




Hewlett Packard
Enterprise

Application Performance Management

Version 9.40, Released August 2017

RTSM Best Practices

Published August 2017



Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise 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. HPE 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 HPE 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 2005 - 2017 Hewlett Packard Enterprise Development LP

Trademark Notices

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

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

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

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

iPod is a trademark of Apple Computer, Inc.

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

Microsoft®, Windows®, Windows NT®, Windows® XP, and Windows Vista® are U.S. registered trademarks of Microsoft Corporation.

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

UNIX® is a registered trademark of The Open Group.

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.hpe.com/group/softwaresupport/search-result?keyword=>

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

Support

Visit the HPE Software Support website at: <https://softwaresupport.hpe.com>

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE 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 website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE 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 HPE Passport user and sign in. Many also require a support contract. To register for an HPE Passport ID, go to <https://softwaresupport.hpe.com> and click **Register**.

To find more information about access levels, go to:

<https://softwaresupport.hpe.com/web/softwaresupport/access-levels>

HPE Software Integrations and Solutions

Visit the Integrations and Solutions Catalog at <https://softwaresupport.hpe.com/km/KM01702731> to explore how the products in the HPE Software catalog work together, exchange information, and solve business needs.

Contents

- Chapter 1: Introduction 5
- Chapter 2: The Role of the RTSM 6
 - RTSM Capabilities 6
 - Enriching RTSM Data 7
 - Manual, Semi-automated, and Automated Enriching Methods 8
- Chapter 3: RTSM Data Feeding Methods 9
- Chapter 4: RTSM Integration Strategy 10
 - Global ID 10
 - Determining which RTSM/UCMDB is Defined as CMS 11
 - Preventing Multiple CMSs 11
 - Initial Synchronization 12
- Chapter 5: APM Hierarchical Deployment 13
 - Global ID in Hierarchical Deployments 13
 - Topology Discovery in a Manager of Managers Hierarchy 14
- Chapter 6: APM-UCMDB Topology Synchronization 15
 - Population, Data Push, and Federation 15
 - When to Use UCMDB-APM Synchronization 16
 - Integration Prerequisites 17
 - Data Flow Probe Installation and Deployment 17
 - Handling Class Model Differences 19
 - Effective Methods for Aligning the Class Model during Synchronization 19
 - Deletion Policy 19
 - Aging Settings 20
 - Importing CI Instances from Previous Versions of RTSM 21
- Chapter 7: Reconciliation 22
 - Auto Complete Reconciliation 22
 - Reconciliation Priorities 22
- Chapter 8: Checking the Logs 26
 - Log Configuration 26
 - Population Synchronization Troubleshooting 26
- Appendix A: Acronym Table 27
- Appendix B: CITs 28
- Send Documentation Feedback 32

Chapter 1: Introduction

The goal of this document is to provide best practices for integrating the HPE APM Run-time Service Model (RTSM) with other products, or with other APM deployments. It describes synchronizing the RTSM with an external Configuration Management Database (CMDB), or with another RTSM.

This document does not explain the basic concepts, or provide detailed information for each data synchronization or other integration flows. We recommend that you use the Data Flow Management Guide, available as part of the HPE APM Help, for more detailed information.

Chapter 2: The Role of the RTSM

For an IT enterprise to be properly aligned with business objectives, it must be focused on ensuring the health and availability of specific IT services. To maintain service health effectively in today's dynamic environments, it is critical that enterprise management technologies be equally dynamic.

APM delivers the ability to capture and leverage data collected from a variety of data sources, and apply it to the challenges of maintaining service health, event correlation, and event reduction. The RTSM is a key component in this process.

The RTSM is an instance of HPE's CMDB product, the HPE Universal CMDB (UCMDB), embedded in APM. This UCMDB instance performs the expected functions of a CMDB, reconciling and storing configuration items (CIs) that represent the IT environment components, to help organizations understand the relationships between these components, and track their configuration.

The RTSM is a UCMDB instance that is managed by APM and is an integral part of the APM Platform. The RTSM leverages the same core technology as the UCMDB, but also serves a different purpose, as the RTSM serves as an operational repository within APM. The various functions performed by the RTSM mean that its role differs significantly from the traditional role fulfilled by a CMDB and its associated use cases. The RTSM reconciles and stores CIs discovered by the different APM data sources. In addition, the RTSM models operational data like downtime, alerts, SLAs, and data related to the configuration of Application Performance Management (APM) capabilities in APM.

RTSM also serves a specific set of use cases supporting the requirements of consolidated event management, provided by APM's Service and Operations Bridge capabilities.

The RTSM is used in the Service and Operations Bridge context to:

- Drive event processing on top of a dynamic, near-real-time model of the IT infrastructure.
- Correlate various monitoring sources against managed services to provide a comprehensive view of the health of the service.
- Serve as a central repository for dynamic changes discovered during the course of real-time operations to facilitate the cases above.

The RTSM is installed as part of APM; APM cannot use any other CMDB instance as its internal CI repository. However, the topologies held in the RTSM can be synchronized with other CMDB instances using the UCMDB multi-instance topology synchronization solution. For details, see ["APM-UCMDB Topology Synchronization" on page 15](#).

RTSM Capabilities

RTSM capabilities include:

- **Service model:** The service model encompasses the entire physical and virtual service infrastructure, including transactions, applications, servers, storage, and networks, fully integrated to present a complete view of each IT service and business service for your entire monitored environment.
- **Service model capturing:** The service model is automatically discovered and created using native APM operational solutions, such as RUM, Diagnostics, SiteScope, and SPIs.

In BPM, the service model is manually created by the user.

- **Manual and automated service model management:** RTSM enables manual CI creation through a designated UI under the IT Universe, as well as automated methods using pattern-based models. This enables creating high level and business CIs according to specific patterns. On top of that, it utilizes special adapters for importing the topology from external Excel CSV and XML files. For details, see ["Manual, Semi-automated, and Automated Enriching Methods" on the next page.](#)
- **Up-to-date maintenance:** The service model is kept up-to-date using APM operational solutions.

Enriching RTSM Data

The various monitoring products included in APM implement discovery services, but this is limited in range and scope by the function that they serve. The combined monitoring information can be used to assemble a topological representation of the managed resources within the RTSM, but this may not give a complete picture of the IT environment.

