reconciliation-by-two-nodes** rule as follows:

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

This tag defines the relation between the **Node** and the **IpAddress** CI types. The relation type is Containment link. The reconciliation is done by the two connected CI Types. The attribute mapping of the connected node (in this case IpAddress) is defined in the **connected-node** attribute.

4. Add the **or** condition between the reconciliation attributes as follows:

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

This tag defines an OR relationship between the reconciliation attributes, meaning the first reconciliation attribute that is **true** sets the whole reconciliation rule to **true**.

5. Add the following attributes:

```
<attribute cmdb-attribute-name="name" column-name="host_name" ignore-case="true"/>
```

This tag sets a mapping between the **node.name** in the BSM to the column **host_name** in the **simpleNode** table.

Do the same with **data_note** attribute:

```
<attribute cmdb-attribute-name="data_note" column-name="note" ignore-case="true"/>
```

Add the connected node attribute:

```
<connected-node-attribute cmdb-attribute-name="name" column-name="ip_address"/>
```

This tag sets a mapping between the **ip_address.name** to the column **ip_address** in the **simpleNode** table.

6. Close the opened tag by order:

```
</or>
```


```
</reconciliation-by-two-nodes>  
</cmdb-class>
```

The contents of the simplifiedConfiguration.xml file now appear as follows:

```
<?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="simpleNode">  
    <primary-key column-name="host_id"/>  
    <reconciliation-by-two-nodes connected-node-cmdb-class-name="ip_address" cmdb-  
link-type="containment">  
      <or is-ordered="true">  
        <attribute cmdb-attribute-name="name" column-name="host_name" ignore-  
case="true"/>  
        <attribute cmdb-attribute-name="data_note" column-name="note" ignore-  
case="true"/>  
        <connected-node-attribute cmdb-attribute-name="name" column-name="ip_  
address"/>  
      </or>  
    </reconciliation-by-two-nodes>  
  </cmdb-class>  
</generic-db-adapter-config>
```

Create the TQL

The TQL is a **node** connected by a containment link to **ip_address**. The node should be marked as **root**, as shown below.



To create the TQL:

1. Go to **Modeling > Modeling Studio**.
2. Click the **New** button and create a new query.
3. Go to the CI Types tab and drag the Node CI Type and IPAddress CI Type to the TQL screen.
4. Connect **Node** and **IPAddress** with a Containment relationship.
5. Right-click on the **Node** element and choose Query Node Properties.
6. Change **Element name** to **Root**.
7. Go to the **Element Layout** tab. Select **Specific Attributes** as the Attributes condition. Choose **Name** and **Note** from the Available Attributes window and move them to the Specific Attributes window.
8. Right-click on the **IPAddress** element and choose Query Node Properties.
9. Go to the **Element Layout** tab. Select **Specific Attributes** as the Attributes condition. Choose **Name** from the Available Attributes window and move it to the Specific Attributes window.
10. Save the TQL.

Create an Integration Point

Create the Integration Point as follows:

1. Go to **Data Flow Management > Integration Studio**, and click the **New Integration Point** button.
2. Insert the details of the Integration Point and click **OK**.
3. In the Population tab, select the **New Integration Job** button, and add the TQL previously created.
4. Save the Integration Point and click the **Run Full Synchronization** button.

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 when preparing the adapter package. For details, see ["Prepare the Adapter Package" on page 74](#).

This task includes the following steps:

- ["Configure the orm.xml File" below](#)
- ["Configure the reconciliation_rules.txt File " on page 86](#)

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 108](#).

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 71](#):

```
<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:

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

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_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 ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118.](#)

Implement a Plug-in

This task describes how to implement and deploy a Generic DB Adapter with plug-ins.

Note: Before writing a plug-in for an adapter, make sure you have completed all the necessary steps in ["Prepare the Adapter Package" on page 74.](#)

1. Option 1 – Write a Java based plug-in

- a. 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 plug-in 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\APIs\DBAdapterFramework_JavaAPI\index.html**
- **C:\hp\UCMDB\UCMDBServer\deploy\ucmdb-docs\docs\eng\APIs\Federation_JavaAPI\index.html**

- b. Write a Java class implementing the plug-in'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 plug-in:

Plug-in Type	Interface Name	Method
Synchronize Full Topology	FcmdbPluginForSyncGetFullTopology	getFullTopology
Synchronize Changes	FcmdbPluginForSyncGetChangesTopology	getChangesTopology
Synchronize Layout	FcmdbPluginForSyncGetLayout	getLayout
Retrieve Supported Queries	FcmdbPluginForSyncGetSupportedQueries	getSupportedQueries
Alter TQL query definition and results	FcmdbPluginGetTopologyCmdFormat	getTopologyCmdFormat
Alter layout request for CIs	FcmdbPluginGetCIsLayout	getCIsLayout
Alter layout request for links	FcmdbPluginGetRelationsLayout	getRelationsLayout
Push Back IDs	FcmdbPluginPushBackIds	getPushBackIdsSQL

The plug-in'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>
```


- c. 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.
 - d. 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.
2. Option 2 – Write a Groovy based plug-in
 - a. Create a Groovy code file (MyPlugin.groovy) in the Adapter Management Menu, under the adapter package configuration files.
 - b. In the Groovy class, implement the appropriate interfaces. The interfaces are defined in the db-interfaces.jar file, see the table above.
3. 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 tqls
names
[getSupportedQueries]
# internal not mandatory plugin to change tq1 definition and tq1 result
[getTopologyCmdFormat]
# internal not mandatory plugin to change layout request and CIs result
[getCisLayout]
# internal not mandatory plugin to change layout request and relations
result
[getRelationsLayout]
# internal not mandatory plugin to change action on pushBackIds
[pushBackIds]
```

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 plug-in class associated, or the fully qualified name of your plug-in class can be listed.

4. 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.

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.

Edit the Adapter

Once you have created and deployed the adapter, you can then edit it within BSM. For details, see "[Adapter Management](#)" in the *Data Flow Management Guide*.

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 (**Administration > 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 *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 *Data Flow Management Guide*.


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 (**Administration > RTSM Administration > Modeling > Modeling Studio**).
2. Create a view. For details, see [Create a Pattern View](#) in the Modeling Guide.

Calculate the Results

In this step you check the results.

1. In BSM, access the Modeling Studio (**Administration > 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.

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 (**Administration > RTSM Administration > Modeling > IT Universe Manager**).
2. Select a CI. The Properties tab displays the results of the federation.

View Reports

In this step you view Topology reports. For details, see [Topology Reports Overview](#) in the Modeling Guide.

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 138](#).

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 plug-in, 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 plug-in 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 the next page](#)
- ["Installation" on the next page](#)
- ["Prepare the Work Environment" on the next page](#)
- ["Create an Adapter" on page 94](#)
- ["Configure the RTSM Plug-in" on page 94](#)
- ["Import the RTSM Class Model" on page 95](#)
- ["Build the ORM File – Map BSM Classes to Database Tables" on page 95](#)
- ["Map IDs" on page 95](#)
- ["Map Attributes" on page 96](#)
- ["Map a Valid Link" on page 96](#)
- ["Build the ORM File – Use Secondary Tables" on page 97](#)

- "Define a Secondary Table" on page 97
- "Map an Attribute to a Secondary Table" on page 97
- "Use an Existing ORM File as a Base" on page 97
- "Importing an Existing ORM File from an Adapter" on page 98
- "Check the Correctness of the orm.xml File – Built-in Correctness Check" on page 98
- "Create a New Integration Point" on page 98
- "Deploy the ORM File to the RTSM" on page 98
- "Run a Sample TQL Query" on page 99

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 26](#).

5. Configure the RTSM Plug-in

- 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 RTSM 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 above in the step ["Build the ORM File – Map BSM Classes to Database Tables" on the previous page](#) 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 ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118](#).

- a. Map the ID attributes of the link class as described in ["Map IDs" on the previous page](#). 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 above in ["Map Attributes" on the previous page](#).
- 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 above in the step ["Build the ORM File – Map BSM Classes to Database Tables" on page 95](#). 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 Plug-in can be used to edit its ORM file graphically. Import the **orm.xml** file into Eclipse, edit it using the plug-in 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 RTSM 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 below in ["Deploy the ORM File to the RTSM" below](#).

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

The Eclipse JPA plug-in 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 *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 ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118](#).

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 *Modeling Guide*.
- b. Create an integration point using the adapter that you created in the step "[Create a New Integration Point](#)" on the previous page. For details, see [New Integration Point/Edit Integration Point Dialog Box](#) in the *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 plug-in, use this sample query name in the Settings dialog box. For details, see the step above "[Configure the RTSM Plug-in](#)" on page 94.
- 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.

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 describes the following configuration files:

- ["The adapter.conf File" on page 101](#)
- ["The simplifiedConfiguration.xml File" on page 102](#)
- ["The orm.xml File" on page 104](#)
- ["The reconciliation_types.txt File" on page 118](#)
- ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118](#)
- ["The transformations.txt File" on page 120](#)
- ["The discriminator.properties File" on page 121](#)
- ["The replication_config.txt File" on page 122](#)
- ["The fixed_values.txt File" on page 122](#)
- ["Persistence.xml File" on page 122](#)

General Configuration

- **adapter.conf.** The adapter configuration file. For details, see ["The adapter.conf File" on the next page](#).

Simple Configuration

- **simplifiedConfiguration.xml.** Configuration file that replaces **orm.xml**, **transformations.txt**, and **reconciliation_rules.txt** with less capabilities. For details, see ["The simplifiedConfiguration.xml File" on page 102](#).

Advanced Configuration

- **orm.xml.** The object-relational mapping file in which you map between RTSM CITs and database tables. For details, see ["The orm.xml File" on page 104](#).
- **reconciliation_rules.txt.** Contains the reconciliation rules. For details, see ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118](#).
- **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 ["The transformations.txt File" on page 120](#).
- **Discriminator.properties.** This file maps each supported CI type to a comma-separated list of possible corresponding values. For details, see ["The discriminator.properties File" on page 121](#).
- **Replication_config.txt.** This file contains a comma-separated list of CI and relationship types whose property conditions are supported by the replication plug-in. For details, see ["The replication_config.txt File" on page 122](#).
- **Fixed_values.txt.** This file enables you to configure fixed values for specific attributes of certain CITs. For details, see ["The fixed_values.txt File" on page 122](#).

Hibernate Configuration

- **persistence.xml.** Used to override out-of-the-box Hibernate configurations. For details, see ["Persistence.xml File" on page 122](#).

Enabling Temporary Table Support for the Adapter

Enabling temporary tables allows the adapter to work more efficiently with the remote database, thus reducing stress on the database and network and also enhancing performance.

To enable temporary table support in General Database Adapter, the following conditions must be met:

- The credentials given to connect to the database, include permission to create, modify, and delete temporary tables.
- Configure the following settings in the adapter.conf configuration file:

temp.tables.enabled=true

performance.enable.single.sql=true

Note: Temporary tables are only supported for Microsoft SQL and Oracle.

The adapter.conf File

This file contains the following settings:

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

Note: Using 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, you can either map the **cmdb_id** to a **string** or **bytes/raw** column type by changing the **META-INF/ adapter.conf** property.
- **performance.enable.single.sql=true.** This is an optional parameter. If it does not appear in the file, its default value is **true**. When **true**, the Generic Database Adapter tries to generate a single SQL statement for each query being executed (either for population or a federated query). Using a single SQL statement improves the performance and memory consumption of the Generic Database Adapter. When **false**, the Generic Database Adapter generates multiple SQL statements, which may take longer and consume more memory than a single one. Even when this attribute is set to **true**, the adapter does not generate a single SQL statement in the following scenarios:

- The database the adapter connects to is not on an Oracle or SQL Server.
- The TQL being executed contains a cardinality condition other than 0..* and 1..* (for example, if there is a cardinality condition like 2..* or 0..2).
- **in.expression.size.limit=950** (default). This parameter splits the 'IN' expression of the executed SQL, when the size limit of the list of arguments is reached.
- **stringlist.delimiter.of.<CIT Name>.<Attribute Name>=<delimiter>**. In order to map a string list attribute to a database column in the generic database adapter, the attribute needs to be mapped to a string column that contains a list of concatenated values. For example, to map the attribute **policy_category** with the CI Type **policy**, and the string column contains a list of values: value1##value2##value3 (that define a list of 3 values value1, value2, value3), use the setting: **stringlist.delimiter.of.policy.policy_category=##**.
- **temp.tables.enabled=true**. Enables using temporary tables to improve performance. Only available when **performance.enable.single.sql** is enabled (only supported in Microsoft SQL and Oracle). Certain permissions in the database server may be required.
- **temp.tables.min.value=50**. Defines the number of condition values (or IDs) that are needed to use temporary tables.

The simplifiedConfiguration.xml File

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

This section includes the following topics:

- ["The simplifiedConfiguration.xml File Template" below](#)
- ["Limitations" on page 104](#)

The 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 RTSM 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.fcmbd.adapters.dbAdapter.dal.transform.
FcmbdDalTransformerFromExternalDB
- com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.
FcmbdDalTransformerToExternalDB

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.impl. GenericEnumTransformer
(generic-enum-transformer-example.xml)" to-CMDB-onverter="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.

The 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 **Administration > RTSM Administration > Adapter Management > db-adapter > Configuration files**.

This section includes the following topics:

- ["The orm.xml File Template" on the next page](#)
- ["Multiple ORM files" on page 108](#)
- ["Naming Conventions" on page 108](#)

- ["Using Inline SQL Statements Instead of Table Names" on page 108](#)
- ["The orm.xml Schema" on page 109](#)
- ["Example of Creating the orm.xml File" on page 114](#)
- ["Configuring a Specific orm.xml for each Remote Product Version" on page 117](#)

The 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>
```

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_class_name]">
        <column updatable="false" insertable="false" name="[column_name]
" />
    </basic>
    <basic name="[CMDB_class_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>.

The 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 mainly explains the specific behavior of the standard Java Persistence API for the Generic Database Adapter.

Element Name and Path	Description	Attributes
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 Name and Path	Description	Attributes
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> .	<ol style="list-style-type: none"> Name: name Description: The name of the UCMDB CI type to which this entity is mapped. If this 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. Is required?: Optional Type: String Name: class 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>). Is required?: Required Type: String

Element Name and Path	Description	Attributes
table (entity-mappings>entity)	This element defines the primary table of the DB entity. Can only appear once. Required.	Name: name 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 valid SELECT statement. If this is a SELECT statement, it must be encapsulated with parentheses. Is required?: Required Type: String
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.	Name: name 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 valid SELECT statement. If this is a SELECT statement, it must be encapsulated with parentheses. Is required?: Required Type: String
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.	Name: name 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. Is required?: Optional Type: String

Element Name and Path	Description	Attributes
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.	Name: strategy Description: Defines the way the inheritance is implemented in your DB. Is required?: Required Type: One of the following values: <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.	Name: name Description: The name of the discriminator column. Is required?: Required Type: String
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.	
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).	Name: name 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 UCMDDB. Is required?: Required Type: String

Element Name and Path	Description	Attributes
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.	Name: name 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. Is required?: Required Type: String
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.	<ol style="list-style-type: none"> Name: name Description: The name of the field. Is required?: Required Type: String Name: table 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. Is required: Optional Type: String
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.	<ol style="list-style-type: none"> Name: name Description: Which of the two ends this field represents. Is required?: Required Type: Either <code>end1</code> or <code>end2</code> Name: target-entity Description: The name of the entity to which the end refers. Is required?: Required Type: One of the entity names defined in the entity mapping document.

Element Name and Path	Description	Attributes
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.	<ol style="list-style-type: none"> 1. Name: name Description: The name of the field in the current table that will be used to perform the one-to-one join. Is required?: Required Type: String 2. Name: name 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. Is required?: Optional Type: String

Example of Creating the orm.xml File

The example presented here demonstrates how to create the **orm.xml file**. In this example, SQL tables in a remote database are mapped to CI types in the UCMDB.

Given tables with the following format in the remote database, populate the **Hosts** table with nodes, the **IP_Addresses** table with IP Addresses, and create links between the nodes and IP Addresses as follows:

Hosts Table

host_name	host_id
Test1	1
Test2	2
Test3	3

IP_Addresses Table

ip_address	ip_id
10.1.1.1	1
10.2.2.2	2

ip_address	ip_id
10.3.3.2	3
10.4.4.4	4

Host_IP_Link Table (links between Nodes and IP Addresses)

host_id	ip_id
1	1
2	2
2	3
3	4

The primary key for the **Hosts** table is the **host_id** field and the primary key in the **IP_Addresses Table** is the **ip_id** field. In the **Host_IP_Link** table the **host_id** and the **ip_id** are foreign keys from the **Hosts Table** and **IP_Addresses Table**.

According to the tables above, create the **orm.xml** file according to the following steps. The entities used in this example are **node**, **ip_address**, and **node_containment_ip_address**

1. Create the **node** entity by mapping the **host_id** from the **Hosts** table as follows:

```
<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
/or_1_0.xsd">
  <description>test_integration</description>
  <package>generic_db_adapter</package>
  <entity class="generic_db_adapter.node">
    <table name="Hosts"/>
    <attributes>
      <id name="id1">
        <column updatable="false" insertable="false" name=
          "host_id"/>
        <generated-value strategy="TABLE"/>
      </id>
      <basic name="name">
```

```
        <column updatable="false" insertable="false" name="
            host_name"/>

    </basic>

</attributes>

</entity>
```

The entity **class** must be a CI Type that already exists in the UCMDB. The table **name** is the table within the database that contains both an ID and the Host information. The ID attribute is required to identify specific hosts and is used later in the mapping. In this example, the **name** attribute of this entity is populated with the **host_name** column in the hosts table.

2. For the next entity, map the IP Addresses from the Interfaces table:

```
<entity name="ip_address" class="generic_db_adapter.ip_address">
    <table name="IP_Addresses"/>
    <attributes>
        <id name="id1">
            <column insertable="false" updatable="false" name="ip_id"/>
            <generated-value strategy="TABLE"/>
        </id>
        <basic name="name">
            <column updatable="false" insertable="false" name="ip_address"/>
        </basic>
    </attributes>
</entity>
```

3. Next the link between the Node and the IP Address must be created by means of the mapping table, and reference the **ip_id** field (though it could reference both the **host_id** and **ip_id** fields if desired).

```
<entity name="node_containment_ip_address"
    class="generic_db_adapter.node_containment_ip_address">
    <table name="Host_IP_Link"/>
    <attributes>
        <id name="id1">
            <column updatable="false" insertable="false" name="ip_id"/>
```

```
<generated-value strategy="TABLE"/>
</id>
<many-to-one target-entity="node" name="end1">
  <join-column name="host_id"/>
</many-to-one>
<one-to-one target-entity="ip_address" name="end2">
  <join-column name="ip_id"/>
</one-to-one>
</attributes>
</entity>
```

The entity name for the container has the format: [end1 CIT]_[link CIT]_[end2 CIT]. So for this example, since the link CI Type is **containment**, the entity name for the container is: **node_containment_ip_address** and the entity class is **generic_db_adapter.node_containment_ip_address**. The ID is required in this block of code, and while this example works with a single ID of the Interface, both columns could reference id1 and id2. The code for that would be:

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

The two ends of this link are 'many-to-one' and 'one-to-one', meaning each IP Address will be linked to 1 node, but a node may be linked to many IP Addresses. The columns included are from the Links table and reference the Hosts and Interfaces tables.

Configuring a Specific orm.xml for each Remote Product Version

It is possible to configure a specific **orm.xml** file so that the adapter uses a specific **orm.xml** for a given remote product version. For example, if the remote data store has two product versions x and y, for each version there can be a different mapping of entities.

To configure a specific orm.xml file per remote product version:

1. Add a parameter to the **adapter.xml** file called **version** and specify possible version values as **valid-values**.
2. In the adapter package, under the META-INF folder, create a folder called **VersionOrm**.
3. In the **VersionOrm** folder, create an **orm.xml** file for each specific version. The filename should contain the version prefix. For example, if the version is called **x**, the filename should be **x_orm.xml**.

Note: The **orm.xml** file in the META-INF folder is loaded for any remote product version, regardless of whether you create a specific **orm.xml** file for a remote product version. It can have entities that are mapped in the same manner for all versions.

The reconciliation_types.txt File

As of UCMDB 10.00, the **reconciliation_types.txt** file is no longer relevant. Any CIT can be used for reconciliation. The federation engine automatically executes the mapping.

The 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), 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	
PK	node_id
	name

DB IP_Address	
PK	ip_id
	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.

The transformations.txt File

This file contains all the converter definitions.

The format is that each line contains a new definition.

The transformations.txt File Template

```
entity[[CMDB_class_name]] attribute[[CMDB_attribute_name]] to_DB_class
[com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.
transform.impl.GenericEnumTransformer(generic-enum-transformer-example.xml)]
from_DB_class[com.mercury.topaz.fcmbd.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.fcmbd.adapters.dbAdapter.dal.transform.FcmbdDalTransformerToExternalDB`. 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.fcmbd.adapters.dbAdapter.dal.transform.FcmbdDalTransformerFromExternalDB` 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 124](#).

The 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**.

The replication_config.txt File

This file contains a comma-separated list of CI and relationship types whose property conditions are supported by the replication plug-in. For details, see ["Plug-ins" on page 129](#).

The 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>, ...
}]
```

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 SQL 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>
```

Connect to Database Using NT Authentication

It is possible to connect to an MS SQL Server that requires NT authentication. To do so a driver that can parse domain is needed (i.e jTDS JDBC Driver).

The authentication is done according to the given parameters (domain, username, password), and not with the current running process NT credentials.

1. In the **persistence.xml** edit the following properties as follows:

```
<!--The driver class name/-->
<property name="hibernate.connection.driver_class" value="net.sourceforge.jtds.jdbc.Driver"/>
<property name="hibernate.connection.url" value="jdbc:jtds:sqlserver://[host name]:
```

```
[port];DatabaseName=[database name];domain=[the domain]"/>
<!--DB login credentials"/-->
<property name="hibernate.connection.username" value="[username]"/>
<property name="hibernate.connection.password" value="[password]"/>
```

2. Place the JDBC driver file under: **<probe installation folder>\lib**.
3. Restart the probe.

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:

- ["Out-of-the-Box Converters" above](#)
- ["The SuffixTransformer Converter" on page 127](#)
- ["The PrefixTransformer Converter" on page 127](#)
- ["The BytesToStringTransformer Converter" on page 128](#)
- ["The StringDelimitedListTransformer Converter" on page 128](#)
- ["The Custom Converter" on page 128](#)

The 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:attribute name="cmdb-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:restriction>
      </xs:simpleType>
    </xs:attribute>
  </xs:complexType>
</xs:element>
```

```
        <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="cmdb-value" type="xs:string" use="required"/>
    <xs:attribute name="external-db-value" type="xs:string" use="required"/>
    <xs:attribute name="is-cmdb-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>
```

Example of Converting an External or RTSM Value to a Null Value:

In this example, a value of **NNN** in the remote database is transformed into a null value in the RTSM database.

```
<value cmdb-value="null" is-cmdb-value-null="true" external-db-value="NNN"/>
```

In this example, the value **000** in the RTSM is transformed into a null value in the remote database.

```
<value cmdb-value="000" external-db-value="null" is-db-value-null="true"/>
```

The 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.

The 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.fcmbd.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.fcmbd.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.

The 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.