RTSM supports integration with other CMDBs as part of a configuration management system (CMS) to synchronize CI topology data. The CMS helps create a coherent, logical model of the IT organization's infrastructure. The RTSM and the UCMDDB complement each other and should be considered, along with Discovery technologies, to be elements of a CMS.

Within the CMS, a UCMDDB instance acts as the central CMDB (this is the CMDB that generates the global IDs – for more information, see ["Global ID" on page 10](#)). Other UCMDDB instances (including the RTSM) are referred to as citizen CMDBs. CMDB implementations often involve federation (the inclusion of data from other sources into the CMDB).

The use of an RTSM operating as an internal CMDB within APM, separate from the rest of the CMS, allows for independent management of downtime, and the upgrade of APM deployments, without affecting the UCMDDB and the rest of the HPE products that are integrated into the CMS. The RTSM class model is an extension of the UCMDDB class model, and allows for the management of the APM private class model extensions that are required for APM operational use cases.

Other use cases, which revolve around traditional CMDB services, such as configuration and change management, may require Universal Discovery (UD) services. When an end-to-end solution is deployed, UD services are typically used to populate a central UCMDDB, and the RTSM can leverage richer data from the UCMDDB where applicable.

The RTSM is a highly optimized service model, created for the high performance needs of the APM use cases, and as such, includes the CIs that are required for complete and accurate monitoring. Other CIs which are not crucial for the APM use cases should be stored in the central UCMDDB instance. The RTSM can also service as an input source of environmental changes to the parent CMS system.

The following are the capacity limitations of RTSM:

- 600K calculated for monitored CIs
- 5M CIs and relationships in general

To avoid over populating RTSM, we recommend that you run Discovery on the CMS side and synchronize only CIs which are relevant to the APM use case. The reasons for this are:

- CMS has a larger capacity than RTSM
- CMS can be upgraded to a new version and contains the latest content packs. When synchronizing with CMS10, it has the advantage of using UD. For information, see APM System Requirements and Support Matrixes.
- CMS is adjusted to better cope with Discovery performance impact

Manual, Semi-automated, and Automated Enriching Methods

RTSM enables manual CI creation through a designated UI under the IT Universe. It enables creating, updating and creating relationships between new and existing CI instances.

Another effective modeling tool in the RTSM modeling studio is the instance-based models, that provide a quick and efficient method to drag and drop CIs and models and create relationships to other CIs on the instance level.

In the RTSM modeling studio, RTSM introduces a semi-automated method using pattern-based models to enable creating high level and business CIs according to a specific pattern. On top of that, it utilizes special adapters for importing the topology from external Excel CSV and XML files.

For fully automated modeling, we recommend using the enrichment manager to create automated scheduled enrichment to update, remove and create relationships between CI instances according to specific criteria and rules based on a TQL.

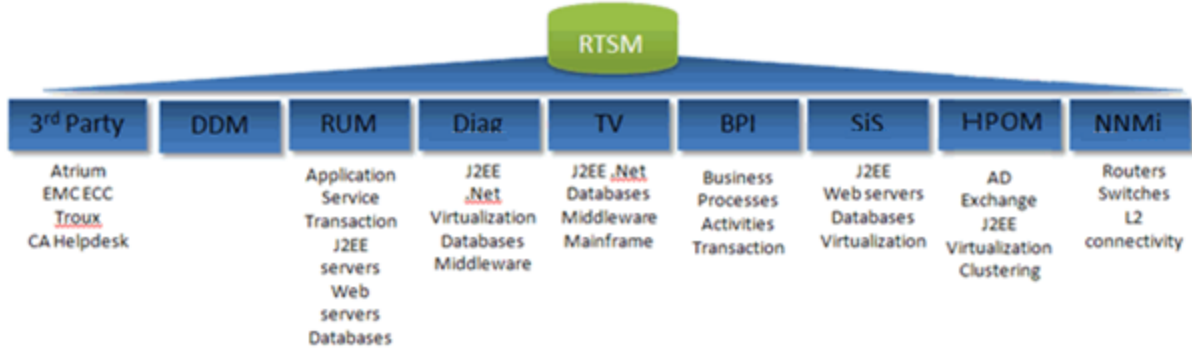
For details on pattern/instance-based modeling and topology enrichment methods, see the Modeling Guide and the Effective Modeling for APM Best Practices.

Chapter 3: RTSM Data Feeding Methods

Note: RTSM is designed to support APM operational use cases only. We recommend that only the topology that is relevant to these operational use cases be populated to the RTSM.

There are several ways of providing the data in an RTSM:

- APM data sources report topology changes, keeping the RTSM up-to-date.



- For change and configuration management and other use cases, we recommend that you integrate the topology via a dedicated UCMDB instance as part of a CMS, instead of integrating directly on top of an RTSM. For more details, see ["RTSM Integration Strategy" on page 10](#).
- RTSM supports Population, Data Push, and Federation integration types. Any of these three integration types can be used to provide relevant topological data in an RTSM. The data source can be either another CMDB instance (dedicated UCMDB or another RTSM), or external data sources (for example, BMC Atrium). For more details, see ["Population, Data Push, and Federation" on page 15](#).

All topology populated to the RTSM goes through the reconciliation engine, providing one consistent service model, regardless of the topology integration method.

Chapter 4: RTSM Integration Strategy

For topology-based integrations between APM and other HPE centers (for example, HPE Service Manager), we recommend that you install a central UCMDB instance as part of a CMS, rather than integrate the other HPE center directly with the RTSM. We also recommend that you run automated discovery directly under the UCMDB instance, and then synchronize the topology relevant for APM operational use cases with APM, as necessary.

APM provides a set of predefined data synchronization Topology Query Languages (TQLs) for pulling the topology from the central UCMDB to APM. These TQLs are used for common synchronization. However, we recommend that you create a specific TQL that is customized for the specific data required to be synchronized in order to avoid overpopulating RTSM.


Direct integration with configuration management products (such as HPE Release Control) is not supported by the RTSM and requires the installation of a CMS or dedicated UCMDB instance.