The StringDelimitedListTransformer Converter

This converter is used to transform a single string list to an integer/string list in the RTSM.

The converter is: **com.mercury.topaz.fcldb.adapters.dbAdapter.dal.transform.impl.StringDelimitedListTransformer.**

The Custom Converter

It is possible to write your own custom converter (transformer) from scratch. This enables you to create any converter required for your needs.

There are two ways to write a custom converter:

1. Write a compiled Java converter
 - a. Create a Java Project in a Java IDE (such as Eclipse, IntelliJ, or Netbeans).
 - b. Add the federation-api.jar and db-interfaces.jar to your class path.
 - c. Create a Java class that implements the following interfaces (from **db-interfaces.jar**):
 - FcldbDalTransformerFromExternalDB
 - FcldbDalTransformerValuesToExternalDB
 - FcldbDalTransformerInit
 - d. Compile the project and create a jar file.
 - e. Place the jar file in the adapter's package (under adapterCode\<Adapter ID>)
 - f. Deploy the package.
 - g. Add the new converter class name to the **transformations.txt** file.
2. Write a Groovy (script based) converter

An example is found in the original GDBA package, **GroovyExampleTransformer.groovy**.

- a. Create a Groovy file in the adapter's package (under adapterCode\<Adapter ID>). You can do this directly using the Adapter Management menu.
- b. Create a Groovy class that implements the following interfaces (from **db-interfaces.jar**):
 - FcmdbDalTransformerFromExternalDB
 - FcmdbDalTransformerValuesToExternalDB
 - FcmdbDalTransformerInit
- c. Add the new converter Groovy class name to the **transformations.txt** file accordingly.

Note: Groovy is a scripting language that extends Java. Regular Java code is valid Groovy code as well.

Plug-ins

The generic database adapter supports the following plug-ins:

- An optional plug-in 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 plug-in for synchronizing layout.
- An optional plug-in to retrieve supported queries for synchronization. If this plugin is not defined, all TQL names are returned.
- An internal, optional plug-in to change the TQL definition and TQL result.
- An internal, optional plug-in to change a layout request and CIs result.
- An internal, optional plug-in to change a layout request and relationships result.
- An internal, optional plug-in to change the action of push Back IDs.

For details about implementing and deploying plug-ins, see ["Implement a Plug-in" on page 87](#).

Configuration Examples

This section gives examples of configurations.

This section includes the following topics:

- ["Use Case" below](#)
- ["Single Node Reconciliation" below](#)
- ["Two Node Reconciliation" on page 133](#)
- ["Using a Primary Key that Contains More Than One Column" on page 135](#)
- ["Using Transformations" on page 137](#)

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

The orm.xml File

Pay attention to the addition of the relationship mapping. For details, see the definition section in ["The orm.xml File" on page 104](#).

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"/>
            </id>
        </attributes>
    </entity>
</entity-mappings>
```

```

        <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>
```

The reconciliation_rules.txt File

For details, see ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118.](#)

multinode[node] expression[node.name]

The 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

The orm.xml File

Since the reconciliation expression is not defined in this file, the same version should be used for any reconciliation expression.

The reconciliation_rules.txt File

For details, see ["The reconciliation_rules.txt File \(for backwards compatibility\)" on page 118.](#)

```
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]
```

The 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

The 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 ["The orm.xml File" on page 104.](#)

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>
  </attributes>
</entity>

<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"
/>
      <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>
```



```
        </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.impl.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.

The 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.
GenericEnumTransformer(generic-enum-transformer-example.xml)] from_DB_class
[com.mercury.topaz.fcmbd.adapters.dbAdapter.dal.transform.impl.
GenericEnumTransformer(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" below](#)
- ["Log Locations" below](#)

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\fcmbd.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 **fcmbd.gdba.dal.sql.log** to **INFO**, change:

```
log4j.category.fcmbd.gdba.dal.SQL=${loglevel},fcmbd.gdba.dal.SQL.appender
```

to:

```
log4j.category.fcmbd.gdba.dal.SQL=INFO,fcmbd.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 67](#).

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

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 RTSMid 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.

Chapter 4: Developing Java Adapters

This chapter includes:

Federation Framework Overview	142
Adapter and Mapping Interaction with the Federation Framework	147
Federation Framework for Federated TQL Queries	147
Interactions between the Federation Framework, Server, Adapter, and Mapping Engine	149
Federation Framework Flow for Population	158
Adapter Interfaces	159
Debug Adapter Resources	161
Add an Adapter for a New External Data Source	161
Create a Sample Adapter	169
XML Configuration Tags and Properties	170
The DataAdapterEnvironment Interface	172

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.

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" below](#)
- ["Data Push" on the next page](#)
- ["Population" on page 145](#)


Federation on the Fly

Federated TQL queries enables data retrieval from any external data repository without replicating its data.

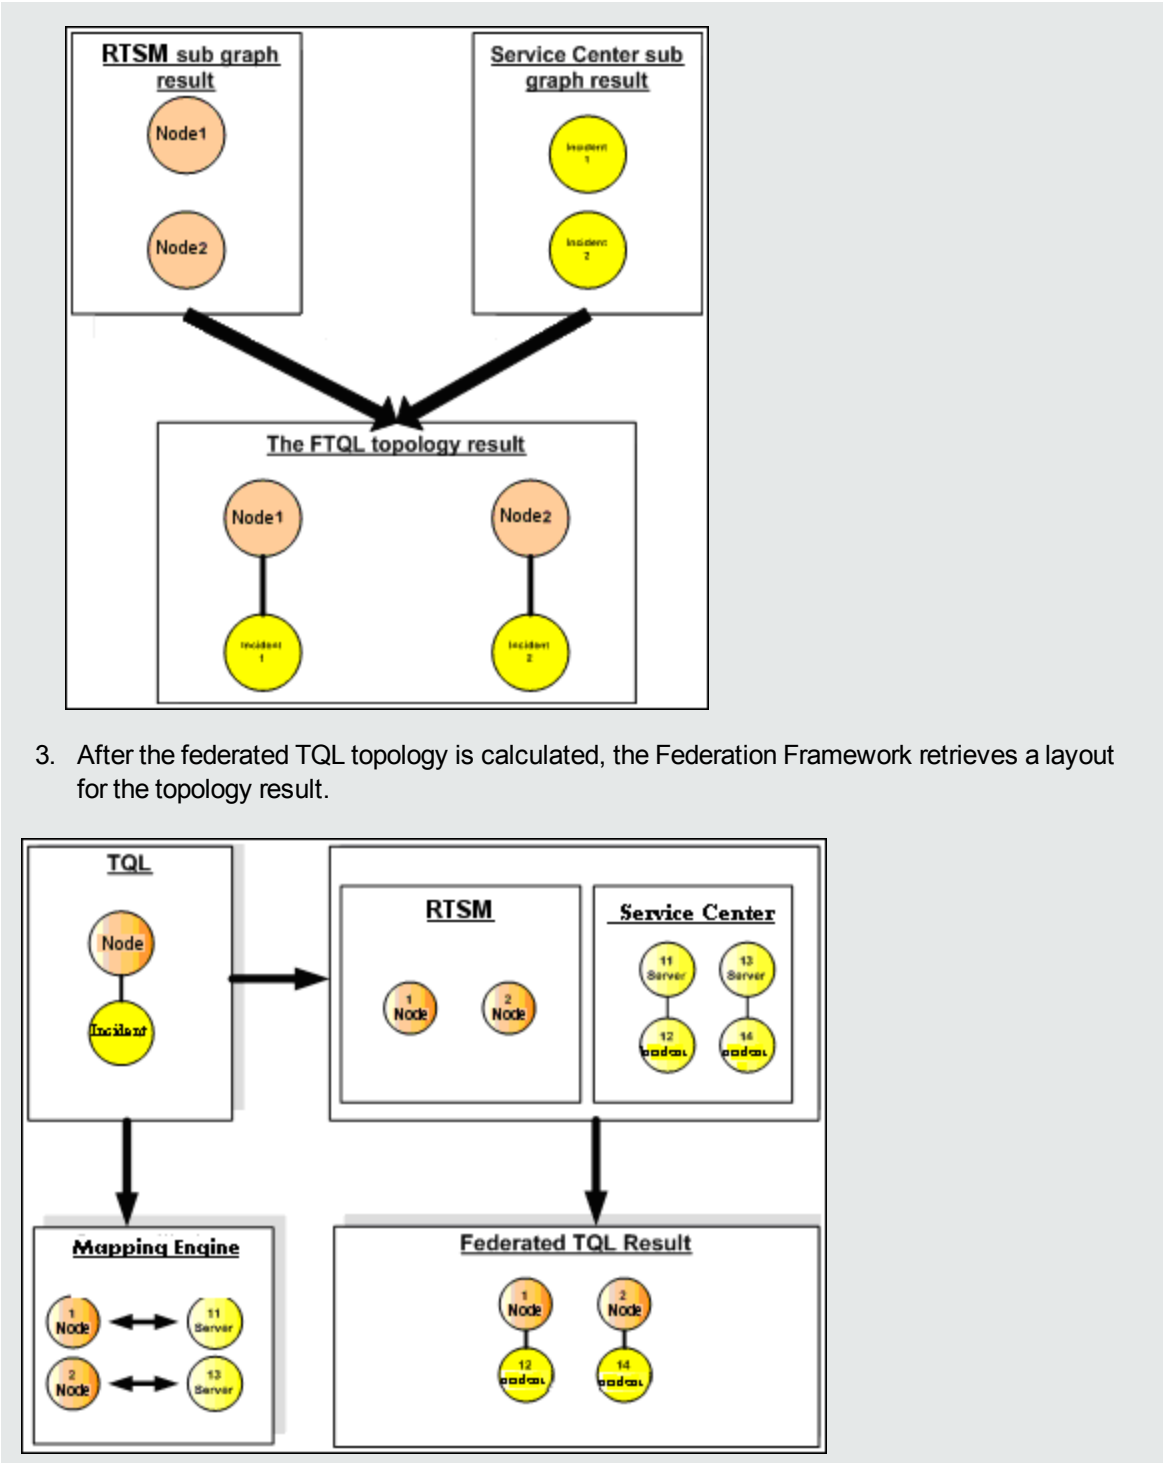
A federated TQL query uses adapters that represent external data repositories, to create appropriate external relationships between CIs from different external data repositories and the BSM CIs.

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.



3. 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:

1. Retrieving the topology result with signatures from the source data repository.
2. Comparing the new results with the previous results.
3. Retrieving a full layout (that is, all CI properties) of CIs and relationships, for changed results only.
4. 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.

Instance-Based Population Flow

If the adapter is defined to support an instance-based flow (by means of the **<instance-based-data>** tag, as described in ["XML Configuration Tags and Properties" on page 170](#)), the population engine automatically finds removed CIs inside the instance and removes them from the UCMDB (assuming deletion is allowed for the specific population job). Each instance must have a Root CI, marked in the TQL definition with the name **Root**. Each time a root CI is passed, its entire instance (all the CIs connected to it) are compared to the last time it was sent to the UCMDB, and any CIs that were connected to the root but are now not connected to it are deleted from the UCMDB. For the adapter to correctly support instance-based flow, any change to any CI or attribute in the entire instance must trigger a resend of the entire instance to the UCMDB.

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" below](#)
- ["Adapter assist Methods" below](#)

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 for Federated TQL Queries

This section includes the following topics:

- ["Definitions and Terms" below](#)
- ["Mapping Engine" below](#)
- ["Federated Adapter" below](#)

See ["Interactions between the Federation Framework, Server, Adapter, and Mapping Engine" on the next page](#) for diagrams illustrating the interactions between the Federation Framework, BSM, adapter, and Mapping Engine.

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.

Interactions between the Federation Framework, Server, Adapter, and Mapping Engine

The following diagrams illustrate the interactions between the Federation Framework, BSM Server, the adapter, and the Mapping Engine. The federated TQL query in the example diagrams has only one virtual relationship, so that only the BSM and one external data repository are involved in the federated TQL query.

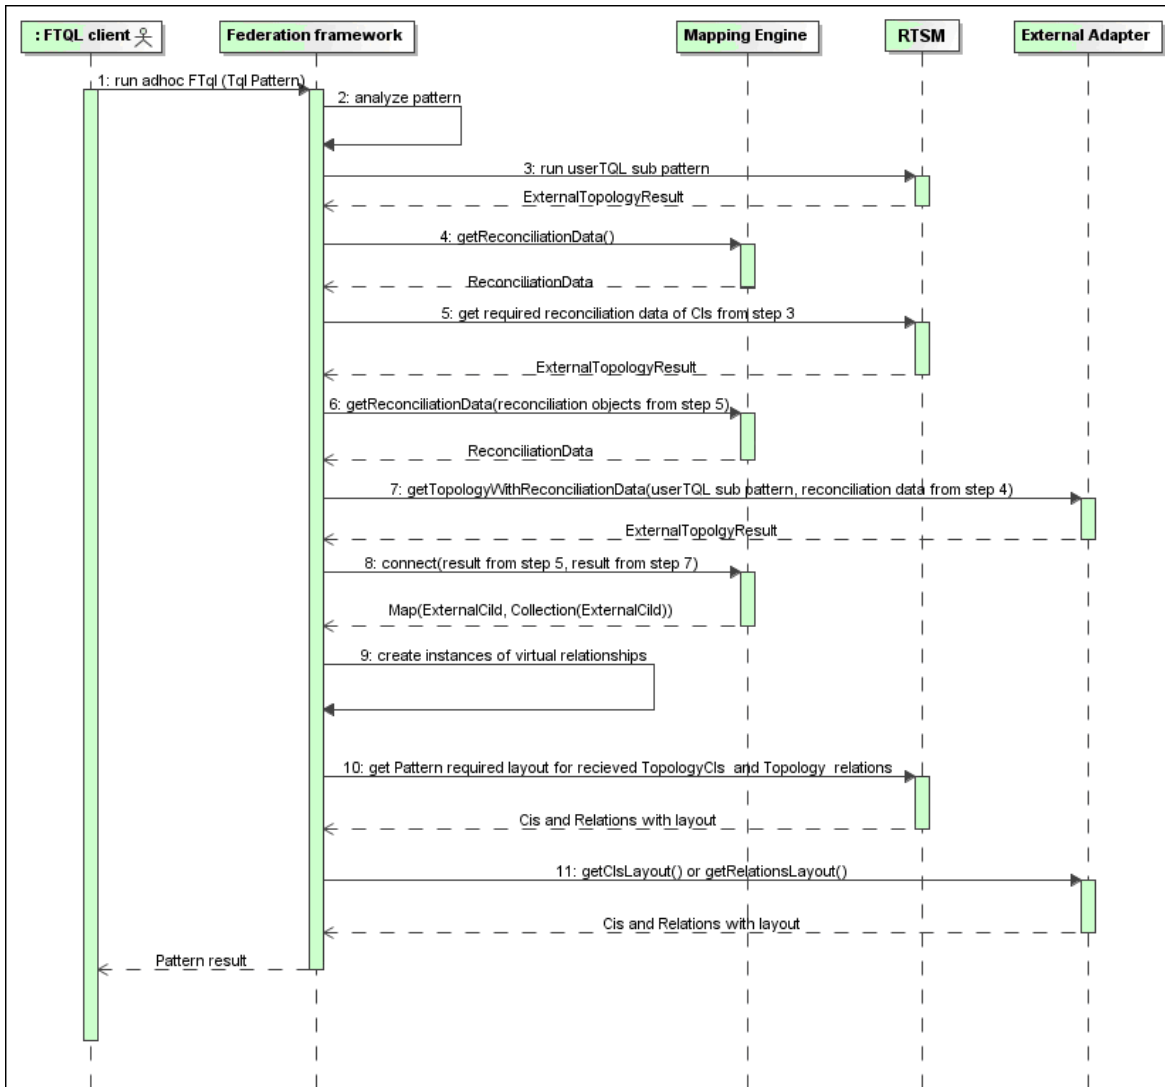
This section includes the following topics:

- ["Calculation Starts at the Server End" below](#)
- ["Calculation Starts at the External Adapter End" on page 152](#)
- ["Example of Federation Framework Flow for Federated TQL Queries" on page 153](#)

In the first diagram the calculation begins in the BSM and in the second diagram in the external adapter. Each step in the diagram includes references the appropriate method call of the adapter or mapping engine interface.

Calculation Starts at the Server 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.

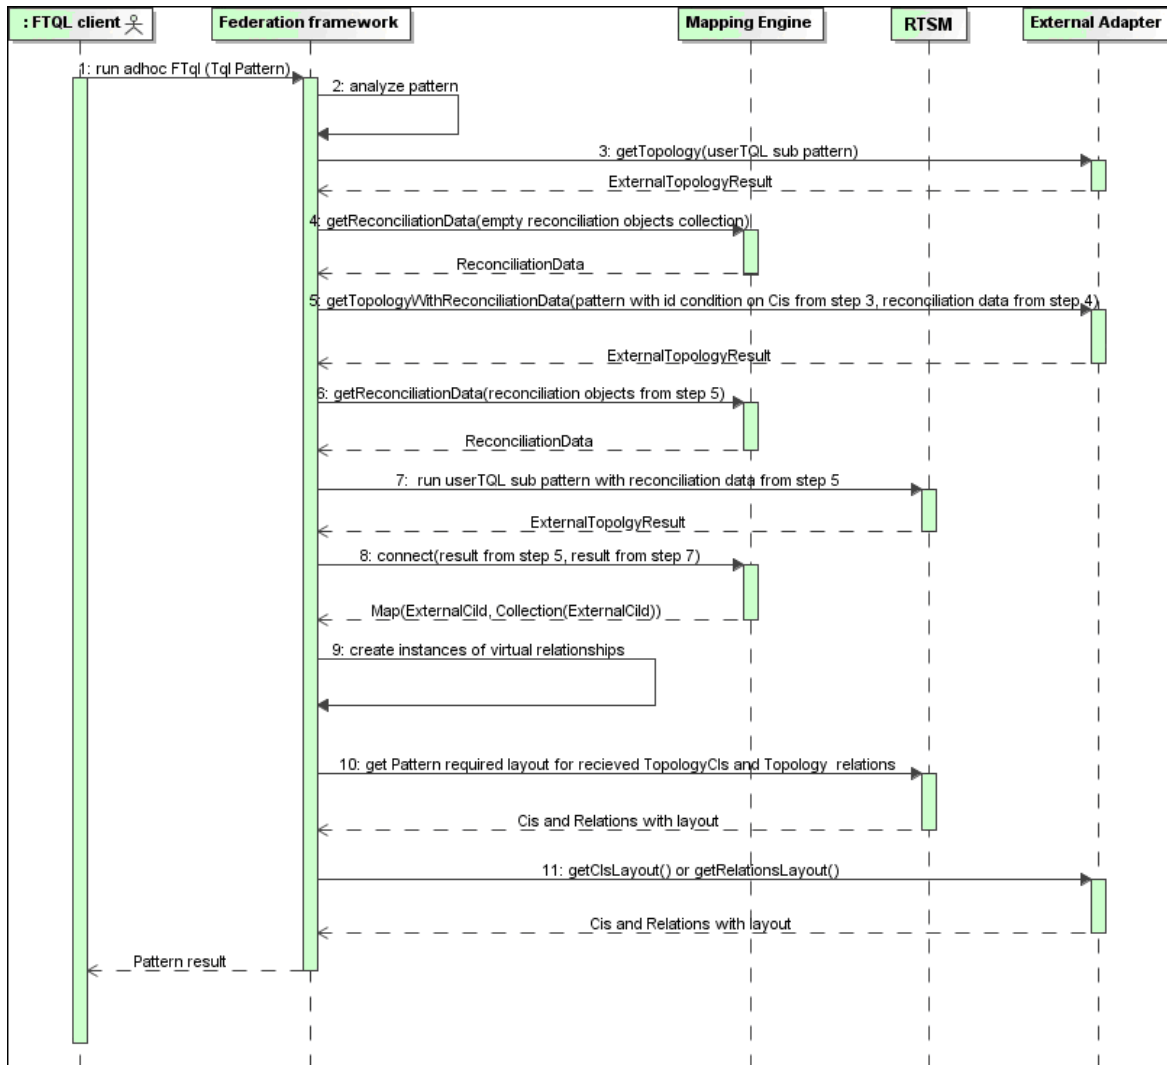


The numbers in this image are explained below:

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.

Number	Explanation
4	<p>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:</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 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" on the previous page from BSM.
6	The Federation Framework calls the Mapping Engine to retrieve the reconciliation data. In this state (by contrast with step "3" on the previous page), the Mapping Engine receives the reconciliation objects from step "5" above 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" above as a parameter.
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" above and the <code>secondResult</code> parameter is the external topology result received from the External Adapter in step "7" above . 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:

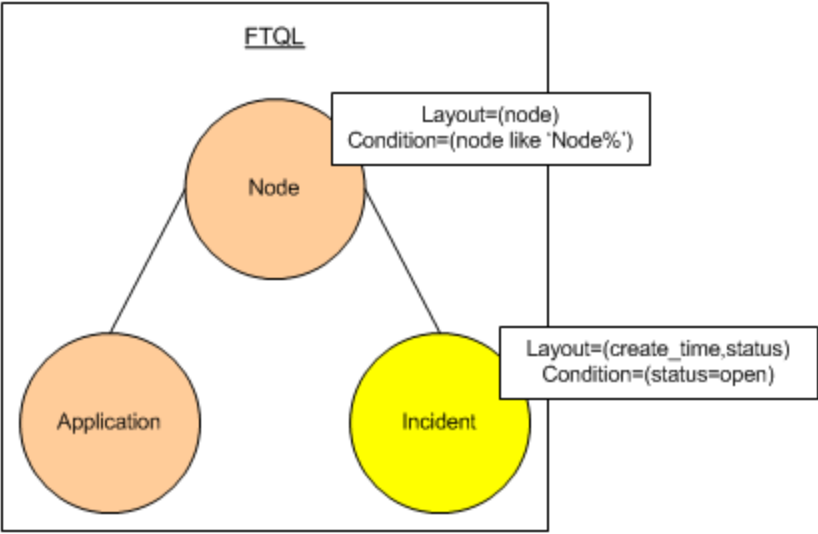
Number	Explanation
1	The Federation Framework receives a call for an federated TQL calculation.
2	The Federation Framework analyzes the adapter, finds the virtual relationship, and divides the original TQL into two sub-adapters – one for BSM and one for the external data repository.
3	The Federation Framework requests the topology of the sub-TQL from the External Adapter. The returned ExternalTopologyResult is not supposed to contain any reconciliation object, since the reconciliation data is not part of the request.

Number	Explanation
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 IP is required too).
5	<p>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 getTopologyWithReconciliationData() 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.</p>
6	<p>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.</p>
7	<p>The Federation Framework requests the topology of the sub-TQL with reconciliation data from step 6 from BSM.</p>
8	<p>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).</p>
9	<p>For each mapping, the Federation Framework creates a virtual relationship.</p>
10	<p>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.</p>

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.