Global ID

The global ID is a unique CI ID, generated by the central UCMDB, which identifies that CI across the entire portfolio, making it easier to work in multiple UCMDB environments. Once a global ID is available, it can be used as a shared context for point-to-point integrations between HPE centers (for example, a new incident is opened in HPE Service Manager when an event is received for a CI).

The UCMDB is configured by default to act as the global ID generator. In certain cases, you may need to change this so that the RTSM acts as the global ID generator; a use case for this may be an APM hierarchical deployment with no CMS, where the top instance of the RTSM is required as the global ID generator.

The following diagram depicts a topology in which UCMDB is the Global ID generator.



For details on configuring the RTSM to act as the global ID generator, refer to the RTSM Data Flow Management Guide.

Determining which RTSM/UCMDB is Defined as CMS

In a specific deployment where several RTSMs and UCMDBs instances are integrated, only one UCMDB instance should be the CMS.

To determine if a specific instance of RTSM/UCMDB is the CMS:

1. Log in to the JMX Console (<http://<UCMDB server>:8080/jmx-console/>).
2. Go to **Multiple UCMDB Instances Services**.
3. Click **getIsGlobalIdGenerator**. If the call returns **true**, the instance is a CMS.

Preventing Multiple CMSs

Since the out-of-the-box UCMDB is set as a CMS, other UCMDB instances should be configured as non-CMS to ensure that there is one and only instance of CMS.

To configure UCMDBs as non-CMS:

1. Log in to the JMX console (<http://<UCMDB server>:8080/jmx-console/>).
2. Go to **Multiple UCMDB Instances Services**.
3. Click **setAsNonGlobalIdGenerator**.

Initial Synchronization

Initial synchronization synchronizes a newly installed RTSM with an existing CMS/RTSM. The newly installed RTSM contains no data. With this type of synchronization, all instances of CIs in the target UCMDB are replicated into the new RTSM, while retaining their original CMDB IDs.

Initial synchronization performs a full replication of an existing CMS/RTSM.

Following the initial synchronization, you can use the Population or Federation method to periodically update specific sets of CI types.

Another use case for initial synchronization is to maintain an up-to-date replication of an active CMS/RTSM, when the replicated side is not receiving any CIs from other feeders and no CIs are manually created.

Initial synchronization can also be performed during upgrades in situations where an RTSM/CMS instance should be replicated.

Note: Initial synchronization is supported in APM.

Initial synchronization is performed via RTSM's JMX console. For details, see Initial Synchronization in the Data Flow Management Guide.

Chapter 5: APM Hierarchical Deployment

APM supports a hierarchical deployment of APM instances that enables the forwarding of events and topology from one APM instance to another. The main motivations for building a hierarchical deployment are:

- **Event related scaling.** When a hierarchy of APM deployments is defined to deal with a very large number of events. The upper instances of the deployment get only the “important” summary events.
- **Topology related scaling.** When a given APM server is planned to reach APM’s capacity limitation in terms of topology and/or calculated CIs, several instances can be shared to populate the topology while the upper instance (Manager of Managers - MoM) will contain the high level and Business CIs. This, together with the event distribution above, allows topology and event-based scaling.
- **Geographical distribution.** When there are several data centers in different geographical locations, each managing their own APM instance. In this use case, the data from different geographical locations can be consolidated in one central instance, the MoM.
- **Organizational structure / Different consumers.** When the structure of the organization includes multiple business units or departments, each with its own APM instance. The consolidated picture is achieved either by two-way synchronization between APM instances, or by defining one central MoM instance.
- **Functional structure.** When the IT Department chooses to manage applications and infrastructure separately, by creating separate APM instances for application owners and for infrastructure owners. In this deployment scenario, there can be several APM instances, each one operated by domain experts.
- **Organizational mergers and acquisitions.** When there are several APM instances as a result of mergers and acquisitions. Consolidation can be achieved by synchronizing the data to one central instance.

We recommend using:

- Population synchronization to synchronize between APM instances under MoM deployment.
- Push synchronization to synchronize between CMS into APM instances, according to the support matrix.

Global ID in Hierarchical Deployments

One of the main motivations for synchronizing the topology in hierarchical deployments is to provide a topological context for the forwarded events. To resolve the incoming event, and find the associated CI, the information available in the event is used. Part of this information is the global CI ID that provides the shared topological context between two APM instances.

The global CI ID can be obtained in one of the following ways: **pushback of IDs** or **two-way topology synchronization**.

- **Deployment with no central UCMDB.** When the deployment does not include a central UCMDB, we recommend that you configure the top APM instance as a global CI ID generator and use pushback of IDs towards each of the child APM instances. This approach is simpler than configuring a two-way synchronization because:

- It requires the installation of only one Data Flow Probe (as opposed to two Data Probes for two-way synchronization).
- It is easier to control the scope of the synchronization since the top instance pushes back the IDs only for the pulled CIs.
- **Deployment includes central UCMDB.** When the deployment includes a central UCMDB, we recommend that you configure a two-way synchronization between the central UCMDB and the top APM instance, and between the top APM instance and its child APMs.
- **Central UCMDB added to existing deployment.** When the initial deployment does not include a central UCMDB, and then one is added, the central UCMDB must generate new global IDs for all the CIs. The new global IDs are broadcast from the central UCMDB to the top APM instance, and then from the top APM instance to its child APMs. This entails a full synchronization between the central UCMDB and the top APM instance, and again between the top APM instance and its child APMs.

Note: The central UCMDB is by default the global CI ID generator. The UCMDB cannot generate global IDs for APM private CI types (such as the End User Group CIs) when they do not exist in the central UCMDB. This is relevant for the use case above, when the central UCMDB is now acting as the global ID generator in place of the top APM instance.

Topology Discovery in a Manager of Managers Hierarchy

A Manager of Managers hierarchy consists of a top APM instance, and its child APMs. The following steps describe the sequence in which topology is discovered by the top APM instance, and is then reported outwards to the child APM instances:

1. The global APM Probe (MoM) pulls the topology from the local (child) APMs, or alternatively, the local APM pushes the topology into the global APM. We recommend using the Push method.
2. The central CMDB pulls the topology from the global APM, and generates global IDs, or alternatively, the global APM pushes topology into the CMS. We recommend using the Push method.
3. The global APM pulls the topology from the central CMDB, or alternatively, the CMS pushes the topology into the global APM. Push is the recommended method when using the latest versions.
4. The local APMs pull topology from the global APM, or alternatively, the global APM pushes the topology to the local APM.
5. Due to the global ID existence per CI, CI resolution from events can be precisely based on the ID instead of using hints based approximation. Events can now be forwarded, and correct CIs are assigned.

Note: Each integration path (pulling or pushing topology from APM to CMS) is configured independently of the other paths with ongoing data synchronization; there is no need to configure any special task scheduling for the described sequence to occur. This sequence simply illustrates how data is transferred from one UCMDB/RTSM to another, and which information is added at each step.

Chapter 6: APM-UCMDB Topology Synchronization

The topology stored in the RTSM can be synchronized with the central UCMDB (or another UCMDB instance) using the UCMDB multi-instance synchronization technology.

For details on topology synchronization between different versions UCMDB of APM, see the APM-CMS Synchronization Integration Matrix in the APM System Requirements and Support Matrixes Guide.

Population, Data Push, and Federation

There are three types of supported integrations: Population, Data Push, and Federation. These are core integration capabilities in RTSM and UCMDB.

- **Population.** Performs a scheduled retrieval of data from the external repository and creates a persistent copy of the data in the local CMDB.

Note: By default, if data required for the reconciliation of a particular CI cannot be retrieved (for example, if the data is missing in the source), that CI is ignored without causing the entire job to fail. You can change this behavior in the CmdbAdapter configuration. For details, see Adapter Configuration Tab in Adapter Configuration Tab in the Data Flow Management Guide.

Keep in mind that warnings are sent to the log file when a CI is ignored. To prevent these warnings from appearing in the log, change the log warning level to Error only (see Log Severity Levels in the RTSM Administration Guide), or add the CIs to the integration TQL.

Note: The capacity limitation of population synchronization is a maximum of 500K CIs for all sync jobs which are pulling the topology from a specific server.

- **Federation.** Performs an on-demand retrieval of data from the external repository and creates a transient image of the data in the local CMDB.
- **Data Push.** Copies data from the local CMDB to an external repository, where the latter typically records a persistent copy of the data.

Data Push integration is the preferred method for topology synchronization from CMS to APM.

Population integration should generally be used for topology integrations between a central UCMDB and the RTSM, or between RTSM instances (see ["APM Hierarchical Deployment" on page 13](#)).

A Federation integration type is used when APM needs to retrieve information from other HPE centers that are integrated with the central UCMDB, without replicating the data; for example, incidents from HPE Service Manager, or planned changes from HPE Release Control. For these use cases, there is no reason to replicate the data as APM does not control the data lifecycle. This flow is supported only under specific E2E solutions, such as Closed Loop Incident Process (CLIP). CLIP is an integrated solution that brings together HPE offerings for APM, SM, UCMDB, and OO. For details, see [Closed Loop Incident Process 9.30](#) (http://helpfiles.intactcloud.com/EC/1.10/CLIP_9.30_OMi_Solution_Configuration_Guides/Content/CLIP_

Solution.htm). For any other use case, we recommend using Data Push or Population for topology synchronization.

Both the Federation and Population integration types can be configured using the same integration point. APM provides an out-of-the-box integration point that includes all required configuration for out-of-the-box APM flows. We recommend starting with this integration point, and modify the configuration based on specific customer requirements. For more details, see “Set Up Integrations between CMS and APM” in the RTSM Data Flow Management Guide.

When to Use UCMDB-APM Synchronization

The following sections describe use cases for topology synchronization.

Synchronizing the Topology from UCMDB to APM

One of the main use cases for synchronizing topology from the UCMDB to APM is to obtain global CI IDs that can be used later for integrations between APM and other HPE centers.

The following are additional use cases where UCMDB to APM synchronization may be required:

- APM flows that are integrated with topology reported by UD:
 - **Monitor Deployment Wizard (MDW) in SiteScope:** The MDW uses the topology discovered by UD when configuring SiteScope monitors. The MDW retrieves relevant information from CIs reported by UD, and use it to help configure the monitors and assign the monitors to the relevant CIs.
 - **Consume UD credentials:** Once UD connects to a discovered source using credentials, the reference to the credentials is stored on the reported CIs. The credentials are then used in several APM flows: MDW, integration with Operations Orchestration, monitor assignment for vertical solutions.
- Enrich the topology in APM:

Some APM data collectors report “shallow” topology. Synchronizing topology discovered by UD, or other UCMDB topology reporting sources, enriches the CIs reported by APM. For example, APM data collectors can report the nodes (hosts) using IP. Once those same nodes are discovered by UD, additional information is available, such as DNS name, OS, OS version, and so forth. This can help APM users manage their CIs in the RTSM.
- There are different practices for the modeling of logical CIs (such as Business Services, Business Applications, and so forth). Any CI that is populated to the UCMDB can be synchronized with RTSM, thus reducing the overhead of manually redefining logical CIs.

Synchronizing the Topology from APM to UCMDB

The following are use cases where APM to UCMDB synchronization may be required:

- APM data collectors (such as RUM, Diagnostics) have the ability to discover the relationships between Business Applications and Transactions, and their underlying infrastructure. These relationships can be synchronized to complete the Business Service definition in the UCMDB, avoiding the need for manual modeling.
- APM data collectors can serve as additional sources populating the UCMDB. The topology is then utilized when building a CLIP (Closed Loop Incident Process) solution, or as a complementary topology source for UD; APM discovers a more “shallow” topology that is used as a trigger for UD deep discovery.

- RTSM contains up-to-date topology that is updated at runtime and can provide a more current picture than other UCMDB data sources (such as UD).

Setting Up UCMDB-APM Topology Synchronization

To set up a two-way topology synchronization, refer to the following sections in RTSM Data Flow Management Guide:

- For instructions of pushing data from CMS to RTSM, see "Work with Push Jobs" in the Data Flow Management Guide.
- For instructions on populating data from the CMS (UCMDB) to RTSM using a pre-configured integration point on the APM side, see "Set Up Integrations Between CMS and APM" in the Data Flow Management Guide.
- For instructions on populating the data from APM to the CMS (UCMDB), see "Work with Population Jobs" in the Data Flow Management Guide.