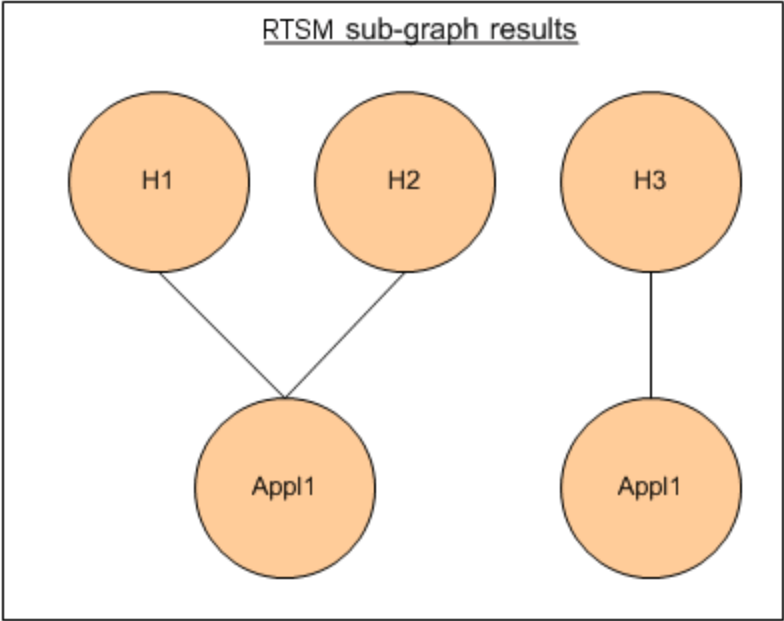


Note: For attribute federation, the adapter's **getTopology** method is called. The reconciliation data is adapted in the user TQL (in this case, the CI element).


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:



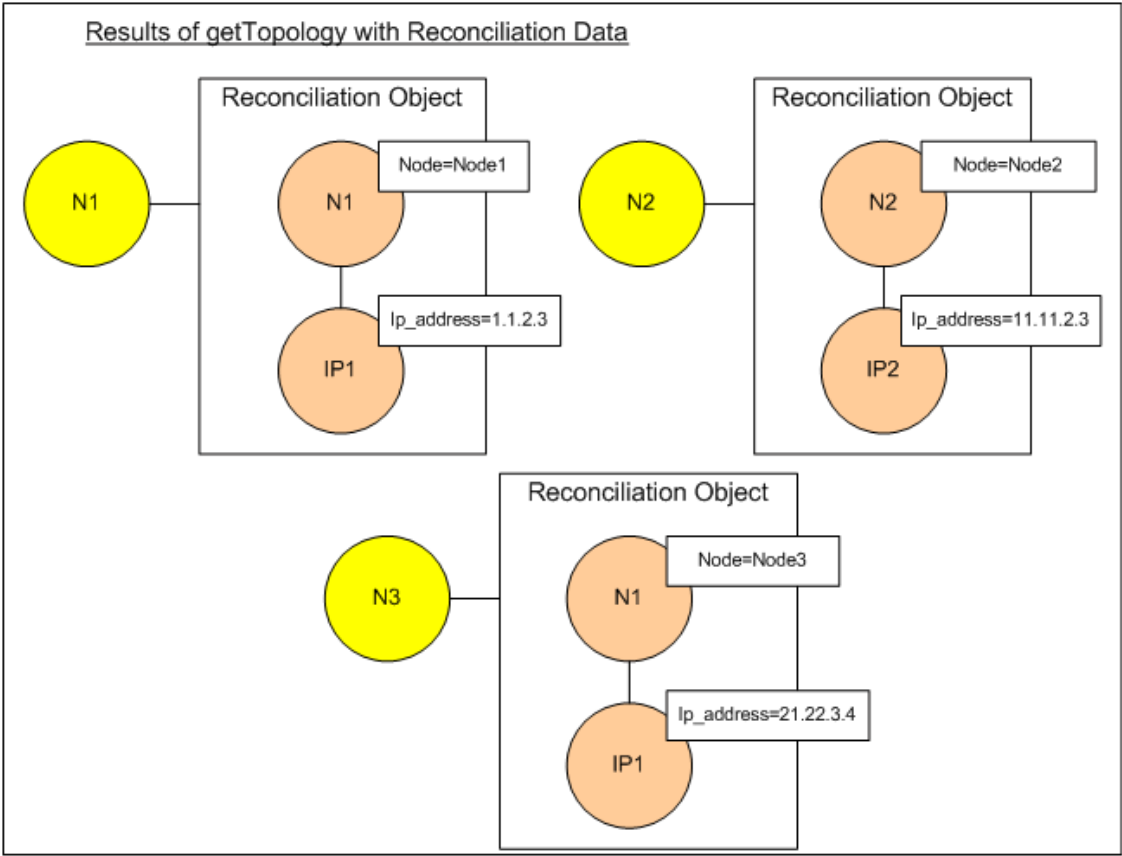
2. The Federation Framework runs the BSM subgraph to request the topology, and receives the following results:




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:



4. 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:

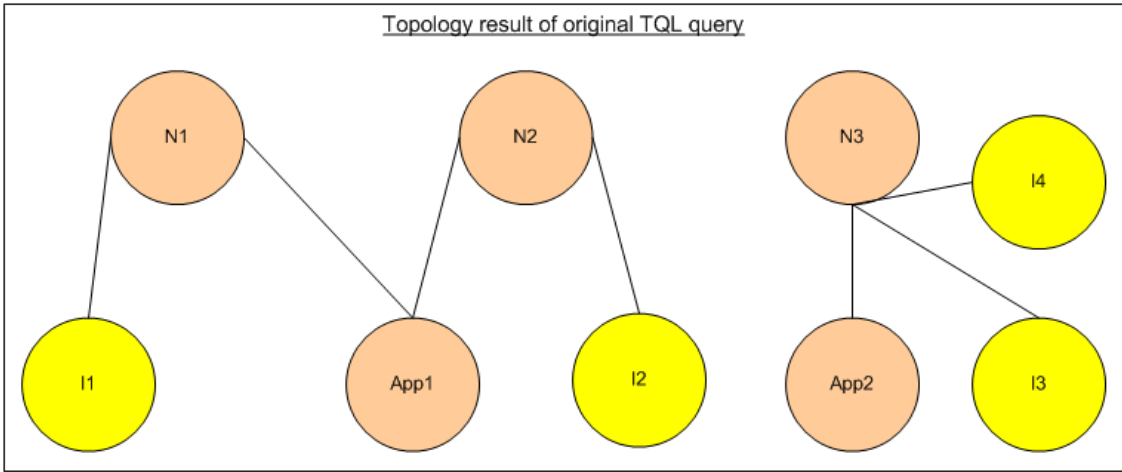


6. The Federation Framework runs the ServiceCenter subgraph with the reconciliation data to

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:

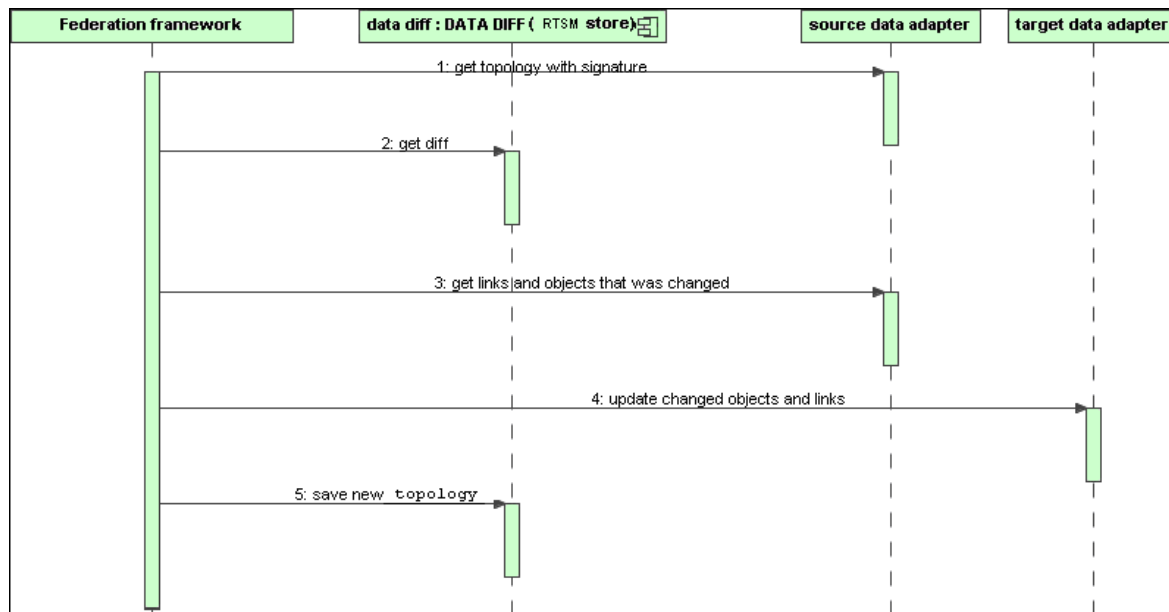
- ["Definitions and Terms" below](#)
- ["Flow Diagram" below](#)

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" below](#)
- ["Adapter Interfaces for Federated TQL Queries" below](#)

Definitions and Terms

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 **Single 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. The reconciliation data needed to complete the operation can be described as complex query (see [SingleNodeFederationTopologyReconciliationAdapter](#) below).

All SingleNode interfaces are created to simplify the workflow; for those cases where you need to use a more extensive query, use the **FederationTopologyAdapter** interface.

- A **FederationTopologyAdapter interface** is used to define adapters that support complex federated queries. The reconciliation request in these adapters is part of the **QueryDefinition** parameter.

The Federation engine uses reconciliation data in order to connect the federated data to the proper local CIs. Reconciliation data may be fetched in more than one request (calculated recursively according to results). In this case, the adapter receives a request with only reconciliation data.

SingleNode Interfaces

The following interfaces have different types of reconciliation data:

- **SingleNodeFederationIdReconciliationAdapter.** Use if the adapter supports a **single-node TQL** and the reconciliation between data repositories is calculated by the ID.
- **SingleNodeFederationPropertyReconciliationAdapter.** Use if the adapter supports a **single-node TQL** and the reconciliation between data repositories is done by the properties of one CI.
- **SingleNodeFederationTopologyReconciliationAdapter.** Use if the adapter supports a **single-node TQL** and the reconciliation between data repositories is done by topology. The adapter should support the case where the query element is empty and only reconciliation topology is requested.

Data Adapter Interfaces

- **FederationTopologyAdapter.** Use this adapter to support complex federated TQL queries. Allows the most diversity. The adapter should support the case where the query definition is describing only reconciliation data.
- **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.

Note: When developing an adapter that may return large data sets of data, its important to allow chunking by implementing the ChunkGetter Interface. See the Java document of the specific adapter for more information.

Resource Reporting Interfaces

The following interfaces enable the adapter to report the resources that can be configured to customize the adapter's behavior. This enables you to edit these resources directly from the Integration Studio. These interfaces should be used in addition to the regular adapter interfaces above.

- **PopulationQueriesResourcesLocator.** Defines which resources may be edited for each specific Population query.
- **PushQueriesResourceLocator.** Defines which resources may be edited for each Data Push query.
- **GeneralResourcesLocator.** Defines which general resources may be edited in this adapter.

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.

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:29000/jmx-console`
 - For the Probe: `http://localhost:1977`

You may have to log in with a user name and password.
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**.

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 the next page](#)
- ["Define Valid Relationships for Virtual Relationships" on the next page](#)
- ["Define an Adapter Configuration" on page 163](#)
- ["Define Supported Classes" on page 167](#)
- ["Implement the Adapter" on page 167](#)
- ["Define Reconciliation Rules or Implement the Mapping Engine" on page 168](#)
- ["Add Jars Required for Implementation to the Class Path" on page 168](#)

- ["Deploy the Adapter" on page 168](#)
- ["Update the Adapter" on page 169](#)

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 RTSM CI types to understand how external CITs are related to the BSM CITs
- model the external CITs in the RTSM class model
- add the definitions for new CI types and their relationships
- define valid relationships in the RTSM class model for the valid relationships between adapter inner classes. (The CITs can be placed at any level of the RTSM 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 RTSM class model, see [Working with the CI Selector](#) in the 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 only relevant 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.

- Create a valid links XML file that contains these links (if they do not already exist).
- 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 **Administration > RTSM Administration > Adapter Management**.
- b. Click the **Create new resource**  button and select **New Adapter**.
- c. In the New adapter dialog box, select **Integration** and **Java Adapter**.
- d. Right-click 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.BaseDestinationDa
ta">

<!-- 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-->

<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 170](#).

4. Define Supported Classes

Define supported classes either in 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.

If the adapter implements **getTopologyWithReconciliationData** and adapter capabilities include the ability to be used as a starting point, the adapter should also support requesting topology with reconciliation data without any conditions (only type). In this case the adapter should return full reconciliation data of the found results.

Adapter reconciliation support can be defined according to **global_id**, in which case **global_id** must be defined as part of the reconciliation attributes in the adapter supported classes. If adapter reconciliation support is defined according to **global_id**, then **getTopologyWithReconciliationData()** should return the **global_id** as part of the reconciliation object properties. The RTSM uses **global_id** for reconciliation of federation results for a CIT instead of the identification rule.

Part of the federation API is the `DataAdapterEnvironment` interface. This interface represents the environment of the data adapter. It contains the environment API needed for the adapter to work. For more information on the `DataAdapterEnvironment` interface, see ["The DataAdapterEnvironment Interface" on page 172](#).

6. Define Reconciliation Rules or Implement the Mapping Engine

If your adapter supports federated TQL queries, you have two options for defining your Mapping Engine:

- Use the default mapping engine, which uses the RTSM's internal reconciliation rules for mapping. To use it, leave the **<default-mapping-engine>** XML tag empty.
- 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 RTSM.
- 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" on page 1](#) in the *RTSM Administration Guide*.

Create a Sample Adapter

This example illustrates how to create a sample adapter. This task includes the following steps:

- ["Select Adapter Logic" below](#)
- ["Load the Project" on the next page](#)

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. 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 option *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\APIs\DBAdapterFramework_JavaAPI\index.html

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.
	<start-point-adapter>	Specifies that this adapter is the start point for TQL query calculation.
	<one-node-topology>	The ability to federated queries with one federated query node.
	<pattern-topology>	The ability to federate complex queries.
	<support-replication-data>	Defines the capability to run data push and population flows.
	<source>	This adapter may be used for population flows.
	<push-back-ids>	Push back the global ID of the CI to the global_id column of the table (must be defined in the orm.xml). The behavior can be overridden by implementing the FcmdbPluginPushBackIds plug-in.
	<changes-source>	This adapter may be used for population changes flows.
	<instance-based-data>	This tag defines that the adapter supports an instance based population flow.
	<target>	This adapter may be used for data push flows.

	<default-mapping-engine>	Allows defining 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
	<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.
	<adapter-settings>	The list of settings of the adapter.
	<list.attributes.for.set>	Determines which attributes override the previous value (if one exists).

The DataAdapterEnvironment Interface

OutputStream openResourceForWriting(String resourceName) throws FileNotFoundException;

This method opens a resource with a given name for writing. It is used for saving persistent data for the integration. This method should be used instead of trying to load files using java methods. The user should ensure that the stream is closed when finished writing to the stream. `close()/flush()` will save the resource. This method creates a runtime resource (it may not overwrite files that came in the adapter package).

Parameter

- **resourceName:** The name of the resource to retrieve. This name should be unique across all integrations of the same adapter.

Return Value

Returns a stream to which to write.

Exceptions

- This method throws the *FileNotFoundException* if the resource type is file and the file does not exist, if the resource is a directory rather than a regular file, or for some other reason the resource cannot be opened for reading.
- This method throws the *SecurityException* if a security manager exists and its *checkRead* method denies access to the file.

InputStream openResourceForReading(String resourceName) throws FileNotFoundException;

This method opens a resource with a given name for reading. It is used for reading persistent data for the integration. This method should be used instead of trying to load a file using java methods. The user should ensure that the stream is closed when finished reading it. It first attempts to load files that came in the adapter package. If not found, it attempts to load a runtime created resource from *DataAdapterEnvironment.openResourceForWriting(String)*. The runtime resources can be viewed using JMX (of the probe and server accordingly).

Parameter

- **resourceName:** The name of the resource to retrieve. This name should be unique across all integrations of the same adapter.

Return Value

Returns a stream to read.

Exceptions

- This method throws the *FileNotFoundException* if the resource type is **file** and the file does not exist, if the resource is a directory rather than a regular file, or for some other reason the resource cannot be opened for reading.
- This method throws the *SecurityException* if a security manager exists and its *checkRead* method denies read access to the file.

Properties openResourceAsProperties(String propertiesFile) throws IOException;

This method opens a resource with a given name and loads it as a *Properties* structure. It is used for reading persistent data for the integration. This method should be used instead of attempting to load the **.properties** files using java methods. It first attempts to load files that came in the adapter package. If not found, it attempts to load a runtime created resource from *DataAdapterEnvironment.openResourceForWriting(String)*. The runtime resources can be viewed using JMX (of the probe and server accordingly).

Parameter

- **propertiesFile:** The name of the resource to retrieve. This name should be unique across all integrations of the same adapter.

Return Value

Returns the file content represented in Properties.

Exceptions

- This method throws the *FileNotFoundException* if the resource type is **file** and the file does not exist, if the resource is a directory rather than a regular file, or for some other reason the resource cannot be opened for reading.
- This method throws the *SecurityException* if a security manager exists and its *checkRead* method denies read access to the file.
- This method throws the *IOException* if the properties file failed to convert to the *Properties* Object.

String openResourceAsString(String resourceName) throws IOException;

This method opens a resource with a given name and loads it as a string. It is used for reading persistent data for the integration. This method should be used instead of trying to load files using java methods.

It first attempts to load files that came in the adapter package. If not found, it attempts to load a runtime created resource from *DataAdapterEnvironment.openResourceForWriting(String)*. The runtime resources can be viewed using JMX (of the probe and server accordingly).

Parameter

- **resourceName:** The name of resource to retrieve. This name should be unique across all integrations of the same adapter.

Return Value

Returns the file content represented in String format.

Exceptions

- This method throws the *FileNotFoundException* if the resource type is **file** and the file does not exist, if the resource is a directory rather than a regular file, or for some other reason the resource cannot be opened for reading.
- This method throws the *SecurityException* if a security manager exists and its *checkRead* method denies read access to the file.
- This method throws the *IOException* if an I/O error occurs.

public void saveResourceFromString(String relativeFileName, String value) throws IOException;

This method receives a String and saves it as a resource. It is used for saving persistent data for the integration. This method should be used instead of trying to save files using java methods. This method converts the String into a stream and saves it to the resource. It creates a runtime resource, but cannot overwrite files that came in the adapter package). The runtime resources can be viewed using JMX (of the probe and server accordingly).

Parameter

- **relativeFileName:** The name of resource to retrieve. This name should be unique across all integrations of the same adapter.
- **value:** The String to save as a resource

Exceptions

This method throws the IOException if an I/O error occurs.

boolean resourceExists(String resourceName);

This method checks if the given resource name exists. It looks for files that came in the adapter package and for runtime created resources from *DataAdapterEnvironment.openResourceForWriting(String)*.

Parameter

- **resourceName:** The name of the resource to retrieve. This name should be unique across all integrations of the same adapter.

Return Value

Returns **True** if *resourceName* exists.

boolean deleteResource(String resourceName);

This method deletes the given resource from persistent data. It deletes a runtime resource, and may not delete files that came in the adapter package. The runtime resources can be viewed using JMX (for the probe and server accordingly).

Parameter

- **resourceName:** The name of the resource to delete. This name should be unique across all integrations of the same adapter.

Return Value

Returns **True** if the resource is successfully deleted.

Collection<String> listResourcesInPath(String path);

This method retrieves a list of resources in the given resource path. It looks for files that came in the adapter package and for a runtime created resources from *DataAdapterEnvironment.openResourceForWriting(String)*. The runtime resources can be viewed using JMX (for the probe and server accordingly).

Parameter

- **path:** The resource path. For example, "META-INF/myfiles/"

Return Value

Return a list of resources in the path.

DataAdapterLogger getLogger();

Retrieves the logger to be used by the adapter. This logger is used for logging events in your adapter.

Return Value

Returns the logger that is used by the DataAdapter.

DestinationConfig getDestinationConfig();

This method retrieves the destination configuration of the integration. This configuration holds all connection and running settings for the integration.

Return Value

Returns the DestinationConfig of the Adapter.

int getChunkSize();

This method retrieves the population chunk size requested for this integration.

Return Value

Returns the population chunk size.

int getPushChunkSize();

This method retrieves the push chunk size requested for this integration.

Return Value

Returns the push chunk size.

ClassModel getLocalClassModel();

This method retrieves a class model for querying information about the local RTSM class model. This method brings an updated ClassModel. Once the ClassModel object is returned, it is not updated for any class model changes. In order to retrieve an updated class model, use this method again to retrieve it.

Return Value

Returns the RTSM class model.

CustomerInformation getLocalCustomerInformation();

This method retrieves customer information for the customer that is executing the adapter.

Return Value

Returns customer information for the customer that is executing the adapter.

Object getSettingValue(String name);

This method retrieves a specific adapter setting.

Parameter

name: The name of setting.

Return Value

Returns the Object setting value.

Map<String, Object> getAllSettings();

This method retrieves all adapter settings.

Return Value

Returns the adapter settings.

boolean isMTEnabled();

This method checks if the server environment supports Multiple Tenancy (MT).

Return Value

Returns **true** if the server environment supports MT, otherwise returns **false**.

String getUcldbServerHostName();

This method returns the local RTSM server host name.

Return Value

Returns the local RTSM server host name.

Chapter 5: Developing Push Adapters

This chapter includes:

Developing and Deploying Push Adapters	178
Build an Adapter Package	179
Create Mappings	182
Write Jython Scripts	186
Support Differential Synchronization	190
Generic XML Push Adapter SQL Queries	192
Generic Web Service Push Adapter	192
Mapping File Reference	211
Mapping File Schema	213
Mapping Results Schema	224
Customization	228

Developing and Deploying Push Adapters

Generic Push Adapters provide a common platform that enables rapid development of integrations that push BSM data to external data repositories (databases and third-party applications). Generic Push Adapters are categorized according to the protocol used to push the data. For details on pushing via XML, using the Generic XML Push Adapter, see ["Generic XML Push Adapter SQL Queries" on page 192](#). For details on pushing via Web Service, using the Generic Web Service Push Adapter, see ["Generic Web Service Push Adapter" on page 192](#).

Developing a custom integration based on a Generic Push Adapter requires:

- Building a new adapter package from the appropriate Generic Push Adapter template files. For details, see ["Build an Adapter Package" on the next page](#).
- Mappings between the BSM CI link types and the external data items. The mappings are stored as XML and are customized to each external data repository. For details, see ["Create Mappings" on page 182](#).
- A Jython script to push the data items into the external data repository. For details, see ["Write Jython Scripts" on page 186](#).
- Additional adapter-specific steps. For example, choosing the path of the file to be written for the XML push adapter, or creating a data receiver for the Web Service push adapter.

Build an Adapter Package

To create a new, MDR-specific push adapter, you should make a copy of the generic adapter and then edit it to customize it as an adapter for a specific push target.

Generic adapters packages can be found in one of the following two locations:

- Generic XML push adapter: **<HP BSM root directory\odb\content\adapters\push-adapter.zip**
- Generic web service adapter: **<HP BSM root directory\odb\content\adapters\service-push-adapter.zip**

To create a new push adapter from the generic push adapter:

1. Extract the content of the selected package zip file to a work folder.
 2. Review the following directories in preparation for the rename and replace phase:
 - **adapterCode:** Contains the directory that is deployed to the **C:\hp\UCMDB\UCMDBServer\runtime\fcmdb\CodeBase** directory. Jars deployed here do not automatically restart the probe and do not appear automatically in the probe's CLASSPATH.
 - **discoveryConfigFiles:** Contains the adapter's mappings definitions and points to the correct Jython script (**push.properties**)
 - **discoveryPatterns:** Contains the adapter's XML definition that is deployed on the UCMDB server
 - **discoveryScripts:** Contains the adapter's Jython scripts via which the connection to the third party data store is made and data is pushed
 - **discoveryResources:** Contains the **UCMDBDataReceiver.jar** containing the Java integration classes for the web service.
- Note:** When you deploy this package, the probe is restarted to include this **.jar** in the probe's CLASSPATH. No action is required beyond deploying the package.
3. Make the following changes within the unzipped adapter directory structure:
 - a. **discoveryConfigFiles\<Your_Push_Adapter_Name>:** rename the directory "PushAdapter" or "XMLtoWebService" to the name of the new push adapter (for example, "myPushAdapter").
 - b. **discoveryConfigFiles\<Your_Push_Adapter_Name>\push.properties:** In the **push.properties** file, do the following:

- Update the name of **jythonScript.name** to the name of the Jython script to be used by the new push adapter (for example, **pushToMyService.py**).
- Update the name of the mappings file to be used by the new push adapter (for example, **myPushAdapter_mappings**). Do not add the **.xml** extension, this is filled in automatically.
- c. **discoveryPatterns\<push adapter name>.xml**: Rename this file to the name of the new adapter's definition XML file (for example, **my_push_adapter.xml**).
- d. **discoveryPatterns\<your_push_adapter>.xml**: Update this file as follows:
 - For the XML element **<pattern>**: set the id and description attributes accordingly. For example:

```
<pattern id="PushAdapter"
xsi:noNamespaceSchemaLocation="../../Patterns.xsd" description="Discovery
Pattern Description" schemaVersion="9.0"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
```

is changed to:

```
<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">
```

- For the XML element **<parameters>**: update the children elements according to the needs of your adapter. By default, the following children elements are used to define a push adapter. These values are assigned when the integration point is defined in the Integration Studio after the adapter is configured. Update the parameter list so that the list of parameters reflects the required connection attributes. Do not remove the **probeName** attribute.
 - **host**: the server name hosting the web service
 - **port**: port of the listening UCMDDB Data Receiver service
 - **Web Service Push Adapter: uri** - the remainder of the URL to form the service endpoint address of the data receiver.
 - **probeName**: defines on which data flow probe the push job runs
- For the XML element **<integration>**: update the value of the child element **<category>** to something other than Generic. By default integration adapters belonging to the Generic category are not shown in the Integration Studio. If you are integrating with a third-party data store, set this value to "Third Party". If you are integrating with an HP BTO product, set this value to "HP BTO Products".

- e. **adapterCode\PushAdapter**: rename this folder with the adapter ID used in the previous step (for example, **adapterCode\MyPushAdapter**).
- f. **discoveryScripts\<your_Jython_push_script>.py**: create a file with the same name as the one defined in the **push.properties jythonScript.name** property. In the **discoveryScript** file, there is a script that inserts the CIs and links to an external Oracle database. Replace **discoveryScripts\pushScript.py** with the script you wrote (for details, see ["Write Jython Scripts" on page 186](#)). If you rename the script, the **jythonScript.name** property in **adapterCode\<adapter ID>\push.properties** should be updated accordingly.
 - XML Push Adapter: **pushScript.py**
 - Web Service Push Adapter: **XMLtoWebService.py**
- g. **tql\<your_integration_TQLs>**: like a regular package, place the TQL XML definition of your integration TQLs in this directory. All TQLs in this folder are deployed when the adapter package is deployed.
- h. **discoveryConfigFiles\<Your_Push_Adapter_Name>\mappings**: create an XML mapping file per TQL that you want to use in your integration. Note that the push adapter applies the transformations in the mapping file to the results of the integration TQLs and then sends that data in three parameters (**addResult**, **updateResult**, and **deleteResult**) of an ad hoc task to the data flow probe.
- i. **adapterCode\<adapter ID>\mappings**: replace the **mappings.xml** file with the mapping files you prepared (for details, see ["Create Mappings" on the next page](#)).

XML Push Adapter: This mapping example corresponds to the example of the tables created in ORACLE in the **sql_queries** file.

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 is used as by 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**.

- 4. After making all the above changes, create a **.zip** file by selecting the folders and files specified in step 3 above (for example, **my_Push_Adapter.zip**).
- 5. Deploy the newly created **.zip** file on the UCMDDB server via the Package Manager (go to **Administration > Package Manager**).
- 6. Create an integration point in **Data Flow Management > Integration Studio** and define the integration TQLs that the integration point uses. Set a schedule for automatic data push.

Troubleshooting

The procedure for building a new push adapter requires complete and correct re-naming and replacing. Any error will likely affect the adapter. The package must be unzipped and re-zipped correctly to act as

a UCMDB package. Refer to the out-of-the-box packages as examples. Common errors include:

- Including another directory on top of the package directories in the ZIP file.
Solution: ZIP the package in the same directory as the package directories such as **discoveryResources**, **adapterCode**, etc. Do not include another directory level on top of this in the ZIP file.
- Omitting a critical re-name of a directory, a file, or a string in a file.
Solution: Following the instructions in this section very carefully.
- Misspelling a critical re-name of a directory, a file, or string in a file.
Solution: Do not change your naming convention in mid-stream once you begin the re-naming procedure. If you realize that you need to change the name, start over completely rather than trying to retroactively correcting the name, as there is a high risk of error. Also, use search and replace rather than manually replacing strings to reduce risk of errors.
- Deploying adapters with the same file names as other adapters, especially in the **discoveryResources** and **adapterCode** directories.
Solution: You may be using a UCMDB version with a known issue that prevents mappings files from having the same name as any other adapter in the same UCMDB environment. If you attempt to deploy a package with duplicates names, the package deployment will fail. This problem may occur even if these files are in different directories. Further, this problem can occur regardless of whether the duplicates are within the package or with other previously deployed packages.

At this point you can create a new push adapter job in the Integration Studio using the new adapter you just deployed.

TQL Best Practices for Push Adapters

1. Create a folder structure in the TQL and View trees, and keep all new TQLs and views there. Use a naming convention.
2. Unless the TQL is small, first copy the most similar TQL.
3. Make one change at a time. Save, test, and preview after each change. Repeat until the results comply with your requirements.

Create Mappings

The raw TQL result data is in the form of the RTSM class model schema. It is likely that the consumer uses a different data model. The push adapter provides a mapping mechanism to transform the data into a format more suitable for consumption. Mappings perform both direct and complex transformations, from direct, naming-type conversion, to parent/child aggregation and referencing functions.

The mapping specification can be found in the section ["Mapping File Reference" on page 211](#). Use the reference to create a mapping file.

Note: The adapter properties file refers to the name of the mapping file. In adapter configuration files, the adapter implements a folder structure using the name of the adapter. Rename this folder when implementing an adapter to maintain uniqueness as required by the Package Manager.

Build a Mapping File

To build a mapping file:

1. Start with a default mapping file.
2. Deploy the adapter and run it once.
3. Observe the results.
4. Identify and note what should be changed.
5. Make the changes identified in the previous step. The following list can help serve as a guide for the order of the changes.
 - a. Start with the top, non-transformative section. Make sure the adapter runs after each change.
 - b. Change the source CIs section to the UCMDB names in the TQL result.
 - c. First map the keys.
 - d. Then add all the direct mappings.
 - e. Add the complex mappings.
 - f. Add the link mappings.

Repeat steps 2-5 until the mapped data is suitable for consumption. Select the appropriate Generic adapter package from which to create the new push adapter.

The mapping files work the same way for all types of push adapters. The Generic XML push adapter writes the mapped results to a file. The Generic Web Service Push Adapter sends the XML results to a data receiver. For more details, see ["Generic Web Service Push Adapter" on page 192](#).

Prepare the Mapping Files

Note: You can retrieve all of the CIs and relationships as they are in the CMDB without mapping, by not creating the **mappings.xml** file. This returns all of the CIs and relationships with all of their attributes.

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 mapping file is called **<query name>.xml**.

For details, see ["Mapping File Schema" on page 213](#).

This task includes the following steps:

- ["Create a mappings.xml File" below](#)
- ["Map CIs" below](#)
- ["Map Links" on the next page](#)

1. Create a mappings.xml File

The mapping file structure is created as follows (use an existing file as a template):

```
<?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 are processed. CIs of inherited types are not 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 so that links of that type and all inherited links are mapped in the same way:

```
<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 so that only links of that type are processed. Links of inherited types are not processed unless their type is also mapped (in one of the two ways):

```
<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>
```

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 37](#).

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 the step on editing the **discoveryPatterns\push_adapter.xml** file in ["Build an Adapter Package" on page 179](#)), 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" below](#)
- ["Handling Test Connection in the Script" on page 190](#)

Working with the Mapping's Results

Generic push adapters create 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.

In the XML, the `id` attribute of a CI holds the `cmdbId` as an external id or the `ExternalId` of that CI if the CI got an `ExternalId` one when the CI was sent to the script. The `end1Id` and `end2Id` fields of the link hold for each of the link's ends the `cmdbId` as an external id or the `ExternalId` of that link's end if the CI at the link's end got an `ExternalId` when it was sent to the script.

When processing the CIs in the Jython script, the return value of the script is a mapping between the CI's CMDB id and the given id (the id given to each CI in the script). If a CI is pushed for the first time, the id that is in the XML of that CI is the CMDB id. If a CI is not pushed for the first time, the CI's id is the same id that was given to that CI in the script when it was first pushed.

The id is retrieved from the CI XML script as follows:

1. From the CI Element in the XML, retrieve the id from the id attribute. For example: `id = objectElement.getAttributeValue('id')`.
2. After retrieving the id from the XML, restore the id from the attribute (string). For example: `objectId = CmdbObjectID.Factory.restoreObjectID(id)`.
3. Check if the `objectId` received in the previous step is the CMDB id. You can do this by checking if the `objectId` has the new id that is given to it by the script. If it does, the returned id is not the CMDB id. For example:
`newId = objectId.getPropertyValue(<the name of the id attribute which is given by the script>).`

If `newId` is null, then the id that was returned in the XML is a CMDB id.

4. If the id is a CMDB id (that is, `newId` is null), perform the following (if the id is not a CMDB id, go to step 5):

- a. Create a property for that CI that holds the new id. For example: `propArray = [TypesFactory.createProperty('<the name of the id attribute which is given by the script>', '<new id>')]`.
- b. Create an `externalId` to that CI. For example:

```
cmdbId = extI.getPropertyValue('internal_id')
className = extI.getType()
externalId = ExternalIdFactory.createExternalCiId(className, propArray)
```
- c. Map the CMDB id to the newly created `externalId` (and in the next step return that mapping to the adapter). For example: `objectMappings.put(cmdbId, externalId)`
- d. When all of the CIs and links are mapped:

```
updateResult = DataPushResultsFactory.createDataPushResults(objectMappings,
linkMappings);
return updateResult
```

5. If the id is the new id (that is, `newId` is not null), then the `externalId` is the `newId`.

It is also possible to report on the push status for each CI and link as follows:

1. `updateStatus = ReplicationActionDataFactory.createUpdateStatus();`
where `updateStatus` is an instance of the `UpdateStatus` class that contains statuses of the CIs and links.
2. Add a status to `updateStatus` by calling the `reportCIStatus` or `reportRelationStatus` method.

For example:

```
status = ReplicationActionDataFactory.createStatus(Severity.FAILURE, 'Failed',
ERROR_CODE_CI, errorParams, Action.ADD);
updateStatus.reportCIStatus(externalId, status);
```

Where `ERROR_CODE_CI` is the number of the error messages as they appear in adapter **properties.errors** file (for details on the **properties.errors** file, see ["Error-Writing Conventions" on page 1](#)), and `errorParams` contains the parameters to pass to the message. See **ReplicationActionDataFactory** javadoc for more details.

3. Create a push result with the statuses as follows:

```
updateResult = DataPushResultsFactory.createDataPushResults(objectMappings,
linkMappings, updateStatus);
return updateResult
```

Example of the XML result

```
<root>
```

```
<data>
  <objects>
    <Object mode="update_else_insert" name="UCMDB_UNIX" operation="add"
mamId="0c82f591bc3a584121b0b85efd90b174"
id="HiddenRmiDataSource%0Aunix%0A1%0Ainternal_
id%3DSTRING%3D0c82f591bc3a584121b0b85efd90b174%0A">

      <field name="NAME" key="false" datatype="char" length="255">UNIX5</field>
      <field name="DATA_NOTE" key="false" datatype="char" length="255"></field>
    </Object>
  </objects>
  <links>
    <link targetRelationshipClass="TALK" targetParent="unix" targetChild="unix"
operation="add" mode="update_else_insert"
mamId="265e985c6ec51a8543f461b30fa58f81"
id="end1id%5BHiddenRmiDataSource%0Aunix%0A1%0Ainternal_
id%3DSTRING%3D41372a1cbcaba27b214b84a2ec9eb535%0A%5D%0Aend2id%
5BHiddenRmiDataSource%0Aunix%0A1%0Ainternal_
id%3DSTRING%3D0c82f591bc3a584121b0b85efd90b174%0A%5D%0AHiddenRmi
DataSource%0Atalk%0A1%0Ainternal_
id%3DSTRING%3D265e985c6ec51a8543f461b30fa58f81%0A">

      <field name="DiscoveryID1">41372a1cbcaba27b214b84a2ec9eb535</field>
      <field name="DiscoveryID2">0c82f591bc3a584121b0b85efd90b174</field>
      <field name="end1Id">HiddenRmiDataSource%0Aunix%0A1%0Ainternal_
id%3DSTRING%3D41372a1cbcaba27b214b84a2ec9eb535%0A</field>
      <field name="end2Id">HiddenRmiDataSource%0Aunix%0A1%0Ainternal_
id%3DSTRING%3D0c82f591bc3a584121b0b85efd90b174%0A</field>
      <field name="NAME" key="false" datatype="char" length="255">TALK4</field>
      <field name="DATA_NOTE" key="false" datatype="char" length="255"></field>
    </link>
  </links>
</data>
</root>
```

Note: If datatype="BYTE", the returned result's value is a **String** that is generated as: new String([the byte array attribute]). The byte[] object can be reconstructed by: <the received String>.getBytes(). If there are differences in the default locale between the server and the probe, the reconstruction is performed according to the server's default locale.

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**.

Support 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.

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 ["Build an Adapter Package" on page 179](#).

This task enables the Push adapter to perform differential synchronization.

The Jython script returns the **DataPushResults** object which contains two Java maps - one for object ID mappings (keys and values are **ExternalCild** 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.fcmbd.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):

```
externalCIId = ExternalIdFactory.createExternalCmdbCiId(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.createExternalCiId(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 **restoreExternalCIId** 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.

Generic XML Push Adapter SQL Queries

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.

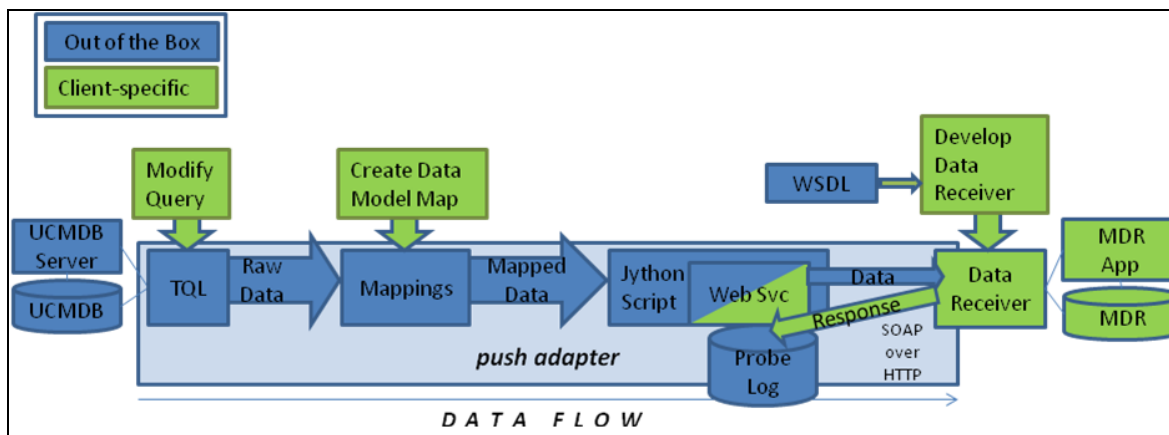
Generic Web Service Push Adapter

The generic web service push adapter provides a UCMDB-initiated push of SOAP messages containing query data to a web service data receiver. The mapped results are sent in standard SOAP messages via the HTTP POST protocol to the data receiver. The data receiver must understand the SOAP messages produced by the push adapter. To facilitate the development of the proper data receiver, a WSDL is provided with this push adapter.

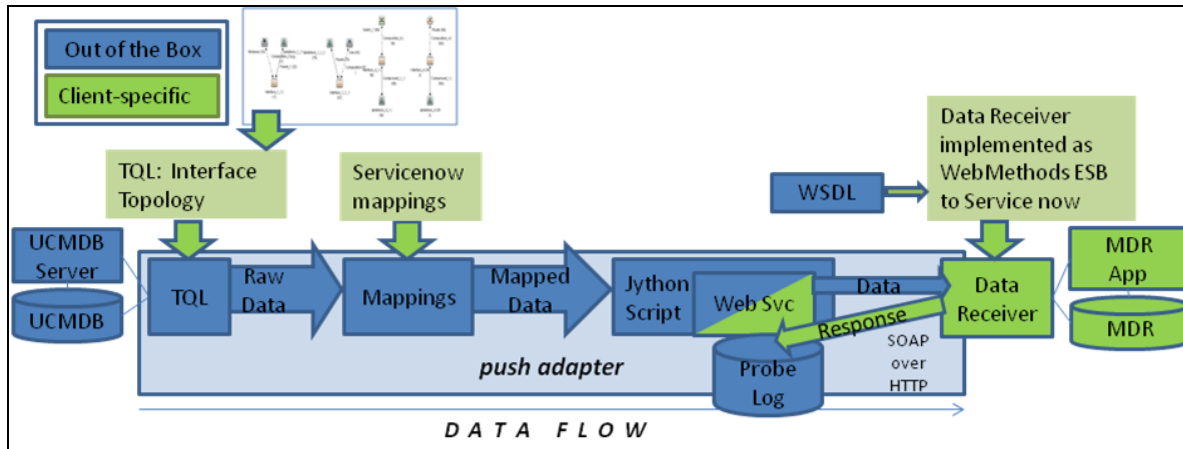
Custom processing of the SOAP message response XML is possible in the Jython script.

To understand the format of the incoming mapped data, the developer of the data receiver should communicate with the developer of the mappings file. An **.xsd** is currently not provided with this version of the web service push adapter, so data must be processed in a way reflective of the incoming data, which is a combination of the original TQL and the applied mappings.

The web service push adapter functions for pushing data to the client are shown below. The items in green are customized or supplied by the client to implement the adapter for a specific push target. The items in blue are out-of-the-box components.

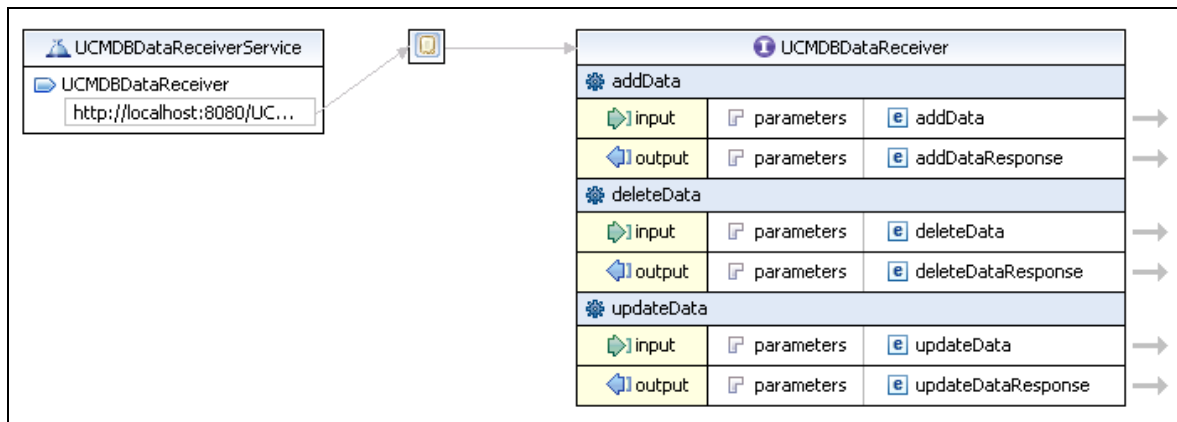


An example of an implementation of the generic web service push adapter to an MDR-specific push adapter using an Enterprise Service Bus (ESB) is shown here:



WSDL

A WSDL is supplied to the client developer to create a data receiver capable of communicating with the UCMDB push adapter via a web service. The **UCMDBDataReceiver.wsdl** describes the SOAP messages that are used to communicate data from the UCMDB to the data receiver. The design diagram from the WSDL is shown here:



The data receiver (which is in practice a server or "service endpoint" in SOAP terminology) should implement three methods: **addData**, **deleteData**, and **updateData**, corresponding to the data sets that the UCMDB pushes. The HTTP headers contain the correct **SoapAction** keyword that indicates the type of data that is being sent. The data receiver is responsible for implementing the business logic and processing the data.

The default WSDL URL is:

- <http://localhost:8080/UCMDBDataReceiver/services/UCMDBDataReceiver?wsdl>

As implemented by the Data Receiver, the URL could look similar to the following:

- <http://testWSPAserver:4444/MyCo.IT.SvcMgt.ws.us:provider/UCMDBDataReceiver?wsdl>

The URL of the web service is the same as the WSDL URL without the “?wsdl” at the end.

The source for the WSDL is included below:

```
<?xml version="1.0" encoding="UTF-8"?>

<wsdl:definitions targetNamespace="http://ucmdb.hp.com"
  xmlns:apachesoap="http://xml.apache.org/xml-soap"
  xmlns:impl="http://ucmdb.hp.com" xmlns:intf="http://ucmdb.hp.com"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
  xmlns:wsdlsoap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">

  <!--WSDL created by Apache Axis version: 1.4 Built on Apr 22, 2006 (06:55:48
  PDT)-->

  <wsdl:types>

    <schema elementFormDefault="qualified"
      targetNamespace="http://ucmdb.hp.com"
      xmlns="http://www.w3.org/2001/XMLSchema">

      <element name="addData">

        <complexType>

          <sequence>

            <element name="xmlAdded" type="xsd:string"/>

          </sequence>

        </complexType>

      </element>

      <element name="addDataResponse">

        <complexType/>

      </element>

      <element name="deleteData">

        <complexType>

          <sequence>

            <element name="xmlDeleted" type="xsd:string"/>

          </sequence>

        </complexType>

      </element>

      <element name="deleteDataResponse">
```

```
        <complexType/>
      </element>
      <element name="updateData">
        <complexType>
          <sequence>
            <element name="xmlUpdate" type="xsd:string"/>
          </sequence>
        </complexType>
      </element>
      <element name="updateDataResponse">
        <complexType/>
      </element>
    </schema>
  </wsdl:types>