You can also find an overview of the topology synchronization mechanism in Integrating Multiple CMDBs in the Data Flow Management.

Integration Prerequisites

Pay attention to the following prerequisites before beginning your synchronization.

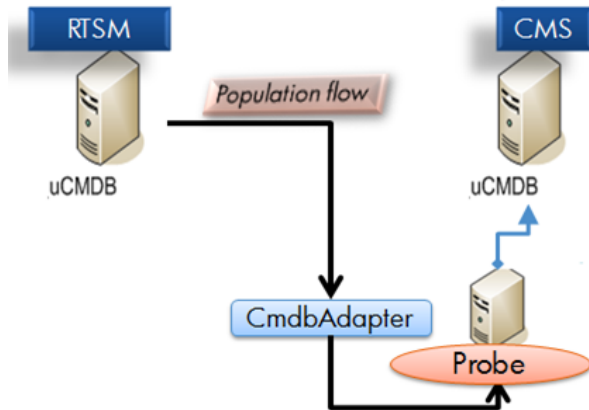
- Review the support matrix.
- We recommend that you deploy the latest supported content packs on both synchronization sides.
- Install a Data Flow Probe according to the instructions in "[Data Flow Probe Installation and Deployment](#)" [below](#)
- For data population, prepare an integration TQL or use one of the pre-defined TQLs. Synchronization TQL should be defined on the source (the system from which the CIs are pulled).
- Examine the class model differences according to the information in section "[Handling Class Model Differences](#)" on [page 19](#) and define the CITs and attributes to be synchronized accordingly.

Data Flow Probe Installation and Deployment

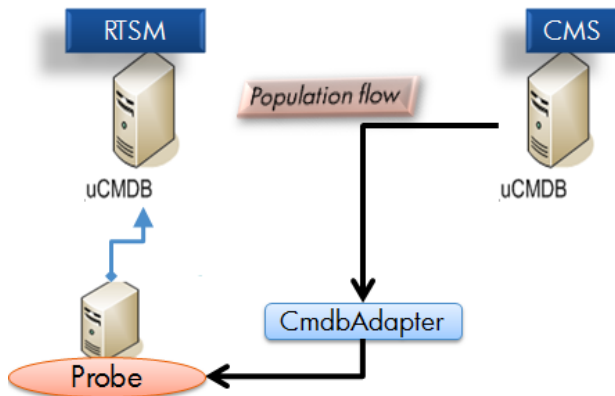
Topology synchronization requires installation of the Data Flow Probe. A single Data Flow Probe can be used to pull the topology from multiple UCMDB sources. For example, when using a hierarchical APM deployment, the same Data Flow Probe can be used to pull the topology from several child APM instances.

The Data Flow Probe can report topology to only one destination. This means that, when configuring a two-way topology synchronization, two separate Data Flow Probes must be used.

In the following diagram, CMS initiates data pull and uses population-flow to populate CIs from RTSM into CMS using the data flow probe connected to the CMS to perform the synchronization.



In the following diagram, RTSM initiates data pull and uses population-flow to populate CIs from CMS into RTSM using the data flow probe connected to RTSM to perform the synchronization.



The side initiating (defining the integration point) the population or the push synchronization needs a running dedicated probe connected to it, as shown in the diagrams above.

Install the probe version that is included with the specific installation media. The probe version must be the same version as the target to which it connects. It is crucial not to mix probe and APM versions.

Although technically, the same Data Flow Probe can be used for running both data synchronization and Discovery jobs, we recommend having a dedicated Data Flow Probe for data synchronization flow. The Data Flow Probe that is used for data synchronization can be installed on one of the APM server machines. If a customer has a two-server APM deployment (Gateway and Data Processing), we recommend installing the Data Flow Probe on the Gateway machine. If there are several Gateway machines, the probe should be installed on one of them.

The Data Flow Probe can be downloaded from the APM Downloads page (Select **Admin > Platform > Setup and Maintenance > Downloads**), and does not require an additional license unless you are running discovery.

Linux UD setup is also supported for CMS-RTSM synchronization.

Note: Federation synchronization does not use a Data Flow Probe, and is run directly from the Data Processing Server (DPS) side.

Handling Class Model Differences

The topology synchronization mechanism assumes that the class model used for defining the synchronization scope is identical at both the source and the destination. It is permissible to have additional attributes/classes/relationships on either the source or destination RTSM/UCMDB, as long as the additions are not part of the data synchronization scope defined by the data synchronization TQLs.

- If you pull the topology from UCMDB to APM 9.1x and above, all class model differences are ignored. In other words, CITs and attributes that do not exist on one of the synchronization sides are not synchronized, and the integration job continues running.
- For all other versions, class extensions should be specifically excluded from the data synchronization TQL. This means that if you have extended class A on source UCMDB by creating class B that inherits from A, you can no longer synchronize class A without specifically excluding class B from the data synchronization TQL.

The synchronization feature enables you to control which CI attributes are synchronized. In order to modify the default settings, select the desired attributes under the integration TQL layout for the TQL used in the synchronization. For more information, see *Layout Setting Dialog Box* in the *Modeling Guide*.

Effective Methods for Aligning the Class Model during Synchronization

When performing any type of synchronization, we highly recommend aligning the class model (CI type definitions and valid relationships) on the synchronizing side. This should be performed when CI Types, attributes, or valid relationships that are required to be synchronized, are not aligned or do not exist on both synchronization sides. This is sometimes the result of modification/customization to the class model which is not consistent on both sides, or large gaps in content packs between the synchronization sides.