  <wsdl:message name="addDataRequest">
    <wsdl:part element="impl:addData" name="parameters">
    </wsdl:part>
  </wsdl:message>

  <wsdl:message name="deleteDataResponse">
    <wsdl:part element="impl:deleteDataResponse" name="parameters">
    </wsdl:part>
  </wsdl:message>

  <wsdl:message name="updateDataResponse">
    <wsdl:part element="impl:updateDataResponse" name="parameters">
    </wsdl:part>
  </wsdl:message>

  <wsdl:message name="deleteDataRequest">
    <wsdl:part element="impl:deleteData" name="parameters">
    </wsdl:part>
  </wsdl:message>

  <wsdl:message name="addDataResponse">
```

```
<wsdl:part element="impl:addDataResponse" name="parameters">
</wsdl:part>
</wsdl:message>
<wsdl:message name="updateDataRequest">
  <wsdl:part element="impl:updateData" name="parameters">
  </wsdl:part>
</wsdl:message>
<wsdl:portType name="UCMDBDataReceiver">
  <wsdl:operation name="addData">
    <wsdlsoap:operation soapAction="addDataRequest"/>
    <wsdl:input message="impl:addDataRequest" name="addDataRequest">
    </wsdl:input>
    <wsdl:output message="impl:addDataResponse" name="addDataResponse">
    </wsdl:output>
  </wsdl:operation>
  <wsdl:operation name="deleteData">
    <wsdlsoap:operation soapAction="deleteDataRequest"/>
    <wsdl:input message="impl:deleteDataRequest"
name="deleteDataRequest">
    </wsdl:input>
    <wsdl:output message="impl:deleteDataResponse"
name="deleteDataResponse">
    </wsdl:output>
  </wsdl:operation>
  <wsdl:operation name="updateData">
    <wsdlsoap:operation soapAction="updateDataRequest"/>
    <wsdl:input message="impl:updateDataRequest"
name="updateDataRequest">
    </wsdl:input>
    <wsdl:output message="impl:updateDataResponse"
name="updateDataResponse">
    </wsdl:output>
  </wsdl:operation>
```

```
</wsdl:portType>

<wsdl:binding name="UCMDBDataReceiverSoapBinding"
type="impl:UCMDBDataReceiver">

  <wsdlsoap:binding style="document"
transport="http://schemas.xmlsoap.org/soap/http" />

  <wsdl:operation name="addData">

    <wsdl:input name="addDataRequest">

      <wsdlsoap:body use="literal" />

    </wsdl:input>

    <wsdl:output name="addDataResponse">

      <wsdlsoap:body use="literal" />

    </wsdl:output>

  </wsdl:operation>

  <wsdl:operation name="deleteData">

    <wsdl:input name="deleteDataRequest">

      <wsdlsoap:body use="literal" />

    </wsdl:input>

    <wsdl:output name="deleteDataResponse">

      <wsdlsoap:body use="literal" />

    </wsdl:output>

  </wsdl:operation>

  <wsdl:operation name="updateData">

    <wsdl:input name="updateDataRequest">

      <wsdlsoap:body use="literal" />

    </wsdl:input>

    <wsdl:output name="updateDataResponse">

      <wsdlsoap:body use="literal" />

    </wsdl:output>

  </wsdl:operation>

</wsdl:binding>

<wsdl:service name="UCMDBDataReceiverService">

  <wsdl:port binding="impl:UCMDBDataReceiverSoapBinding"
name="UCMDBDataReceiver">
```

```
<wsdlsoap:address
  location="http://localhost:8080/UCMDBDataReceiver/services/
    UCMDBDataReceiver"/>

</wsdl:port>

</wsdl:service>

</wsdl:definitions>
```

Response Handling

The data receiver should return a string in the **addDataResponse**, **deleteDataResponse**, or **updateDataResponse** structures. The adapter passes the response data unprocessed to the probe's **probeMgr-adaptersDebug.log**. The receiver can return any string data, and the responses are wrapped in SOAP-compliant XML. In the Jython script you can use the **SOAPMessage** and related Java classes to parse the response messages. The following is an example of a response message from the data receiver:

```
<2012-03-16 15:47:38,080> [INFO ] [Thread-110] - XMLtoWebService.py:addData
received response:

<soapenv:Body xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/">

<intf:addDataResponse xmlns:intf="http://ucmdb.hp.com">

  <xml>&lt;result&gt;&lt;status&gt;error&lt;/status&gt;
    &lt;message&gt;Error publishing config item changes&lt;/message&gt;
    &lt;/result&gt;</xml>

</intf:addDataResponse>

</soapenv:Body>
```

The message shown is an error message **<Error publishing config item changes>**, but the content can be anything that the data receiver is designed to respond with. The response is an error message simply because that is the intent, because the designer says it is an error message and the push adapter expects the response to be some indication of success or failure. The content can be reconciliation IDs of all the successfully added CIs, or error messages for specific CIs. Customization of the GWSPA could include parsing the response message and taking actions such as resending certain CIs or performing other logging.

Testing the WSDL

The SOAPUI Eclipse plug-in is used to test web service layers during development. You can use SOAPUI to assist with customization of a web service. SOAPUI offers an integrated development environment (IDE) to test building, sending, and receiving of SOAP messages. In the SOAPUI perspective, the WSDL on pages [194-198](#) generated the following sample message:

```
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
```

```
xmlns:ucm="http://ucmdb.hp.com">

  <soapenv:Header/>

  <soapenv:Body>

    <ucm:addData>

      <ucm:xmlAdded>?</ucm:xmlAdded>

    </ucm:addData>


  </soapenv:Body>

</soapenv:Envelope>
```

The “?” in the **xmlAdded** element above is the location of the data, which is supplied by the web service push adapter integration.

Observing Results

When the push adapter is operating normally, in non-debug mode, the data is never written to a file until the final result is written (the intermediate TQL results and mapped data results are not normally visible in any log file). However, the results can be written to the probe's debug file by un-commenting the **logger.debug** statements (remove the “#” character) in the DiscoveryMain section as shown here:



```
# get referenced data - unused in this adapter implementa
#addRefResult = Framework.getTriggerCIData('referencedAdd
#updateRefResult = Framework.getTriggerCIData('referenced
#deleteRefResult = Framework.getTriggerCIData('referenced

# uncomment out the logger statements to see the data
empty = isEmpty(addResult, "addResult")
if not empty:
    addResult = cleanUp(addResult)
    # send to ESB web service
    logger.info(SCRIP_NAME+":sending addData Result")
    sendDataToReceiver("add", URL, addResult)
    #logger.debug(addResult)

empty = isEmpty(updateResult, "updateResult")
if not empty:
    updateResult = cleanUp(updateResult)
```

Ensure the logger statement begins on the same column as the other preceding and following lines. Jython is indent-sensitive and the script will fail if the indentation of all lines is not correct.

The debug log file **probeMgr-adaptersDebug.log** on the probe here shows the contents of the output:

```
<2011-12-07 14:02:23,019> [INFO ] [Thread-273] - XMLtoWebService.py started
<2011-12-07 14:02:23,019> [DEBUG] [Thread-273] - ESB Push parameters:
<2011-12-07 14:02:23,019> [DEBUG] [Thread-273] - Wshost=harpy.trtc.com
<2011-12-07 14:02:23,019> [DEBUG] [Thread-273] - WShostport=5555
```

```
<2011-12-07 14:02:23,019> [DEBUG] [Thread-273] -  
WSuri=ws/DtITServiceManagement.esla.v1.ws.provider:UMDBDataReceiver  
  
<2011-12-07 14:02:23,019> [INFO ] [Thread-273] - URL is  
http://harpy.trtc.com:5555/ws/DtITServiceManagement.esla.v1.ws.  
provider:UMDBDataReceiver  
  
<2011-12-07 14:02:23,035> [DEBUG] [Thread-273] - Connected to  
http://harpy.trtc.com:5555/ws/DtITServiceManagement.esla.v1.ws.  
provider:UMDBDataReceiver  
  
<2011-12-07 14:02:23,035> [ERROR] [Thread-273] - sending results  
  
<2011-12-07 14:02:23,035> [DEBUG] [Thread-273] - <?xml version="1.0"  
encoding="UTF-8"?>  
  
<root>  
  <data>  
    <objects>  
      <Object mode="" name="u_imp_ip_switch" operation="add"  
mamId="9e8c2f6bdfe4b7d0864c79e70833902c">  
        <field name="Correlation ID" key="true" datatype="char"  
length="">9e8c2f6bdfe4b7d0864c79e70833902c</field>  
        <field name="name" key="false" datatype="char" length="">nma_  
09sw</field>  
        <field name="location" key="false" datatype="char" length="" />  
        <field name="u_chassis_vendor_type" key="false" datatype="char"  
length="">ciscoCat2960-24TT</field>  
        <field name="serial_number" key="false" datatype="char"  
length="" />  
        <field name="ram" key="false" datatype="char" length="" />  
        <field name="os_version" key="false" datatype="char" length=""  
/>  
      </Object>  
    </objects>  
  </data>  
</root>
```

Modifying the Jython Script

XMLtoWebService.py

The Jython script used by the Web Service push adapter is very similar to the XML push adapter. The script uses **UCMDBDataReceiver.jar**, included with the adapter. The script implements the **SendDataToReceiver()** method. **SendDataToReceiver()** uses three parameters:

1. Action (add, update, or delete)
2. The URL of the Data Receiver
3. The data

For example, the add block looks like: **SendDataToReceiver("add", URL, addResult)**

All web service and SOAP layers are wrapped. The URL is the service endpoint address of the UCMDB data receiver. This is the same URL used to obtain the wsdl via the "?wsdl" suffix.

The source of the Jython script is shown below. The web service integration wrapper lines are highlighted in green.

```
#####  
# script: XMLtoWebService.py  
#####  
  
# This jython script accepts TQL data results (adds, updates, and deletes) from  
# the Integration adapter.  
  
# and sends it to a web service. The web service is called UCMDBDataReceiver.  
  
# A web service client of this name must be addressable at the URL provided by  
# the parameters.  
  
# The SendDataToReceiver.jar exposes the SendDataToReceiver function, as well as  
# the service locator.  
  
# examples of the service locator are in the testconnection section.  
  
# regular expressions  
import re  
  
# logging  
import logger  
  
# web service interface  
from com.hp.ucmdb import SendDataToReceiver  
from com.hp.ucmdb.SendDataToReceiver import locateService  
from com.hp.ucmdb.SendDataToReceiver import SendData  
  
#####  
#####      VARIABLES      #####  
#####  
  
SCRIPT_NAME = "XMLtoWebService.py"  
logger.info(SCRIPT_NAME+" started")
```

```
def cleanUp(str):

    # replace mode=""
    str = re.sub("mode=\"\w+\\"s+", "", str)

    # replace mamId with id
    str = re.sub("\smamId=\"", " id=\"", str)

    # replace empty attributes
    str = re.sub("[\n|\\s|\\r]*<field name=\"\w+\" datatype=\"\w+\" />", "", str)

    # replace targetRelationshipClass with name
    str = re.sub("\stargetRelationshipClass=\"", " name=\"", str)

    # replace Object with object with name
    str = re.sub("<Object mode=\"", "<object mode=\"", str)
    str = re.sub("<Object operation=\"", "<object operation=\"", str)
    str = re.sub("<Object name=\"", "<object name=\"", str)
    str = re.sub("</Object>", "</object>", str)

    # replace field to attribute
    str = re.sub("<field name=\"", "<attribute name=\"", str)
    str = re.sub("</field>", "</attribute>", str)

    #logger.debug("String = %s" % str)
    #logger.debug("cleaned up")

    return str

def isEmpty(xml, type = ""):
    objectsEmpty = 0
    linksEmpty = 0
```

```
m = re.findall("<objects />", xml)

if m:
    #logger.warn("\t[%s] No objects found" % type)
    objectsEmpty = 1

m = re.findall("<links />", xml)

if m:
    #logger.warn("\t[%s] No links found" % type)
    linksEmpty = 1

if objectsEmpty and linksEmpty:
    return 1

return 0

#####
#####      MAIN      #####
#####

def DiscoveryMain(Framework):
    #fix this for web service export
    errMsg = "UCMDBDataReceiver Service not found."
    testConnection = Framework.getTriggerCIData("testConnection")
    # Get Web Service Push variables
    WHostName = Framework.getTriggerCIData("Host Name")
    WShostport = Framework.getTriggerCIData("Protocol Port")
    WSuri = Framework.getTriggerCIData("URI")

    logger.info(SCRIPT_NAME+":ESB Push parameters:")
    logger.info("Host Name="+WHostName)
    logger.info("Protocol Port="+WShostport)
    logger.info("URI="+WSuri)
    URL = "http://" + WHostName + ":" + WShostport + "/" + WSuri
    logger.info("URL="+URL)
    if testConnection == 'true':
```

```
# locate the service
test_receiver = SendDataToReceiver()
locator = test_receiver.locateService(URL)
#locator = locateService(URL)
if(locator):
    logger.info(SCRIPT_NAME+":Test connection was successful")
    return
else:
    raise Exception, errMsg
    return

# do same thing here if not just a test connection -
receiver = SendDataToReceiver()
locator = receiver.locateService(URL)
if(locator):
    logger.info(SCRIPT_NAME+":Connected to "+URL)
else:
    logger.error(SCRIPT_NAME+":no locator")
    raise Exception, errMsg
    return

# get add/update/delete result objects from the Framework
addResult = Framework.getTriggerCIData('addResult')
updateResult = Framework.getTriggerCIData('updateResult')
deleteResult = Framework.getTriggerCIData('deleteResult')
logger.debug(deleteResult)

# get referenced data - unused in this adapter implementation
#addRefResult = Framework.getTriggerCIData('referencedAddResult')
#updateRefResult = Framework.getTriggerCIData('referencedUpdateResult')
#deleteRefResult = Framework.getTriggerCIData('referencedDeleteResult')
# uncomment out the logger statements to see the data
empty = isEmpty(addResult, "addResult")
```

```
if not empty:
    addResult = cleanUp(addResult)
    # send to ESB web service
    logger.info(SCRIPT_NAME+":sending addData Result")
    rcvr = SendDataToReceiver()
    resp = rcvr.SendData("add", URL, addResult)
    logger.info(SCRIPT_NAME+":addData received response:"+resp)
    #logger.debug(addResult)
empty = isEmpty(updateResult, "updateResult")
if not empty:
    updateResult = cleanUp(updateResult)
    # send to ESB web service
    #logger.debug(updateResult)
    logger.info(SCRIPT_NAME+":sending updateData Result")
    rcvr = SendDataToReceiver()
    resp = rcvr.SendData("update", URL, updateResult)
    logger.info(SCRIPT_NAME+":received response:"+resp)

empty = isEmpty(deleteResult, "deleteResult")
if not empty:
    deleteResult = cleanUp(deleteResult)
    # send to ESB web service
    #logger.debug(deleteResult)
    logger.info(SCRIPT_NAME+":sending deleteData Result")
    rcvr = SendDataToReceiver()
    resp = rcvr.SendData("delete", URL, deleteResult)
    logger.info(SCRIPT_NAME+":received response:"+resp)
logger.info(SCRIPT_NAME+" ended")
```

Customizing Response Message Processing

The data receiver should return a string containing any response or status desired. The web service push adapter passes by default the response to the probe's info-level log. The response message is SOAP-formatted XML containing the returned response string(s) inside. Any data can be returned by the receiver such as grouped or individual error or success messages. If additional processing is

desired, the response can be processed by the adapter's Jython script. No Java programming is required.

An example of a return response message, sent using the following:

```
// stub example for building your own UCMDBDataReceiver
public class UCMDBDataReceiver {

    public String addData (String xmlAdd){
        System.out.println(xmlAdd); // do something with the data
        // send back a response message based on what you did
        String tr = new String("a test response from addData!");
        return tr;
    }
}
```

is shown here:

```
<soapenv:Body xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/">
    <addDataResponse xmlns="http://ucmdb.hp.com">
        <addDataReturn>a test response from addData!</addDataReturn>
    </addDataResponse>
</soapenv:Body>
```

Modifying the Data Receiver

A Java client can implement the classes contained in **UCMDBDataReceiver.jar** and call the web service in the same manner as Jython. In addition, the unwrapped methods may also be called. A Javadoc exists for the **UCMDBDataReceiver.jar** classes. The source code below shows how to use these essential methods to wrap the data in a SOAP message and send it to the receiver over HTTP.

The process is to create a **UCMDBDataReceiverServiceLocator** object, then assign the **UCMDBDataReceiverEndPointAddress** to the URL of the data receiver.

To send data, the locator's **getUCMDBDataReceiver** method is called to create a **UCMDBDataReceiver** object. The **UCMDBDataReceiver** object implements the methods to actually send the add/change/delete data. There are three identical code blocks to process each type of request.

The source code for the **SendDataToReceiver** class is listed below. Highlighted objects and methods are the essential elements to use.

```
/**
```

```
* Test SendData for the UCMDB Data Receiver for the UCMDB Web Service Push
Adapter
*/

package com.hp.ucmdb;

import com.hp.ucmdb.SendDataToReceiver;

/**
 * TestSendData can be used to verify the SOAP classes are working.
 * TestSendData creates a SendDataToReceiver class and invokes its SendData
 * method.
 * a response String is returned.
 * The test URL is typically appended with "?wsdl" to get the WSDL of the
 * service.
 */
public class TestSendData {

    /**
     * @param args - test SOAP message.
     * optional arguments [0] a test string [1] a service endpoint URL of a
     * Data Receiver.
     * the default URL is sent the incoming argument as a test message.
     * the default URL is
     * "http://localhost:8080/UCMDBDataReceiver/services/UCMDBDataReceiver".
     * If any errors are encountered, TestClient will attempt to throw
     * exceptions.
     */

    public static void main(String[] args) {

        // use test message if supplied, otherwise supply a default test string
        String teststring = new String("Test SOAP message from
        UCMDBDataReceiver TestSendData.");

        if(args.length > 0) {
            teststring = args[0];
        }

        // use test URL if supplied, otherwise supply the default URL
        String URL = new String("");
        if(args.length > 1) {
```

```
        URL = args[1];
    }
    // return response
    String response = new String("");
    // perform the tests
    try {
        if(URL.equals("")) {
            UCMDBDataReceiverServiceLocator locator = new
            UCMDBDataReceiverServiceLocator();

            UCMDBDataReceiver receiver = locator.getUCMDBDataReceiver();
            URL = locator.getUCMDBDataReceiverAddress();

            System.out.println("TestClient: tested
            URL="+locator.getUCMDBDataReceiverAddress());

            System.out.println("TestClient: receiver="+receiver.toString
            ());
        }

        SendDataToReceiver sdtr = new SendDataToReceiver();

        // this sends a test push and gets a response message
        response = sdtr.SendData("add", URL, args[0]);

        System.out.println("Response received was:"+response);
    } catch (Exception e) {
        System.out.println("TestClient: Remote Error:");
        e.printStackTrace();
    }
}
}
```

Source code is also included in the **UCMDBDataReceiver.jar** file for the other classes:

- TestClient.java
- UCMDBDataReceiver.java
- UCMDBDataReceiverProxy.java
- UCMDBDataReceiverService.java

- `UCMDBDataReceiverServiceLocator.java`
- `UCMDBDataReceiverSoapBindingStub.java`

The source was generated in the Eclipse IDE, then modified. Exercise caution when modifying the UCMDB code, as much of it is auto-generated to match the SOAP specification and the UCMDB data receiver.

Javadoc

A fully-commented **javadoc** is provided with the generic web service push adapter. The **javadoc** is included in the docs folder **javadoc**. Start with **index.html**. The overview page provides access to the documentation for all classes and methods in the SDK.

All Classes

- **SendDataToReceiver**: API for the web service wrapper
- **TestClient**: test client to verify connectivity to an service endpoint
- **UCMDBDataReceiver**: web service wrapper

The rest are automatically generated by the web service builder:

- `UCMDBDataReceiverProxy`
- `UCMDBDataReceiverService`
- `UCMDBDataReceiverServiceLocator`
- `UCMDBDataReceiverSoapBindingStub`

Overview

Basic usage of the SDK, including source code examples, is explained in the documentation in the package. This **javadoc** is for the UCMDB web service push adapter. The API may be called from Jython or Java.

The SDK provides two source samples, **TestClient** and **SendDataToReceiver**. **TestClient** provides a very limited test of the responding local client. **SendDataToReceiver** is the main class used to send data to a web service.

First, use this SDK (mainly the enclosed WSDL) to implement a UCMDB data receiver to communicate with this web service. Then use this SDK to create a push adapter in the UCMDB to push UCMDB TQL result data to the data receiver. Basic usage of this API is described below, with both Jython and Java implementations.

Implementing `SendDataToReceiver()`

SendDataToReceiver() wraps all functions with a single method:

- **Jython:** `SendDataToReceiver("add",yourURL,"Hello!")`
- **Java:** `SendDataToReceiver("add",yourURL,"Hello!");`

Or, create a **SendDataToReceiver** object (for example, to manipulate other settings) and then call the **SendData** method separately, as shown here:

- **Jython:**

```
rcvr = SendDataToReceiver()
responseMsg = rcvr.SendData("add", yourURL, "Hello!")
```

- **Java:**

```
SendDataToReceiver rcvr = new SendDataToReceiver();
String responseMsg = rcvr.SendData("add", yourURL, "Hello!");
```

Or, if you need to do it a step at a time, you can do the following:

1. Create a new **UCMDBDataReceiverServiceLocator()** object x, then set the object's endpoint address later, shown here:

- **Jython:**

```
x = UCMDBDataReceiverServiceLocator()
x.setUCMDBDataReceiverEndPointAddress(URL)
```

- **Java :**

```
UCMDBDataReceiverServiceLocator x = new UCMDBDataReceiverServiceLocator();
x.setUCMDBDataReceiverEndPointAddress(URL);
```

2. Then, create a **UCMDBDataReceiver** with

- **Jython:** `y = x.getUCMDBDataReceiver()`
- **Java:** `UCMDBDataReceiver y = x.getUCMDBDataReceiver();`

3. Then, send the data via the SOAP web service like this:

- **Jython:**

- `y.addData(yourData)`
- or `y.updateData(yourData)`
- or `y.deleteData(yourData)`

- **Java:**

- `y.addData(yourData);`
 - `or y.updateData(yourData);`
 - `or y.deleteData(yourData);`
4. It may be necessary to test connectivity, then if successful reuse the same locator object to return **UCMDBDataReceiver** to use for data transfer.

The classes contain no destructors and do not perform memory management.

Mapping File Reference

Using Mappings

A mapping must be created for each target attribute in the transformed XML output. The mappings specify where and how to obtain the data. If the data is in another corresponding attribute in the UC MDB, then a direct mapping is used.

To pull data from multiple attributes, or attributes from the UC MDB CI's child or parent CI's attributes, other complex mappings may be necessary. The mapping schema below shows all possible mappings.

The mapping file is an XML file which defines which CI/Relationship types in UC MDB are mapped to which CI/Relationship types in the target data store. The format is explained in detail below. The mapping file controls which CI and relationship types are pushed, as well as controlling exactly which attributes are pushed.

A mapping entry exists for each attribute to be pushed to the target MDR. Each mapping entry may consist of one or more attributes in the raw UC MDB push data. Mapping entries allow completely granular control of the final structure and naming of the data to be pushed to the target MDR.

Direct Mappings

Mappings transform one data model to another (in this case, the UC MDB to the push target MDR). Transformations may be simple, in the case of a 1:1 relationship between the UC MDB attribute and the target, they differ only by name and perhaps type.

Most attribute mappings are direct. For example, the server name "ServerX", may be represented in UC MDB as a CI of type **unix** with an attribute name of **primary_server_name**, of type **string** with a length of 50. The target MDR's data model may specify the same logical entity with a CI type of **linux**, with an attribute name of **hostname** with a type of **char[]** with a maximum length of 250. Direct mappings can accomplish all these aforementioned types of translation tasks.

Here is an example of a direct mapping:

```
<target_attribute name="dns_domain" datatype="char">  
<map type="direct" source_attribute="domain_name" />  
</target_attribute>
```

This direct mapping maps the UCMDB attribute **dns_domain** to the **domain_name** attribute in the target data model.

Use the **char** data type regardless of the actual data type, unless it is necessary to use the actual data type.

Complex Mappings

More complex mappings enable additional transformations:

- To map attribute values from multiple CIs to one target CI.
- To map attributes of children CIs (those having a **container_f** or contained relationship) to the parent CI in the target data store. For example, setting a value called **Number of CPUs** on a target Host CI. Another example could be setting the value **Total Memory** (by adding up the memory size values of all memory CIs of a host CI in UCMDB) on a target Host CI.
- To map attributes of parent CIs (those having a **container_f** or contained relationship) on the target data store's CI. For example, setting a value called **Container Server** on a target attribute called **Installed Software** CI, by getting the value from the containing host of the software CI in the UCMDB.

Below is an example of a complex mapping, using two source attributes separated by a comma character, to create the target attribute **os**:

```
<target_attribute name="os" datatype="char">  
  <map type="compoundstring">  
    <source_attribute name="discovered_os_name" />  
    <constant value="," />  
    <source_attribute name="host_osinstalltype" />  
  </map>  
</target_attribute>
```

Reversing Link Directions

It is possible that the UCMDB contains data that differs in structure from source to source. For example, the relationship between an IpAddress CI and an Interface CI may be a **parent**, as may occur with the HP Network Node Manager integration. Or it may be a **containment** link as is commonly created by Universal Discovery. Furthermore, the direction of these links are opposite to each other.

It is currently not possible to reverse the direction of links in the mappings file. Reversal of the **_end1** and **_end2** variables either switches the order of the data in the transformed XML or the link is missing in the source data.

One possible solution to this problem is to define an Enrichment rule as follows:

1. The enrichment's TQL part is a subset of a TQL that is used by the push adapter. This TQL in particular selects all the links that are in the opposite direction of what is desired in the transformed xml.
2. The enrichment part defines a new link of the correct direction and desired type.
3. Enrichment is activated and then creates the correct links.
4. The integration job TQL now refers to the enriched link rather than the original link.
5. The <link> mappings in the push adapter then refer to the enriched link as well and produce a set of links consistent in type and direction.

Mapping File Schema

Element Name and Path	Description	Attributes
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.	<ol style="list-style-type: none"> 1. Name: type Description: Name of the source data repository. Is required?: Required Type: String 2. Name: versions Description: Version(s) of the source data repositories. Is required?: Required Type: String 3. Name: vendor Description: Vendor of the source data repository. Is required?: Required Type: String

Element Name and Path	Description	Attributes
target (integration > info)	Defines information about the target data repository.	<ol style="list-style-type: none"> 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.	
source_ci_type_tree (integration > targetcis)	Defines a source CIT and all of the CI types which inherit from it.	<ol style="list-style-type: none"> Name: name Description: Name of the source CIT. Is required?: Required Type: String Name: mode Description: The type of update required for the current CI type. Is required?: Required Type: One of the following strings: <ol 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 Name and Path	Description	Attributes
source_ci_type (integration > targetcis)	Defines a source CIT without the CI types which inherit from it.	<ol style="list-style-type: none"> Name: name Description: Name of the source CIT. Is required?: Required Type: String Name: mode Description: The type of update required for the current CI type. Is required?: Required Type: One of the following strings: <ol 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.
target_ci_type (integration > targetcis > source_ci_type -OR- integration > targetcis > source_ci_type_tree)	Defines a target CIT.	<ol style="list-style-type: none"> 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 Name and Path	Description	Attributes
<p>targetprimarykey</p> <p>(integration > targetcis > source_ci_type)</p> <p>-OR-</p> <p>(integration > targetcis > source_ci_type_tree)</p> <p>-OR-</p> <p>(integration > targetrelations > link)</p> <p>-OR-</p> <p>(integration > targetrelations > source_link_type_tree)</p>	<p>Identifies target CIT primary key attributes.</p>	
<p>pkey</p> <p>(integration > targetcis > source_ci_type > targetprimarykey)</p> <p>-OR-</p> <p>integration > targetcis > source_ci_type_tree > targetprimarykey</p> <p>-OR-</p> <p>(integration > targetrelations > link > targetprimarykey)</p> <p>-OR-</p> <p>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 Name and Path	Description	Attributes
<p>target_attribute</p> <p>(integration > targetcis > source_ci_type</p> <p>-OR-</p> <p>integration > targetcis > source_ci_type_tree</p> <p>-OR-</p> <p>integration > targetrelations > link</p> <p>-OR-</p> <p>integration > targetrelations > source_link_type_tree)</p>	<p>Defines the target CIT's attribute.</p>	<ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> uppercase – Convert to uppercase lowercase – Convert to lowercase <p>If this attribute is empty, no conversion function will be applied.</p>

Element Name and Path	Description	Attributes
<p>map</p> <p>(integration > targetcis > source_ci_ type > target_attribute</p> <p>-OR-</p> <p>integration > targetcis > source_ci_type_tree > target_attribute)</p> <p>-OR-</p> <p>(integration > targetrelations > link > target_attribute</p> <p>-OR-</p> <p>integration > targetrelations > source_link_type_tree > target_attribute)</p>	<p>Specifies how to obtain the source CIT's attribute value.</p>	<ol style="list-style-type: none"> Name. type Description. The type of mapping between the source and target values. Is required. Required Type. One of the following strings: a. direct – Specifies a 1-to-1 mapping from source attribute's value to target attribute's value. b. compoundstring – Sub-elements are joined into a single string and the target attribute value is set. c. childattr – Sub-elements are one or more child CIT's attributes. Child CITs are defined as those with composition or containment relationship. d. constant – Static string Name. value Description. Constant string for type=constant Is required. Only required when type=constant Type. String Name. attr Description. Source attribute name for type=direct Is required. Only required when type=direct Type. String

Element Name and Path	Description	Attributes
<p>aggregation</p> <p>(integration > targetcis > source_ci_type > target_attribute > map</p> <p>-OR-</p> <p>integration > targetcis > source_ci_type_tree > target_attribute > map</p> <p>-OR-</p> <p>(integration > targetrelations > link > target_attribute > map</p> <p>-OR-</p> <p>integration > targetrelations > source_link_type_tree > target_attribute > map)</p> <p>Only valid when the map's type is childattr</p>	<p>Specifies how the source CI's child CI attribute values are combined into a single value to map to the target CI attribute. Optional.</p>	<p>Name: type Description. The type of aggregation function Is required?: Required Type. One of the following strings:</p> <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 the sum of all numeric 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 Name and Path	Description	Attributes
<p>source_child_ci_type</p> <p>(integration > targetcis > source_ci_type > target_attribute > map</p> <p>-OR-</p> <p>integration > targetcis > source_ci_type_tree > target_attribute > map</p> <p>-OR-</p> <p>(integration > targetrelations > link > target_attribute > map</p> <p>-OR-</p> <p>integration > targetrelations > source_link_type_tree > target_attribute > map)</p> <p>Only valid when the map's type is childattr.</p>	<p>Specifies from which connected CI the child attribute is taken.</p>	<ol style="list-style-type: none"> Name. name Description. The type of the child CI Is required. Required Type. String Name. source_attribute Description. The attribute of the child CI that is mapped. Is required. Required only if the childAttr aggregation type (which is on the same path) is not =count. Type. String

Element Name and Path	Description	Attributes
<p>validation</p> <p>(integration > targetcis > source_ci_type > target_attribute > map</p> <p>-OR-</p> <p>integration > targetcis > source_ci_type_tree > target_attribute > map</p> <p>-OR-</p> <p>(integration > targetrelations > link > target_attribute > map</p> <p>-OR-</p> <p>integration > targetrelations > source_link_type_tree > target_attribute > map)</p> <p>Only valid when the map's type is childatt</p>	<p>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.</p>	<ol style="list-style-type: none"> 1. Name. minlength Description. Excludes strings shorter than the given value. Is Required?: Not required Type. Integer 2. Name. maxlength Description. Excludes strings longer than the given value. Is Required?: Not required Type. Integer 3. Name. minvalue Description. Excludes numbers smaller than the specified value. Is Required?: Not required Type. Numeric 4. Name. maxvalue Description. Excludes numbers greater than the specified value. Is Required?: Not required Type. Numeric
<p>targetrelations</p> <p>(integration)</p>	<p>Container element for all relationship mappings. Optional.</p>	

Element Name and Path	Description	Attributes
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.	<ol style="list-style-type: none"> Name: name Description: Source relationship name. Is required?: Required Type: String Name: target_link_type Description: Target relationship name Is required?: Required Type: String Name: nameSpace Description: The namespace for the link that will be created on the target. Is required?: Not required Type: String Name: mode Description: The type of update required for the current link. Is required?: Required Type: One of the following strings: <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. 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 Name and Path	Description	Attributes
link (integration > targetrelations)	Maps a source Relationship to a target Relationship. Mandatory only if targetrelation is present.	<ol style="list-style-type: none"> Name: source_link_type Description: Source relationship name. Is Required?: Required Type: String Name: target_link_type Description: Target relationship name. Is required?: Required Type: String Name: nameSpace Description: The namespace for the link that will be created on the target. Is required?: Not required Type: String Name: mode Description: The type of update required for the current link. Is required?: Required Type: On the following strings: <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. 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 Name and Path	Description	Attributes
target_ci_type_end1 (integration > targetrelations > link -OR- integration > targetrelations > source_link_type_tree)	Target relationship's End1 CI type.	<ol style="list-style-type: none"> 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
target_ci_type_end2 (integration > targetrelations > link -OR- integration > targetrelations > source_link_type_tree)	Target relationship's End2 CI type.	<ol style="list-style-type: none"> Name: name Description: Name of the target relationship's End2 CI type. Is required?: Required Type: String Name: superclass Description: Name of the End2 CI type's super-class. Is required?: Not required Type: String

Mapping Results Schema

Element Name and Path	Description	Attributes
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 Name and Path	Description	Attributes
Object (root > data > objects)	Describes the update operation for a single object and all of its attributes.	<ol style="list-style-type: none"> Name: name Description: Name of the CI type Is required?: Required Type: String Name: mode Description: The type of update required for the current CI type. Is required?: Required Type: One of the following strings: <ol 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. Name: operation Description: The operation to perform with this CI. Is required: Required Type: One of the following strings: <ol style="list-style-type: none"> add – The CI should be added update – The CI should be updated delete – The CI should be deleted If no value is set, then the default value of add is used. Name: mamId Description: The ID of the object on the source RTSM. Is required?: Required Type: String

Element Name and Path	Description	Attributes
<p>field (root > data > objects > Object -OR- root > data > links > link)</p>	<p>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>).</p>	<ol style="list-style-type: none"> 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 Name and Path	Description	Attributes
links (root > data)	The root element for the links to update.	<ol style="list-style-type: none"> Name: targetRelationshipClass Description: The name of the relationship (link) in the target system. Is required?: Required Type: String Name: targetParent Description: The type of first end of the link (parent). Is required?: Required Type: String Name: targetChild Description: The type of the second end of the link (child). Is required?: Required Type: String Name: mode Description: The type of update required for the current CI type. Is required?: Required Type: One of the following strings: <ol 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. Name: operation Description: The operation to perform with this CI. Is required?: Required Type: One of the following strings: <ol style="list-style-type: none"> add – The CI should be added update – The CI should be updated delete – The CI should be deleted If no value is set, then the default value of add is used. Name: mamId Description: The ID of the object on the source RTSM. Is required?: Required Type: String

Customization

This section explains some of the basic procedures for common types of customization for push adapters.

Adding an Attribute

1. Ensure that the attribute is included in the TQL result.
2. Add the attribute mapping to the mappings file in the correct CI mapping section.
3. Ensure the data receiver is prepared to receive the additional attribute in the data.

Removing an Attribute

To remove an attribute, remove the attribute from the mapping file. You should also remove the attribute from the TQL if it is no longer used in the result or as a conditional node.

Adding a CI Type

1. Add the CI type to the TQL.
2. Ensure the CI type and its attribute data appear in the TQL result (use calculate and preview).
3. Add the CI type's mapping in the mappings file. Copy another CI type's mappings to quickly create a new CI type.
4. Modify the copied XML's name and attribute mappings to correspond to the new CI type and its attributes. See "[Mapping File Reference](#)" on [page 211](#) for the available types of mappings.

Removing a CI Type

1. Remove the CI Type from the TQL
2. Remove the mapping section for that CI type in the mappings file.

Adding Links

1. Ensure the two end CIs are present in the data.
2. Ensure the link you need to add is in fact a valid link (check in the CI type manager).
3. Add the link elements in the relationships section of the mappings xml.

Removing Links

1. Remove the link section of the link you want to remove in the mappings file.
2. If possible, remove the link from the TQL (unless it affects the efficiency or function of the TQL).

Chapter 6: Developing Enhanced Generic Push Adapters

This chapter includes:

Overview	229
The Mapping File	229
The Groovy Traveler	232
Write Groovy Scripts	235
Implement PushConnector Interface	236
Build an Adapter Package	237
Mapping File Schema	238

Overview

An enhanced push adapter is a data structure that represents the TQL query result. Each adapter built over the enhanced push adapter will handle this data structure and push it to its required target.

The data structure is named **ResultTreeNode (RTN)**. The RTN is created according to the mapping file of the adapter and the results of the TQL query. The queries used for the enhanced push adapter must be root based, that is the query must contain one query node with element name **root**, or one or more relationship elements beginning with the prefix **root**. This CI or relationship serves as the root element of the query. For details, see [Data Push](#) in the *Data Flow Management Guide*.

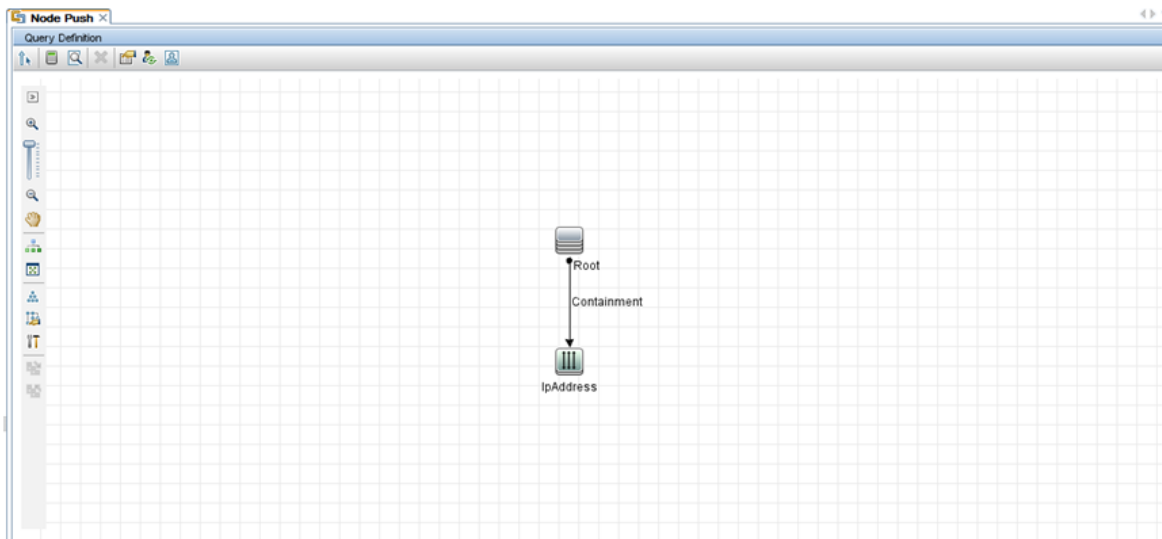
There are two basic steps involved in developing enhanced push adapters:

1. Implementing the PushConnector interface – this interface receives the data to add, update and delete as a list of RTNs, and perform the push into the target.
2. Creating the mapping file – the mapping file determines the creating of the RTN structure, by mapping CIs and attributes of the TQL result.

The Mapping File

The following example demonstrates how to create the mapping file.

In this example, we will simulate a push of a node and an IP address. We will create a TQL query called: **Node Push**, as follows:



In the mapping file we create two target CI types: **Computer** and **IP**. Computer has one variable and two attributes. IP has one attribute.

The following is the mapping XML file:

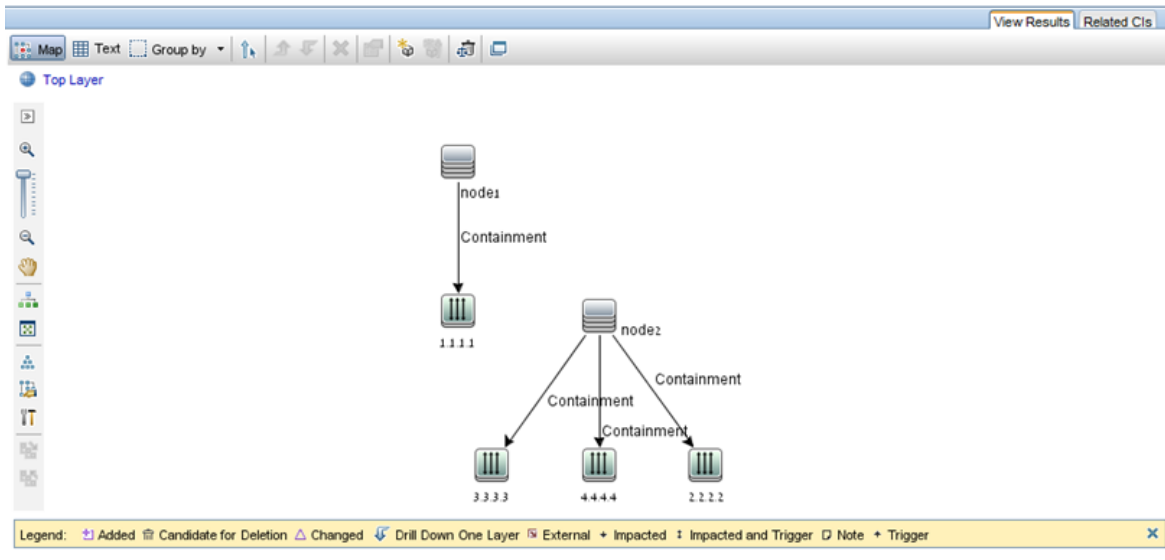
```

<integration xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="http://schemas.hp.com/ucmdb/1/pushAdapter">
  <info>
    <source name="UCMDB" versions="10.0" vendor="HP" />
    <target name="PushProduct" versions="9.3" vendor="HP" />
  </info>
  <import>
    <!--imports the Groovy script file. -->
    <scriptFile path="mappings.scripts.PushFunctions" />
  </import>
  <targetcis>
    <source_instance_type query-name="Node Push" root-element-name="Root">
      <target_ci_type name="Computer" is-
        valid="(Root['root_iscandidatefordeletion']==null) ? true :
        !Root['root_iscandidatefordeletion']">
        <variable name="isVirtual" datatype="BOOLEAN"
          value="PushFunctions.isVirtual(Root['root_class'])" />

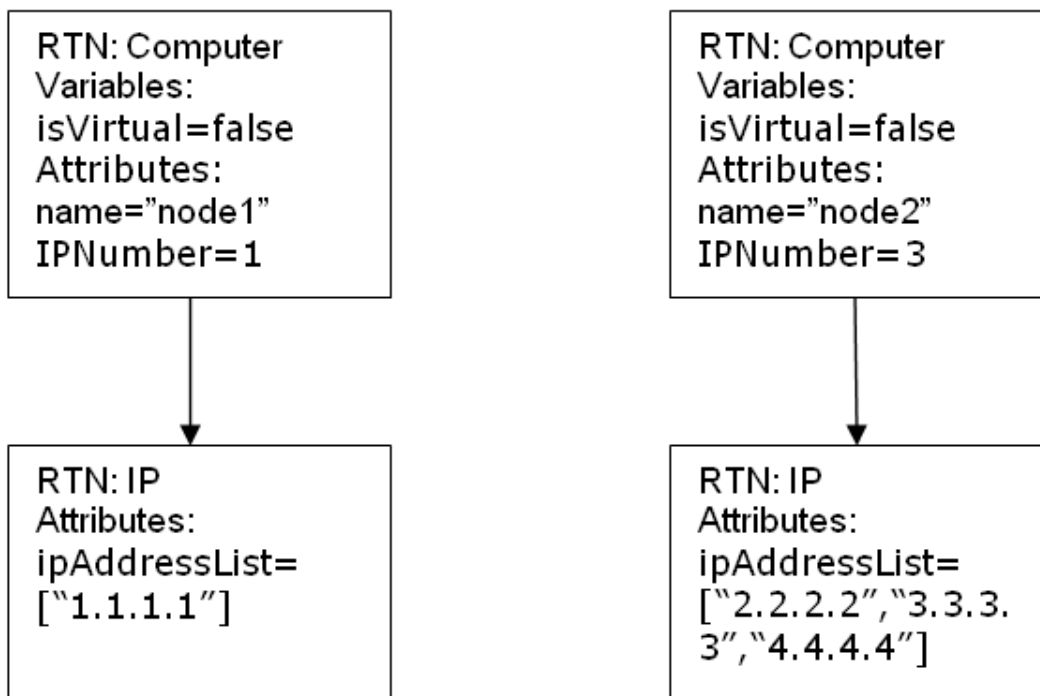
        <target_mapping name="name" datatype="String" value="Root['name']" />
        <target_mapping name="ipNumber" datatype="INTEGER"
          value="Root.IpAddress.size()" />
        <target_mapping name="description" datatype="STRING" value="PushFunctions
          .getDescription(isVirtual)" />
        <target_ci_type name="IP">
          <target_mapping name="ipAddressList" datatype="STRING_LIST"
            value="Root.IpAddress*.getAt('name') " />
        </target_ci_type>
      </target_ci_type>
    </targetcis>
  </integration>

```

The query results appear as follows:



Here is the RTN list built according to this mapping file:



Each root instance is mapped separately using the mapping file. Thus in this example, the PushConnector receives a list of two RTN roots.

Note: The previous push adapter had the ability to create a general mapping for a CI type. The new push adapter mapping is per TQL query. While running a push job that uses a query named **x**, the adapter looks for the relevant mapping file (the one that has attribute: query-name=x).

You can calculate the values in the mapping file using groovy script language. For details, see ["The Groovy Traveler" below](#).

The Groovy Traveler

Access the TQL query results in the following manner:

- **Root[*attr*]** returns the attribute ***attr*** of the Root element.
- **Root.Query_Element_Name** returns a list of the CI instances named in the TQL Query_Element_Name and that are linked to the current root CI.
- **Root.Query_Element_Name[2][*attr*]** returns the attribute ***attr*** of the third Query_Element_Name that is linked to the current root CI.
- **Root.Query_Element_Name*.getAt(*attr*)** returns a list of the attributes ***attr*** of the CI instances named Query_Element_Name in the TQL and that are linked to the current root CI.

There are additional attributes that can be accessed by the groovy traveler:

- **cmdb_id** – returns the UCMDB ID of the CI or relationship as a string.
- **external_cmdb_id** – returns the external ID of the CI or relationship as a string.
- **Element_type** – returns the element type of the CI or relationship as a string.

The import tag:

```
<import>
<scriptFile path="mappings.scripts.PushFunctions"/>
</import>
```

This means that we are declaring an import for all the groovy scripts in the mapping file. In this example, **PushFunctions** is a groovy script file that contains some static functions, and we can access them during the mapping (i.e. value="PushFunctions.foo()")

source_instance_type

The mapping is done per TQL, the query-name value is the related TQL of the current mapping. The '*' means that this mapping file is associated with all TQL queries beginning with the prefix: **Node Push**.