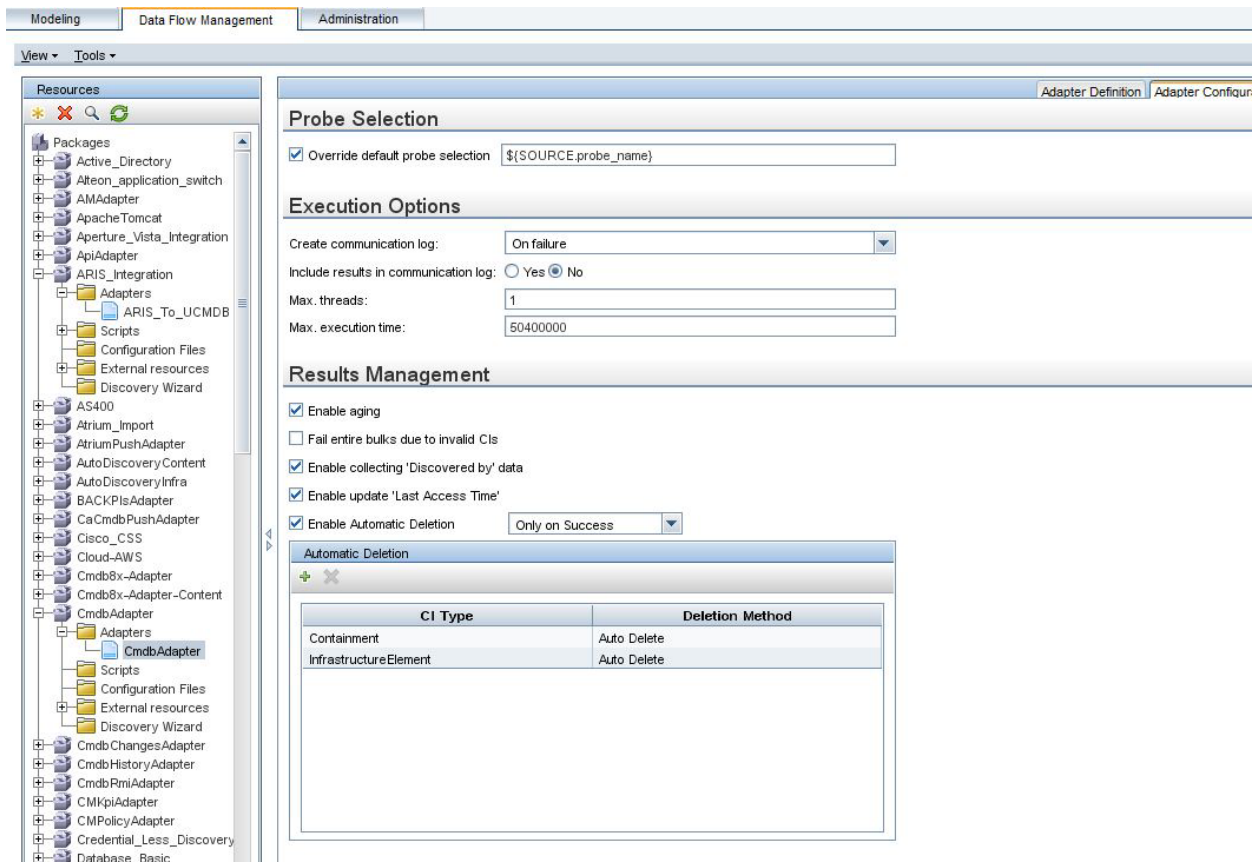
The best method for performing and maintaining this alignment, is to export the required content into a package using the RTSM or CMS package manager. Once a package is created, it should be deployed on the other synchronization side using the same package manager.

For more information about using the package manager, see *Package Manager* in the *RTSM Administration Guide*.

Deletion Policy

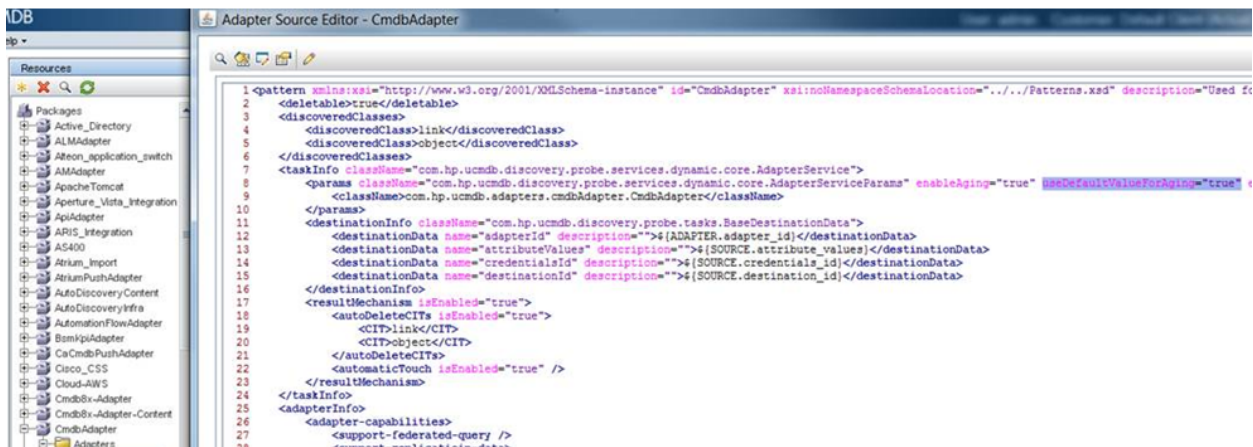
Deletion policy is part of the data synchronization configuration. It defines what to do when a CI that has been part of the data synchronization scope is deleted. Because APM has a tight integration with its service model, the deletion of some of the CIs can badly affect existing configurations of SLAs, alerts, downtime, and more.

The default deletion policy provided in the RTSM synchronizes the deletion of infrastructure CIs and containment relationships only. The default behavior can be modified by editing the configuration provided as part of the *CmdAdapter* that is deployed in the RTSM.



Aging Settings

By default, the aging settings of each CIT is used (the *enable aging* attribute is not synchronized). You can override the default settings, and force aging for every CIT which is synchronized, by setting `enableAging=true` and `useDefaultValueForAging=false` in the adapter XML or turn off aging altogether.



Importing CI Instances from Previous Versions of RTSM

You can manually import CI instances from previous versions. However, keep in mind that there are CI type differences from previous APM versions to the current version. Therefore, after exporting the CI instances from the previous version to an Excel file, you will need to modify the Excel file accordingly. Check the CI Type Manager for details.