```
<source_instance_type query-name="Node Push*" root-element-name="Root">
```

The source_instance_type tag designates the root element we are mapping.

root-element-name should be exactly the same as the name of the root in the TQL.

target_ci_type

This tag is used for the creation of the RTN.

The name attribute represents the target_ci_type name: name=Computer

The **is-valid** attribute is a Boolean value that is calculated during the mapping, and determines if the current target_ci is valid. Invalid target_ci_types are not added to the RTN. In this example, we do not want to create a target_ci_type instance for which the **root_iscandidatefordeletion** attribute in UCMDB is true.

The target_ci_type can have variables that are calculated during the mapping:

```
<variable name="vSerialNo" datatype="STRING" value="Root['serial_number']"/>
```

The variable **vSerialNo** gets the value of the **serial_number** of the current root.

The attribute of the RTN is created by the **target_mapping** tag. The result of the execution of the groovy script in the **value** field, is assigned to the RTN attribute.

```
<target_mapping name="SerialNo" datatype="STRING" value="vSerialNo"/>
```

SerialNo assigns the value of the variable **vSerialNo**.

It is possible to define a target_ci_type as child of another target_ci_type as follows:

```
<target_ci_type name="Portfolio">
  <variable name="vSerialNo" datatype="STRING" value="Root['global_id']"/>
  <target_mapping name="CMDBId" datatype="STRING" value="globalId"/>
  <target_ci_type name="Asset">
    <target_mapping name="SerialNo" datatype="STRING" value="vSerialNo"/>
  </target_ci_type>
</target_ci_type>
```

The RTN **Portfolio** will have child RTN named **Asset**.

for-each-source-ci

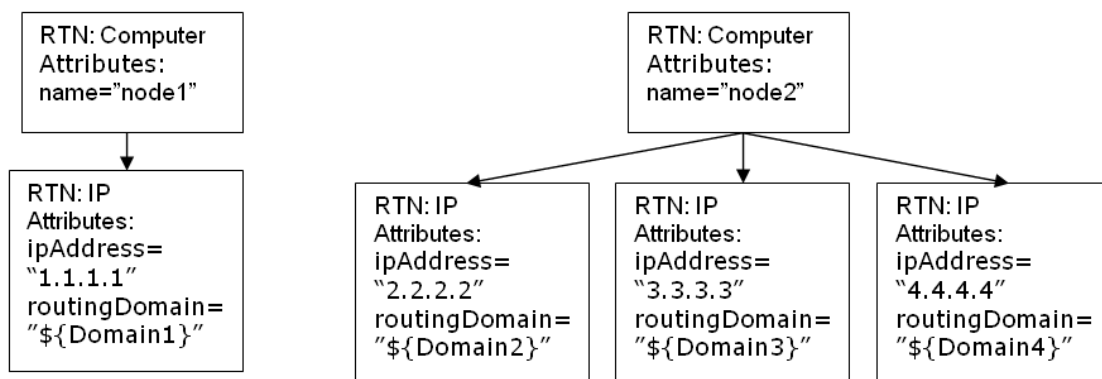
This tag lists the specific CIs of the root instance. It has the following fields:

- source-cis="" – the list of CIs for which a target CI is created. This list is defined by the groovy traveler in the **Root.Address** field.
- count-index="" – a variable that holds the index of the CI in the current iteration of the loop.
- var-name="" – the name of the CI in the current iteration of the loop.

Let's modify our example mapping file:

```
<target_ci_type name="Computer">
  <target_mapping name="name" datatype="String" value="Root['name']" />
  <for-each-source-ci count-index="i" var-name="currIP" source-cis="Root.IpAddress">
    <target_ci_type name="IP">
      <target_mapping name="ipAddress" datatype="STRING"
        value="Root.IpAddress[i]['name']"/>
      <target_mapping name="routingDomain" datatype="STRING"
        value="currIP['routing_domain']"/>
    </target_ci_type>
  </for-each-source-ci>
</target_ci_type>
```

The RTN list that will build according to this mapping file will look like this:



dynamic_mapping

This tag adds the ability to create a mapping of data from the target data store during the creation of the RTN structure.

Example: Assume that the target is a database with a table named **Computer** that has an **id** column and a **name** column that is correlated to **Node.name** in UCMDB. Both columns are unique. The database also has a table named **IP** that has a referenced key to the **parentID** in the Computer table. The 'dynamic_mapping' can create a map that stores the name and id as <name,id>. Based on this map, the Adapter can match ids with computers and can push the correct value to the parentID attribute in the IP table. You can use this map to assign a value to the parentID attribute while creating the RTN.

The mapping is determined by the **map_property**. The dynamic_mapping is executed once for each chunk.

```
<dynamic_mapping name="IdByName " keys-unique="true">
```

The **name** attribute represents the name of the map. The **keys-unique** attribute indicates if the keys are unique (each key is mapped to one value, or to a set of values).

The name of the map in this example is **IdByName** and it has unique keys. In order to access the map in the script, execute the following command:

```
DynamicMapHolder.getMap('IdByName')
```

It returns a reference to that map.

The **map_property** tag creates the property on which the mapping is based.

Example:

```
<map_property property-name="SQLQuery" datatype="STRING"  
property-value="SELECT name, id FROM Computer"/>
```

In this example the name of the property is **SQLQuery** and its value is an SQL statement that creates the map. The implementation of the methods **retrieveUniqueMapping** and **retrieveNonUniqueMapping** for the PushConnector interface will determine the actual content of the returned map.

Global Variables

The following global variables are accessible to the groovy script in the mapping file:

- **Topology** – Type: Topology. An instance of the topology of the current chunk.
- **QueryDefinition** - Type: QueryDefinition. An instance of the query definition of the current TQL.
- **OutputCI** – Type: ResultTreeNode. The RTN of the root element in the current tree mapping.
- **ClassModel** – Type: ClassModel. An instance of the class model.
- **CustomerInformation** – Type: CustomerInformation. Information about the customer running the job.
- **Logger** – Type: DataAdapterLogger. This logger is available in the adapter for writing logs to the UCMDB logging framework.

Write Groovy Scripts

In this section we create the **PushFunctions.groovy** file. This file will contain static functions that are used during the mapping of the root instance.

```
package mappings.scripts  
  
public class PushFunctions {  
  
    public static boolean isVirtual(def nodeRole){  
        return isListContainsOne(def list, "MY_VM", "MY_SIMULATOR");  
    }  
}
```

```
public static String getDescription(boolean isVirtual){
    if(isVirtual){
        return "This is a VM";
    }
    else{
        return "This is physical machine";
    }
}

private static boolean isListContainsOne(def list, ...stringList){
    //returns true if the list contains one of the values.
}
}
```

Implement PushConnector Interface

The implementation should support the following basic steps:

```
public class PushExampleAdapter implements PushAdapterConnector
{

    public UpdateResult pushTreeNodes(PushConnectorInput input) throws
    DataAccessException{

        // 1. build an UpdateResult instance - the UpdateResult is used to return
        mappings between the sent ids to the actual ids that entered the data store.
        // Also has an update status which allows to pass the status of data that was
        actually pushed, detailed status reports on failed IDs, and actions actually
        performed on successful ids.
        // 2. handle the data:
        // a. handle data to add. Can be retrieved by:
        input.getResultTreeNodes.getDataToAdd();
        // b. handle data to update.
        // c. handle data to delete.
        // 3. Return the Update result.
    }

    public void start(PushDataAdapterEnvironment env) throws DataAccessException{
        // this method is called when the integration point created,
        or when the adapter is reloaded
        //(i.e after changing one of the mapping files
        // and pressing 'save').
    }
}
```

```
public void testConnection(PushDataAdapterEnvironment env) throws
DataAccessException {
    // this method is called when pressing the 'test
connection' button in the
    //creation of the integration point.
    // For example if we push data to RDBMS this method
can create a connection
    //to the database and will run a dummy SQL statement.
    // If it fails it writes an error message to the log
and throws an exception.
}

Map<Object, Object> retrieveUniqueMapping(MappingQuery mappingQuery){
//This method will create the map according to the given mappingQuery. It will
be called in the
// mapping stage of the adapter execution, before the 'UpdateResult
pushTreeNodes' method.
// This method is called when the 'keys-unique' attribute of the 'dynamic_
mapping' tag is true.
}

Map<Object, Set<Object>> retrieveNonUniqueMapping(MappingQuery mappingQuery){
// This method is called when the 'keys-unique' attribute of the 'dynamic_
mapping' tag is false.
// In this case a key can be mapped to several values.
}
}
```

Build an Adapter Package

Ensure that the adapter package contains the following folders:

- **adapterCode**. Under this folder, create a folder named **PushExampleAdapter**, which will contain the .jar file we created from the PushExampleAdapter.java. It will also contain a folder named **mappings**, where you can place the mapping file created earlier, **computerIPMapping.xml**. It should also contain another folder named **scripts** that contains the **PushFunctions.groovy** file.
- **discoveryConfigFiles**. To contain configuration files such as the error codes used when reporting an error using UpdateResult. In this example, the folder is empty.
- **discoveryPatterns**. To contain the **push_example_adapter.xml**.

- **tql.** To contain the TQL query created for the example. This folder is optional, but when the package is deployed, the TQL is automatically created.

Mapping File Schema

Element Name and Path	Description	Attributes
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 (info)	The source product	Name: name Description: the name of the product Is required?: required Type: String
		Name: vendor Description: the vendor of the product Is required?: required Type: String
		Name: versions Description: the version of the product Is required?: required Type: decimal

Element Name and Path	Description	Attributes
target (info)	The target product	Name: name Description: the name of the product Is required?: required Type: String
		Name: vendor Description: the vendor of the product Is required?: required Type: String
		Name: versions Description: the version of the product Is required?: required Type: decimal
Import (integration)	A container element for imported script files.	
scriptFile (integration>import)	Defines the groovy script file to be imported.	Name: path Description: the path of the script file Is required?: required Type: string
Targetcis (integration)	A container element for Target CI types.	

Element Name and Path	Description	Attributes
Source_instance_type (integration>targetcis)	Defines the source CI instance type. Must be the root element as defined in the TQL.	Name: query-name. Description: The name of the relevant TQL Is required?: required Type: String
		Name: name Description: The root element as defined in the TQL. Is required?: required Type: String
dynamic_mapping (integration>targetcis>source_instance_type)	Defines a map that is created once per chunk.	Name: name Description: the name of the map Is required?: required Type: String
		Name: keys-unique Description: States if the keys are unique or not. Is required?: required Type: boolean

Element Name and Path	Description	Attributes
target_ci_type (integration>targetcis> source_instance_type) -OR- (integration>targetcis> source_instance_type> for- each-source-ci)	Defines the target CI type to be added to the RTN.	Name: name Description: the name of the target CI type Is required?: required Type: String
		Name: is-valid Description: checks if the current target CI valid according to the given script. Is required?: not required Type: String (groovy script)

Element Name and Path	Description	Attributes
Variable (target_ci_type)	Defines a variable for the target CI type.	Name: name Description: the name of the variable Is required?: required Type: String
		Name: datatype Description: the data type of the variable. Is required?: required Type: type-enum (can be one of the following: INTEGER, LONG, FLOAT, DOUBLE, STRING, BYTES, XML, BOOLEAN, DATE, INTEGER_LIST, STRING_LIST)
		Name: value Description: The value to assign to the variable Is required?: required Type: groovy script

Element Name and Path	Description	Attributes
target_mapping (target_ci_type)	Defines an attribute for the target CI type.	Name: name Description: the name of the attribute Is required?: required Type: String
		Name: datatype Description: the data type of the variable. Is required?: required Type: type-enum (can be one of the following: INTEGER, LONG, FLOAT, DOUBLE, STRING, BYTES, XML, BOOLEAN, DATE, INTEGER_LIST, STRING_LIST)
		Name: value Description: The value to assign to the variable Is required?: required Type: groovy script
		Name: ignore-on-null Description: If the target mapping value is null and this attribute is TRUE, ignore the attribute Is required?: not required Type: Boolean (groovy script)
before-mapping (target_ci_type)	A groovy script that is executed before the mapping of the target CI type.	
after-mapping (target_ci_type)	A groovy script that is executed after the mapping of the target CI type.	

Element Name and Path	Description	Attributes
for-each-source-ci (target_ci_type)	Defines an iteration of specific CIs of the root instance.	Name: count-index Description: the index of the current iterated CI Is required?: required Type: String
		Name: var-name Description: the variable that refers to the current iterated CI. Is required?: not required Type: String
		Name: source-cis Description: The CI name in the query that is to be iterated Is required?: required Type: String (groovy script)

Chapter 7: Viewing KPIs in External Applications

This chapter includes:

Set Up an Adapter to View KPIs in an External Application	245
---	-----

Set Up an Adapter to View KPIs in an External Application

This chapter explains how to retrieve the status and value of HP Operations Manager i 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" below](#)
- ["View the Attributes in the CIT Manager" on the next page](#)
- ["Define a TQL" on the next page](#)
- ["Consume KPIs in External Application" on the next page](#)

1. Create an Adapter Instance

In HP Operations Manager i, in the integration point, create an adapter instance named **BACKPIsAdapter**. Use the following definitions:

- **Adapter.** Enter **BACKPIsAdapter**.
- **Is Integration Activated.** This check box 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 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 Operations Manager i.

The following attributes are used by the **KpiObjective** CIT:

- **Origin.** The origin of the KPIs, for example, HP Operations Manager i.
- **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 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 ["Define a TQL Query"](#) in the Modeling Guide.

4. Consume KPIs in External Application

Using the API of your external application, send a query to run the relevant TQL in HP Operations Manager i, to retrieve the status and value of the federated KPIs.

Part 2: Using APIs

Chapter 8: Introduction to APIs

This chapter includes:

APIs Overview	248
---------------------	-----

APIs Overview

The following APIs are included with HP Operations Manager i:

- **UCMDB Java API.** Explains how third-party or custom tools can use the Java API to extract data and calculations and to write data to the HP Operations Manager iRun-time Service Model (RTSM). For details, see ["RTSM API" on page 249](#).
- **RTSM Web Service API.** Enables writing configuration item definitions and topological relations to the HP Operations Manager iRun-time Service Model (RTSM), and querying the information with TQL and ad hoc queries. For details, see ["RTSM Web Service API" on page 258](#).
- **Data Flow Management Java API.** Enables managing probes, jobs, triggers and credentials for Data Flow Management. For details, see ["Data Flow Management Java API" on page 296](#).
- **Data Flow Management Web Service API.** Enables managing probes, jobs, triggers and credentials for Data Flow Management. For details, see ["Data Flow Management Web Service API" on page 298](#).

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.

Chapter 9: RTSM API

This chapter includes:

Conventions	249
Using the RTSM API	249
General Structure of an Application	251
Put the API Jar File in the Classpath	253
Create an Integration User	253
UCMDB API Use Cases	255
Examples	256

Conventions

Note: The RTSM Web Service API can be used with the HP Operations Manager i 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 [RTSM API Reference](#).

These files are located in the following folder:

`\\<BSM Gateway Server root directory>\AppServer\
webapps\site.war\amdocs\eng\APIs\UCMDB_JavaAPI \index.html`

Using the RTSM API

Note: Use this chapter in conjunction with the API Javadoc, available in the online Documentation Library.

The RTSM 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 Modeling Guide.

Users of the RTSM API should be familiar with:

- The Java programming language
- RTSM

This section includes the following topics:

- ["Uses of the API" below](#)
- ["Permissions" below](#)

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 ["UCMDB API Use Cases" on page 255](#).

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.

In addition, the user must have the **Access to SDK** general action permission in order to log in.

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 ["Create an Integration User" on page 253](#).

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
<code>getAuthorizationModelService</code>	Perform authorization operations (create users and user groups, assign roles to users and groups, and so on).
<code>getClassModelService</code>	Information about types of CIs and relations
<code>getConfigurationService</code>	Infrastructure settings management, for server configuration
<code>getDataStoreMgmtService</code>	Query data store information including which CIs and attributes will be federated.
<code>getDDMConfigurationService</code>	Configure the Data Flow Management system
<code>getDDMManagementService</code>	Analyze and view the progress, results, and errors of the Data Flow Management system

Service Methods	Use
getDDMZoneService	Import and export management zones (with their activities).
getHistoryService	Information about history of monitored CIs (changes, removals, and so on)
getImpactAnalysisService	Run impact analysis scenario (also known as correlation).
getLicensingService	Query information about licenses installed in the system.
getMultipleCMDBService	Convert between global IDs and UCMDB IDs.
getMultiTenancyService	Create, read, update, and delete tenants.
getPersistencyService	Persist binary data into key-value pairs.
getQueryManagementService	Manage access to queries - save, delete, list existing. Also provides query validation and discovery of queries dependencies.
getReconciliationService	Supplies identification and merging capabilities.
getResourceBundleManagementService	Resource tagging ("bundling" services). Allows explicit creation of new tags and removal of tags from all tagged resources.
getResourceManagementService	Deploy resource packages (of TQL queries, views, users, and so on) to the system.
getSecurityService	Verify whether credentials are valid.
getServerService	Query generic information about the system.
getSnapshotService	Provide services for managing snapshots (get, save, compare, and so on)
getSoftwareSignatureService	Define software items to be discovered by the Data Flow Management system
getStateService	Provide services for managing states (list, add, remove, and so on)
getSystemHealthService	Provide system health services (basic system performance indicators, capacity and availability metrics)
getTopologyQueryService	Get information about the IT universe
getTopologyUpdateService	Change information in the IT universe
getUcmdbVersion	Query UCMDB and content pack versions and build information.

Service Methods	Use
getViewArchiveService	View result archiving services. Allows saving the current view result and retrieving previously saved results.
getViewService	View execution service (execute definition, execute saved) and management service (save, delete, list existing). Also provides view validation and dependencies discovery.

The client communicates with the server over HTTP(S).

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: Use of the UCMDB Java API Jar requires you to have JRE version 6 or later installed.

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 the user API authentication permission.

To create an integration user:

1. Launch the Web browser and enter the server address, as follows:

`http://localhost.<domain_name>:21212/jmx-console.`

You may have to log in with a user name and password.

2. Under UCMDB, click **service=UCMDB Authorization Services**.
3. Locate the **createUser** operation. This method accepts the following parameters:

- **customerId**. The customer ID.
- **username**. The integration user's name.
- **userDisplayName**. The integration user's display name.
- **userLoginName**. The integration user's login name.
- **password**. The integration user's password.

4. Click **Invoke**.
5. In a single-tenant environment, locate the **setRolesForUser** method and enter the following parameters:
 - **userName**. The integration user's name.
 - **roles**. SuperAdmin.

Click **Invoke**.

6. Either create more users, or close the JMX console.
7. Log on to BSM as an administrator.
8. From the **Administration** tab, run **Package Manager**.
9. Click the **Create custom package** icon.
10. Enter a name for the new package, and click **Next**.
11. In the Resource Selection tab, under **Settings**, click **Users**.
12. Select a user or users that you created using the JMX console.
13. Click **Next** and then **Finish**. Your new package appears in the Package Name list in Package Manager.
14. Deploy the package to the users who will run the API applications.

For details, see the section ["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 (**removeUser**, **resetPassword**, **UserAuthenticate**, and so on).

There are two out-of-the-box system users. It is recommended to change their passwords after installation using the **resetPassword** method.

- **sysadmin/sysadmin**

- **UISysadmin/UISysadmin** (this user is also the **SuperIntegrationUser**).

If you change the UISysadmin password using **resetPassword**, you must do the following:

- i. In the JMX Console, locate the **UCMDB-UI:name=UCMDB Integration** service.
- ii. Run **setCMDBSuperIntegrationUser** with the user name and new password of the integration user.

UCMDB API Use Cases

The use cases listed in this section 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" below](#)
- ["Querying the RTSM" below](#)
- ["Querying the Class Model" on the next page](#)
- ["Analyzing Change Impact " on the next page](#)

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\APIs\JavaSDK_Samples**

Chapter 10: RTSM Web Service API

This chapter includes:

Conventions	258
RTSM Web Service API Overview	259
RTSM Web Service API Reference	262
Call the Web Service	262
Query the RTSM	263
Update the RTSM	266
Query the RTSM Class Model	267
Query for Impact Analysis	269
RTSM General Parameters	269
RTSM Output Parameters	272
RTSM Query Methods	273
UCMDB Update Methods	285
RTSM Impact Analysis Methods	288
RTSM Web Service API Use Cases	290
Examples	291

Conventions

Note: The HP Universal CMDB Web Service API can be used with the HP Operations Manager iRun-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.
- 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:

```
\\<HP Operations Manager i root directory>\AppServer\
webapps\site.war\amdocs\eng\APIs\CMDB_Schema
\webframe.html
```

RTSM Web Service API Overview

Note to HP Software-as-a-Service customers: For details on how to use the RTSM Web Service API in an HP Software-as-a-Service environment, contact HP Software-as-a-Service Support.

Note: Use this chapter in conjunction with the UCMDB schema documentation, available in the online Documentation Library.

The RTSM Web Service 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 impact analysis
- 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 Modeling Guide.

Users of the RTSMWeb Service API should be familiar with:

- The SOAP specification
- An object-oriented programming language such as C++, C# or Java
- RTSM
- Data Flow Management

This section includes the following topics:

- ["Uses of the API" below](#)
- ["Permissions" below](#)

Uses of the API

The RTSM Web Services 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

To access the WSDL file for the web service, go to:

http://localhost:21212/axis2/services/UcmdbService?wsdl. You will need to provide server administrator user credentials to view the WSDL file.

The user must have the **Run Legacy API** general action permission in order to log in.

The following table displays the additional required permissions for each Web Service API command:

Web Service API command	Required Permissions
addCIsAndRelations deleteCIsAndRelations updateCIsAndRelations	General Action: Data Update
executeTopologyQueryByName(AdHoc) executeTopologyQueryByNameWithParameters(AdHoc) executeTopologyQueryWithParameters(AdHoc)	General Action: Run Query by Definition For each query: View permission

Web Service API command	Required Permissions
getTopologyQueryExistingResultByName getTopologyQueryResultCountByName releaseChunks pullTopologyMapChunks getCINeighbours getFilteredCIsByType getCIsById getCIsByType getRelationsById	General Action: View CIs For each query: View permission
getQueryNameOfView	General Action: View CIs For each view: View permission
getChangedCIs	General Action: View History, View CIs
calculateImpact getImpactPath getImpactRulesByGroupName getImpactRulesByNamePrefix	General Action: Run Impact Analysis
getAllClassesHierarchy getClassAncestors getCmdbClassDefinition	None

Note: When the root context has been changed in UCMDB, follow these steps to enable access to the Web Service API:

1. Open the **\UCMDB\UCMDBServer\deploy\axis2\WEB-INF\web.xml** configuration file and locate the following section:

```
<servlet-class>
org.apache.axis2.transport.http.AxisServlet
</servlet-class>
```

Add the following lines after that:

```
<init-param>
<param-name>axis2.find.context</param-name>
<param-value>>false</param-value>
</init-param>
```

2. Open the **\UCMDB\UCMDBServer\deploy\axis2\WEB-INF\conflaxis2.xml** configuration file and locate the following line:

```
<parameter name="enableSwA" locked="false">>false</parameter>
```

Add the following line after it:

```
<parameter name="contextRoot" locked="false">test1/setup1/axis2</parameter>
```

where **test1/setup1** is your root context.

(To remove the root context, remove the text added to the path.)

3. Restart the UCMDb server.

RTSM Web Service API Reference

For full documentation on the request and response structures, refer to the [RTSM Web Service API Reference](#). These files are located in the following folder:

```
\\<HP Operations Manager i root directory>\AppServer\
webapps\site.war\amdocs\eng\APIs\CMDB_Schema
\webframe.html
```

Call the Web Service

Use standard SOAP programming techniques in the HP Universal CMDB Web Service API 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, the 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`

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`

For examples of connecting to the RTSM, see ["RTSM Web Service API Use Cases" on page 290](#).

Query the RTSM

The RTSM is queried using the APIs described in "[RTSM Query Methods](#)" on page 273. The queries and the returned RTSM elements always contain real BSM IDs. For examples of using query methods, see [Query Example](#).

This section includes the following topics:

- "[Just In Time Response Calculation](#)" below
- "[Processing Large Responses](#)" below
- "[Specifying Properties to Return](#)" on the next page
- "[Concrete Properties](#)" on the next page
- "[Derived Properties](#)" on page 265
- "[Naming Properties](#)" on page 265
- "[Other Property Specification Elements](#)" on page 265

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 is a previously calculated result. Therefore, the results of running a query returned to the client application may be different from 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 283](#) and ["releaseChunks" on page 284](#).

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 usually results in returning more properties than needed, and incurs a price in performance.

There are two types of predefined properties: qualifier properties and simple properties.

- **Qualifier properties.** Use this 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.

Update the RTSM

You update the RTSM with the update APIs. For details of the API methods, see "[UCMDB Update Methods](#)" on page 285.

This task includes the following steps:

- "[UCMDB Update Parameters](#)" below
- "[Use of ID Types with Update Methods](#)" below

UCMDB Update Parameters

This topic describes the parameters used only by the service's update methods.

- **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 CITs, 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 `referencedCIs` in ["CIsAndRelationsUpdates"](#) on the [previous page](#).

- **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 285 and ["RTSM Query Methods"](#) on page 273.

- **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 [Troubleshooting](#) 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.

Query the RTSM 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 Modeling Guide.

This section provides information on the following methods that return information about CITs and relations:

- ["getClassAncestors"](#) below
- ["getAllClassesHierarchy"](#) on the next page
- ["getCmdbClassDefinition"](#) on the next page

getClassAncestors

The `getClassAncestors` method retrieves the path between the given CIT and its root, including the root.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
className	The type name. For details, see "Type Name" on page 271 .

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 270 .

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 the next page .
className	The type name. For details, see "RTSM General Parameters " below .

Output

Parameter	Comment
cmdbClass	The class definition, consisting of name, classType, displayLabel, description, parentName, qualifiers, and attributes.
comments	For internal use only.

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](#).

RTSM General Parameters

This section describes the most common parameters of the service's methods.

This section includes the following topics:

- ["CmdbContext" on the next page](#)
- ["ID" on the next page](#)
- ["Key Attributes" on the next page](#)
- ["ID Types" on the next page](#)
- ["CIProperties" on the next page](#)
- ["Type Name" on page 271](#)
- ["Configuration Item \(CI\)" on page 271](#)
- ["Relation" on page 271](#)

CmdbContext

All RTSM 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 Modeling Guide. For information about identifying the key attributes from within the API client application, see ["getCmdbClassDefinition" on page 268](#).

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 if 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 the next page](#) and ["Relation" on the next page](#).

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. "[RTSM 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 Modeling Guide.

The props element is a `CIProperties` collection. For details, see "[RTSM General Parameters](#) " on [page 269](#).

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. "[RTSM 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 "[Query the RTSM Class Model](#) " on [page 267](#)"[Query the RTSM Class Model](#) " on [page 267](#).

For details, see [CI Type Manager](#) in the 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 270](#).

RTSM Output Parameters

This section describes the most common output parameters of the service methods. For more details, refer to the [online schema documentation](#).

This section includes the following topics:

- ["CIs" below](#)
- ["ShallowRelation" below](#)
- ["Topology" below](#)
- ["CINode" below](#)
- ["RelationNode" on the next page](#)
- ["TopologyMap" on the next page](#)
- ["ChunkInfo" on the next page](#)

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 `Relation` 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 the previous page](#)).
- `relationNodes`. This is one or more `RelationNode` (see ["RelationNode" above](#)).

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 263](#).

RTSM Query Methods

This section provides information on the following methods:

- ["executeTopologyQueryByNameWithParameters" on the next page](#)
- ["executeTopologyQueryWithParameters" on the next page](#)
- ["getChangedCIs" on page 275](#)
- ["getCINeighbours" on page 276](#)
- ["getCIsByID" on page 277](#)
- ["getCIsByType" on page 277](#)
- ["getFilteredCIsByType" on page 278](#)
- ["getQueryNameOfView" on page 281](#)

- ["getTopologyQueryExistingResultByName" on page 282](#)
- ["getTopologyQueryResultCountByName" on page 282](#)
- ["pullTopologyMapChunks" on page 283](#)
- ["releaseChunks" on page 284](#)

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 270 .
<code>queryName</code>	The name of the parameterized TQL in the RTSM for which to get the map.
<code>parameterizedNodeList</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 for items of a specific Configuration Item Type.

Output

Parameter	Comment
<code>topologyMap</code>	For details, see "TopologyMap" on the previous page .
<code>chunkInfo</code>	For details, see "ChunkInfo" on the previous page and "Processing Large Responses" on page 263 .

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 `parameterizedNodeList` 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 of 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.

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
<code>cmdbContext</code>	For details, see " CmdbContext " on page 270.
<code>queryXML</code>	An XML string representing a TQL without resource tags.
<code>parameterizedNodeList</code>	The conditions each node must meet to be included in the query results.

Output

Parameter	Comment
<code>topologyMap</code>	For details, see " TopologyMap " on page 273.
<code>chunkInfo</code>	For details, see " ChunkInfo " on page 273 and " Processing Large Responses " on page 263.

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 270 .
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
getChangedCIsResponseList	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 270 .
ID	The ID of the CI for which to retrieve the neighbors. This must be a real RTSM ID.
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 271 .
CIProperties	The data to be returned on each configuration item, called the Query Layout in the user interface. For details, see "TypedProperties" on page 266 .
relationProperties	The data to be returned on each relation (called the Query Layout in the user interface). For details, see "TypedProperties" on page 266 .

Output

Parameter	Comment
topology	For details, see "Topology" on page 272 .
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 270 .
CIsTypedProperties	A typed properties collection. For details, see "Other Property Specification Elements" on page 265 .
IDs	Only real RTSM IDs are valid in this collection.

Output

Parameter	Comment
CIs	A collection of CI elements.
chunkInfo	For details, see "ChunkInfo" on page 273 and "Processing Large Responses" on page 263 .

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 270 .
type	The class name. For details, see "Type Name" on page 271 .

Parameter	Comment
properties	The data to be returned on each configuration item. For details, see "CustomProperties" on page 265 .

Output

Parameter	Comment
CIs	A collection of CI elements.
chunkInfo	For details, see: "ChunkInfo" on page 273 and "Processing Large Responses" on page 263 .

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 24, 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
Greater	Numerical
GreaterEqual	Numerical
In	<p>String, numerical, and list</p> <p>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.</p>
InList	<p>List</p> <p>The condition's value and the property's value are lists.</p> <p>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.</p>
IsNull	<p>String, numerical, and list</p> <p>The item's property has no value. When operator <code>IsNull</code> is used, the value of the condition is ignored, and in some cases can be nil.</p>
Less	Numerical
LessEqual	Numerical
Like	<p>String</p> <p>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, <code>%Bi%</code> matches <code>Bismark</code> and <code>Bay of Biscay</code>, but not <code>biscuit</code>.</p>

Operator	Type of Condition/Comments
LikeIgnoreCase	String Use the LikeIgnoreCase operator as you use the Like operator. The match, however is not case-sensitive. Therefore, %Bi% matches biscuit.
NotEqual	String and numerical
UnchangedDuring	Date 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. 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. Note: The name UnchangedDuring is kept to preserve backward compatibility. In previous versions, the operator was used only with create and modify time properties.

Example of Setting Up a Condition:

```
FloatCondition fc = new FloatCondition();
FloatProp fp = new FloatProp();
fp.setName("attr_name");
fp.setValue(11f);
fc.setCondition(fp);
fc.setFloatOperator(FloatCondition.FloatOperator.EQUAL);
```

Example of Querying for Inherited Properties:

The target CI is `sample` which has two attributes, `name` and `size`. `sampleII` extends the CI with two attributes, `level` and `grade`. This example sets up a query for the properties of `sampleII` that were inherited from `sample` by specifying them by name.

```
GetFilteredCIsByType request = new GetFilteredCIsByType()
request.setCmdContext(cmdContext)
request.setType("sampleII");
CustomProperties customProperties = new CustomProperties();
PropertiesList propertiesList = new PropertiesList();
propertiesList.setPropertyNames(Arrays.asList("name", "size"));
customProperties.setPropertiesList(propertiesList);
request.setProperties(customProperties);
```


Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
type	The class name. For details, see "Type Name" on page 271 . The type can be any of the types defined using the CI Type Manager. For details, see CI Type Manager in the Modeling Guide.
properties	The data to be returned on each CI (called the Query Layout in the user interface). For details, see "CustomProperties" on page 265 .
conditions	A collection of name-value pairs and the operators that relate one to the other. For example, host_hostname like QA.
conditionsLogicalOperator	<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 273 and "Processing Large Responses" on page 263 .

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 270 .
viewName	The name of a view, which 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
<code>cmdbContext</code>	For details, see "CmdbContext" on page 270 .
<code>queryName</code>	The name of a TQL.
<code>queryTypedProperties</code>	A collection of sets of properties to retrieve for items of a specific Configuration Item Type.

Output

Parameter	Comment
<code>topologyMap</code>	For details, see "TopologyMap" on page 273 .
<code>chunkInfo</code>	For details, see "ChunkInfo" on page 273 and "Processing Large Responses" on page 263 .

getTopologyQueryResultCountByName

The `getTopologyQueryResultCountByName` method retrieves the number of instances of each node that matches the specified query.

Input

Parameter	Comment
<code>cmdbContext</code>	For details, see "CmdbContext" on page 270 .
<code>queryName</code>	The name of a TQL.
<code>countInvisible</code>	If true, the output includes CIs defined as invisible in the query.

Output

Parameter	Comment
getTopologyQueryResultCountByNameResponse	The number of instances matching the query.

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 273](#) and ["Query the RTSM" on page 263](#).

The client application must be able to handle the partial maps.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
ChunkRequest	The number of the chunk to retrieve and the <code>ChunkInfo</code> that is returned by the query method.
queryTypedProperties	A collection of sets of properties to retrieve for items of a specific CI type.

Output

Parameter	Comment
topologyMap	For details, see "TopologyMap" on page 273 .
comments	For internal use only.

Example of Handling Chunks:

```
GetCIsByType request =  
    new GetCIsByType(cmdbContext, typeName, customProperties);  
GetCIsByTypeResponse response =  
    ucldbService.getCIsByType(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++) {
            CIs cis =
                res.getTopologyMap().getCINodes().getCINode(m).getCIs();
            for(int i=0 ; i < cis.sizeC IList() ; i++) {
                // your code to process the CIs
            }
        }
    }

    GetCIsByType request =
        new GetCIsByType(cmdbContext, typeName, customProperties);
    GetCIsByTypeResponse response =
        ucmdbService.getCIsByType(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().getCINodes().size();
            m++) {
            CIs cis =
                res.getTopologyMap().getCINodes().getCINodes().get(m).getCIs();
            for(int i=0 ; i < cis.getCIs().size(); i++) {
                // your code to process the CIs
            }
        }
    }
}
```

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 270 .
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" below](#)
- ["addCustomer" on the next page](#)
- ["deleteCIsAndRelations" on the next page](#)
- ["removeCustomer" on page 287](#)
- ["updateCIsAndRelations" on page 287](#)

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
cmdbContext	For details, see "CmdbContext" on page 270 .
updateExisting	Set to true to update items that already exist in the RTSM. Set to false to throw an exception if any item already exists.
CIsAndRelationsUpdates	The items to update or create. For details, see "CIsAndRelationsUpdates" on page 266 .
ignoreValidation	If true, no check is performed before updating the RTSM.
dataStore	Changer information.

Output

Parameter	Comment
createdIDsMapList	The list of client IDs mapped to RTSM IDs. For details, see the description above.
comments	For internal use only.

addCustomer

The `addCustomer` method adds a customer.

Input

Parameter	Comment
CustomerID	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
cmdbContext	For details, see "CmdbContext" on page 270 .
CIsAndRelationsUpdates	The items to delete. For details, see "CIsAndRelationsUpdates" on page 266
dataStore	Changer information.

removeCustomer

The `removeCustomer` method deletes a customer record.

Input

Parameter	Comment
CustomerID	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 270 .
CIsAndRelationsUpdates	The items to update. For details, see "CIsAndRelationsUpdates" on page 266 .
ignoreValidation	If true, no check is performed before updating the RTSM.
dataStore	Changer information.

Output

Parameter	Comment
createdIDsMapList	The list of client IDs mapped to RTSM IDs. For details, see "addCIsAndRelations" on page 285 .

RTSM Impact Analysis Methods

This section provides information on the following methods:

- ["calculateImpact" below](#)
- ["getImpactPath" on the next page](#)
- ["getImpactRulesByNamePrefix" on the next page](#)

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" on the next page](#).

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
impactCategory	The type of event that would trigger the rule being simulated.
IDs	A collection of ID elements.
impactRulesNames	A collection of <code>ImpactRuleName</code> elements.
severity	The severity of the triggering event.

Output

Parameter	Comment
impactTopology	For details, see "Topology" on page 272 .
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" on the previous page](#) is used as the identifier input argument of `getImpactPath`.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
identifier	The key to the server response that was returned by <code>calculateImpact</code> .
relation	A Relation based on one of the "ShallowRelation"s returned by <code>calculateImpact</code> in the <code>impactTopology</code> element.

Output

Parameter	Comment
impactPathTopology	A CIs collection and an <code>ImpactRelations</code> collection.
comments	For internal use only.

An `ImpactRelations` element consists of an ID, type, end1ID, end2ID, rule, and 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 270 .
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.

RTSM Web Service API Use Cases

The following use cases assume two systems:

- HP Operations Manager i server
- A third-party system that contains a repository of configuration items

This section includes the following topics:

- ["Populating the RTSM" below](#)
- ["Querying the RTSM" below](#)
- ["Querying the Class Model" on the next page](#)
- ["Analyzing Change Impact" on the next page](#)

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

See the following code samples:

- [The Example Base Class](#)
- [Query Example](#)
- [Update Example](#)
- [Class Model Example](#)
- [Impact Analysis Example](#)

These files are located in the following directory:

**\\<BSM Gateway Server root directory>\AppServer\
webapps\site.war\amdocs\eng\APIs\WebServiceAPI_Samples**

Chapter 11: Internal-Global ID Conversion API

This chapter includes:

Using the Internal-Global ID Conversion API	292
How to Convert RTSM Internal IDs to Global IDs	292
How to Convert Global IDs to RTSM Internal IDs	294

Using the Internal-Global ID Conversion API

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**

How to Convert RTSM Internal IDs to Global IDs

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/convertRtسمToGlobalIds?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/  
convertRtسمToGlobalIds?ciIds=35845d64ba6250d4d45d0e75e3e59001,  
316e438afc58135f9e87113bd4427d37,175fe30654e9a9dd16ddab66b2237c92,  
cb8a27f8ce900aa380c41ac1c74d8153,8066fe1fcb712926fe3793a75b0e733c,  
bbbc37f7ce900aa380c41ac1c74d8153
```

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 cilds 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

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?ciId=s=35845d64ba6250d4d45d0e75e3e59001,316e438afc58135f9e87113bd4427d37,175fe30654e9a9dd16ddab66b2237c92,cb8a27f8ce900aa380c41ac1c74d8153
```

Result map

```
<map>
  <entry>
    <key>316e438afc58135f9e87113bd4427d37</key>
    <value>316e438afc58135f9e87113bd4427d37</value>
  </entry>
  <entry>
    <key>175fe30654e9a9dd16ddab66b2237c92</key>
    <value>175fe30654e9a9dd16ddab66b2237c92</value>
  </entry>
</map>
```

```
</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 cilds parameter. Number of passed CI IDs reaches the limit, which is 10000.
INTERNAL_SERVER_ERROR	500	General failure to run query against RTSM

Chapter 12: Data Flow Management Java API

This chapter includes:

Using the Data Flow Management Java API	296
---	-----

Using the Data Flow Management Java API

Note: Use this chapter in conjunction with the DFM API Javadoc, available in the online Documentation Library.

This chapter explains how third-party or custom tools can use the HP Data Flow Management Java API to manage Data Flow. 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.

The following services are available in the Discovery Services package:

- **DDMConfigurationService.** Services to configure the Data Flow Probes, clusters, IP Ranges, and Credentials. The Universal Discovery server can be configured with an XML file or via the Data Flow Probe.
- **DDMManagementService.** Services to analyze and view the progress, results, and errors of the Universal Discovery run.
- **DDMSoftwareSignatureService.** Services to define software items to be discovered by the Data Flow Probe components. The definitions are system-wide. If more than one Data Flow probe component is defined, the definitions apply to all of them.
- **DDMZoneService.** Services to manage zone-based discovery.

In addition to these services, there are Data Flow Management client APIs, which are used in creating Jython adapters. For details, see ["Developing Jython Adapters" on page 37](#).

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.

In addition, the user must have the **Access to SDK** general action permission in order to log in.

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 ["Create an Integration User" on page 253](#).

Chapter 13: Data Flow Management Web Service API

This chapter includes:

Data Flow Management Web Service API Overview	298
Conventions	299
Calling the Web Service	299
Data Flow Management Methods and Data Structures	299
Code Sample	310
Adding Credentials Example	313

Data Flow Management Web Service API Overview

This chapter explains how third-party or custom tools can use the HP Data Flow Management Web Service API to manage Data Flow.

The HP Data Flow Management Web Service API is used to integrate applications with HP Business Availability Center. 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 Data Flow Management Web Service should be familiar with:

- The SOAP specification
- An object-oriented programming language such as C++, C# or Java
- HP Business Availability Center
- Data Flow Management

The user must have the **Run Legacy API** general action permission in order to log in.

For full documentation on the available operations, see *HP Universal Discovery Schema Reference*. These files are located in the following folder:

```
\\<HP Operations Manager i root  
directory>\AppServer\webapps\site.war\amdocs\eng\APIs\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.
- Data Flow Management 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.

Calling the Web Service

The HP Data Flow Management Web Service API 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>:21212 /axis2/services/DiscoveryService`
- Default password: `"admin"`
- Default username: `"admin"`

SOAP programmers can access the WSDL at:

- `axis2/services/DiscoveryService?wsdl`

Data Flow Management Methods and Data Structures

This section lists the Data Flow Management Web Service API methods and data structures, and provides a brief summary of their uses. For full documentation of the request and response for each operation, see *HP Universal Discovery Schema Reference*.

This section includes the following topics:

- ["Data Structures" below](#)
- ["Managing Discovery Job Methods" on the next page](#)
- ["Managing Trigger Methods" on page 302](#)
- ["Domain and Probe Data Methods" on page 304](#)
- ["Credentials Data Methods" on page 307](#)
- ["Data Refresh Methods" on page 309](#)

Data Structures

These are some of the data structures used in the Data Flow Management 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 Universal Discovery Schema Reference*.

IPList

A list of IP elements, each of which contains an IPv4 or IPv6 Address.

See `IPList` in the *HP Universal Discovery Schema Reference*.

IPRange

An `IPRange` has two elements, `Start` and `End`. Each element contains an `Address` element, which is an IPv4 or IPv6 Address.

See `IPRange` in the *HP Universal Discovery 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 Universal Discovery Schema Reference*

Managing Discovery Job Methods

activateJob

Activates the specified job.

See " [Code Sample](#)" on page 310.

Input

Parameter	Comment
cmdbContext	For details, see " CmdbContext " on page 270.
JobName	The name of the job.

deactivateJob

Deactivates the specified job.

Input

Parameter	Comment
cmdbContext	For details, see " CmdbContext " on page 270.
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 270.
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 270 .

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 270 .
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 270 .
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 270.
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 270.
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 270.
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 270.
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 270.
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 270.
domainName	The name of the domain.

Output

Parameter	Comment
domainType	The domain type.

getDomainsNames

Returns the names of the current domains.

See "[Code Sample](#)" on page 310.

Input

Parameter	Comment
cmdbContext	For details, see " CmdbContext " on page 270.

Output

Parameter	Comment
domainNames	The list of domain names.

getProbelPs

Returns the IP addresses of the specified probe.

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
domainName	The domain to check.
probeName	The name of the probe used on that domain.

Output

Parameter	Comment
probelPs	The "IPList" of the addresses in the probe.

getProbesNames

Returns the names of the probes in the specified domain.

See [" Code Sample" on page 310](#).

Input

Parameter	Comment
cmdbContext	For details, see "CmdbContext" on page 270 .
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 270.
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 "[Code Sample](#)" on page 310.

Input

Parameter	Comment
cmdbContext	For details, see " CmdbContext " on page 270.
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 270.
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 "[Code Sample](#)" on page 310.

Input

Parameter	Comment
cmdbContext	For details, see " CmdbContext " on page 270.
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 270.
domainName	The domain for which to get the credentials.
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 270 .
domainName	The domain for which to get the credentials.
protocolName	The name of 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 270 .
domainName	The domain.
protocolName	The name of 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 270 .
domainName	The domain in which to update credentials.
protocolName	The name of a protocol used on the domain.

Parameter	Comment
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 270.
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 between 0-1. When the response is 1, the **rediscoverCIs** call has completed.

Input

Parameter	Comment
cmdbContext	For details, see " CmdbContext " on page 270.
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 270.
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 270.
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 propeties 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;
}
}
```

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 = 21212;
    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 propeties 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;  
}  
} // End class
```

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on RTSM Developer Reference Guide (Operations Manager i 10.00)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to ovdoc-asm@hp.com.

We appreciate your feedback!



[Go OMi!](#)