Note: Topology synchronization is not supported for APM and UCMDB versions prior to version 9.0

To import CI instances from previous RTSM versions:

1. Export the required application instances using a CSV/Excel based report. For details, see the appropriate version of the APM Modeling Guide.
2. Manually reformat the exported file. For details, see How to Set Up Import File in Excel in the [HPE UCMDB CP11 Discovery and Integration Content Guide](https://softwaresupport.hpe.com/km/KM1443574) (<https://softwaresupport.hpe.com/km/KM1443574>) (requires HPE Passport login).
3. Use the CSV/Excel adapter to import it into RTSM using BMS 9.x. For details see Import from Excel Workbook Job in the [HPE UCMDB CP11 Discovery and Integration Content Guide](https://softwaresupport.hpe.com/km/KM1443574) (<https://softwaresupport.hpe.com/km/KM1443574>) (requires HPE Passport login).

Chapter 7: Reconciliation

Reconciliation process is designed to update the database model and avoid duplicate CIs in the APM RTSM. Reconciliation is the process of identifying and matching entities from different data repositories (for example, Operations Manager, SiteScope, UCMDB Discovery, Discovery and Dependency Mapping, Network Node Manager, and Real User Monitor).

Many different data collectors can send CIs to RTSM. In actuality, each different source might be providing information about the same CI. The reconciliation engine is responsible for identifying and matching entities from different data collectors and storing them, without duplicating CIs, in RTSM.

For more information, see Reconciliation in the Data Flow Management Guide.

Auto Complete Reconciliation

Auto Complete Reconciliation allows synchronizing additional CIs that are not part of the integration TQL, but are related to CIs that are selected to be synchronized, according to the existing reconciliation rules.

By default, when you synchronize a CI, all related CIs are checked to see if they are needed for reconciliation. The CIs that are needed for reconciliation are synchronized as well. Even CIs that are not explicitly requested to be synchronized as part of the integration TQL are synchronized.

This feature is important for allowing the target system which is pulling the CIs to properly identify and reconcile the CIs based on all available data.

Reconciliation Priorities

Reconciliation Priority specifies how matched CIs are merged in RTSM. When several data collectors (DCs) update the CIT and their attributes, the reconciliation priority defines which DC will be the last updater.

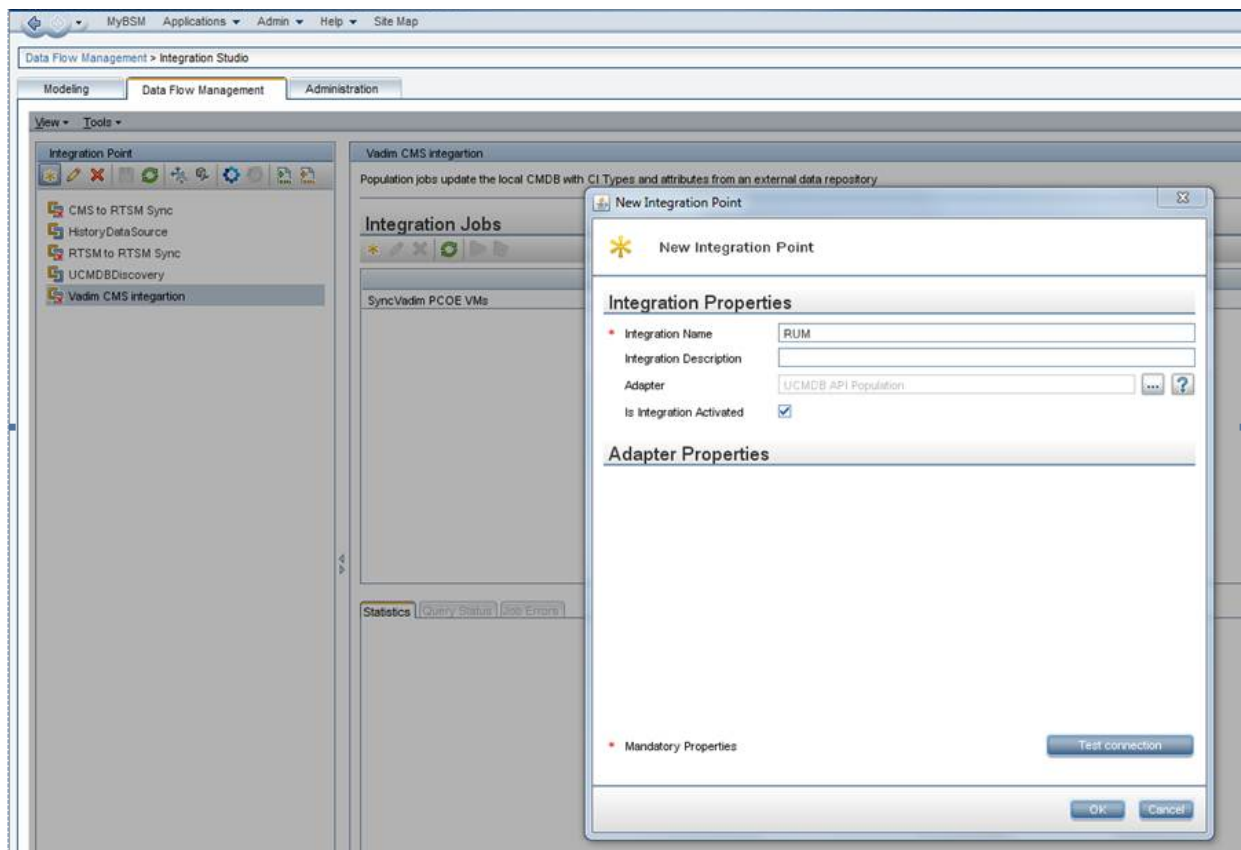
In order to define a priority for a certain APM data collector, you need to define an integration point specifically for the collector. To define an integration point which refers to a specific data collector, the name of the integration point must match the integration point name (also known as the data store origin). The following table lists the data collector and its integration point name.

| Data Collector | Integration point name (data store origin) |
|----------------|--|
| RUM | RUM |
| Diag | diagnostics |

In order to define a priority for a certain APM data collector, define an integration point specifically for the collector.

To define an integration point:

1. Select **Admin > RTSM Administration > Data Flow Management > Integration Studio**.
2. Enter the name of the integration point according to the data collector for which the data integration point is associated.
3. In the Adapter field, select **UCMDB API population**.
4. Click **OK**.

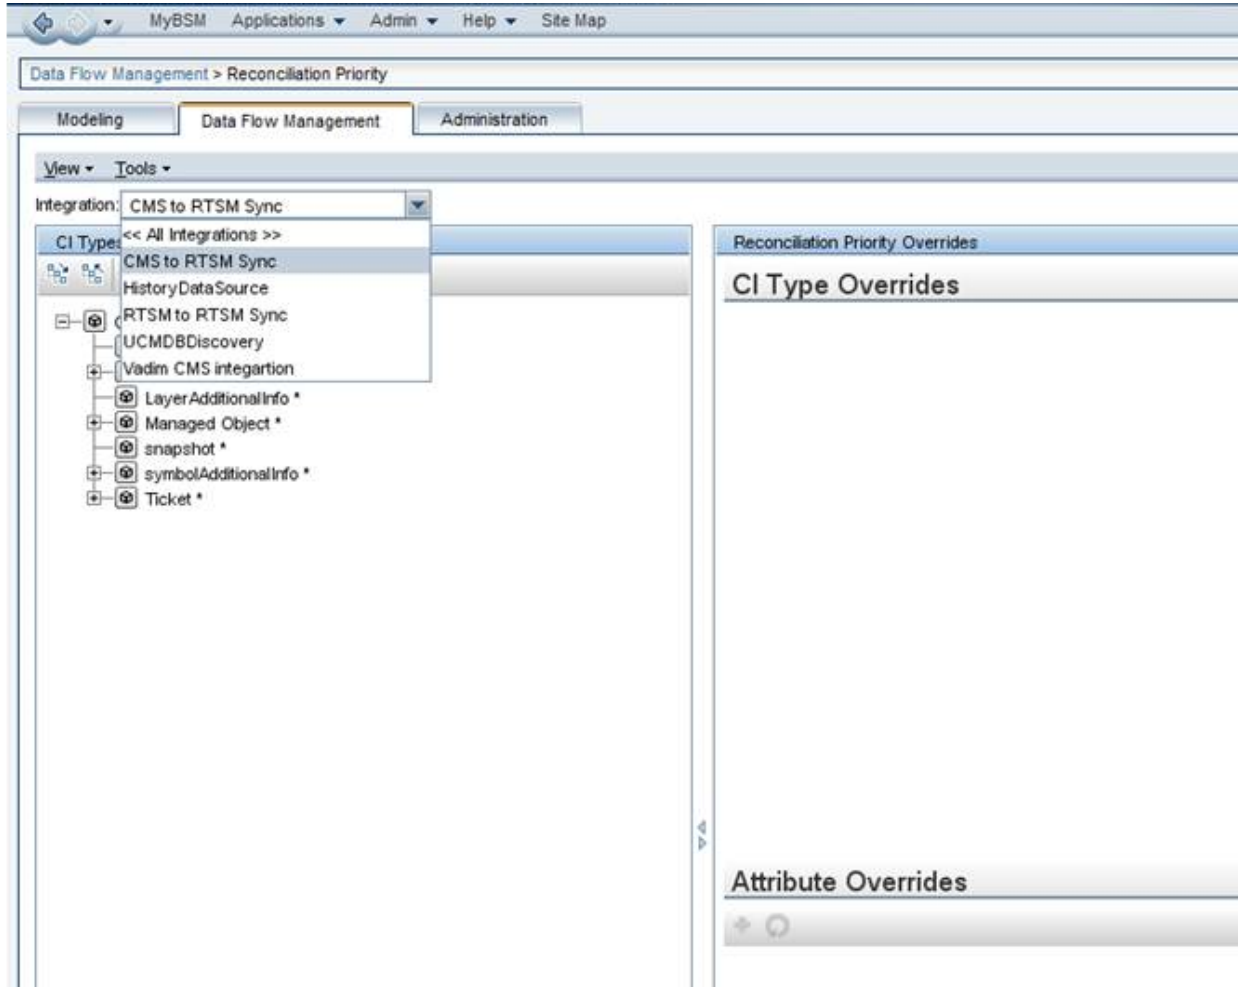


After creating the integration point, you can add it to the reconciliation priorities.

The reconciliation priorities are configured and managed under:

Admin > RTSM Administration > Data Flow Management > Reconciliation Priority.

It is enabled by defining a priority for an integration point.



Priorities are defined on the CIT level and, optionally, on the attribute level.

The screenshot shows the MyBSM Administration console with the 'Administration' tab selected. The left pane shows a tree view of 'CI Types' with 'ConfigurationItem' selected. The main area displays two tables: 'CI Type Overrides' and 'Attribute Overrides'.

CI Type Overrides Table:

| Integration Name | Priority | Inherited |
|-----------------------|----------|-----------|
| HistoryDataSource | 100 | Object |
| UCMDBDiscovery | 100 | Object |
| Vadim CMS integration | 100 | Object |

Attribute Overrides Table:

| Attribute | Integration Name | Priority | Inherited |
|-----------|------------------|----------|-------------------|
| Name | CMS to RTSM Sync | 100 | ConfigurationItem |

Chapter 8: Checking the Logs

The following logs contain information relating to the RTSM processes:

- **UCMDB:**
 - UCMDB\UCMDBServer\runtime\log\fcmdb.*
 - UCMDB\UCMDBServer\runtime\log\error.log
- **Probe:**
 - UCMDB\DataFlowProbe\runtime\log\fcmdb.*
 - UCMDB\DataFlowProbe\runtime\log\error log

On data push, the appropriate log files can be found on the data flow probe connected to the side initiating the push.

Log Configuration

You can configure the logs in the following files:

- UCMDB\- UCMDB\

Population Synchronization Troubleshooting

If an population integration does not work correctly, first make sure the data flow probe is up and connected to the UCMDB. To troubleshoot population synchronization, examine the contents of the probe's log connected to the side initiating the synchronization.

Occasionally, synchronization causes reconciliation issues. In those occasions, reconciliation log files can help pinpoint the problem.

The following are the reconciliation logs:

- cmdb.reconciliation.audit.log
- cmdb.reconciliation.log
- cmdb.reconciliation.datain.ignored.log

The location of the reconciliation log files is **<APM DPS install root>\odb\runtime\log**.

To change the log level go to:

<APM DPS install root>\odb\runtime\log\reconciliation.properties.

Appendix A: Acronym Table

The following acronyms are used in this document:

| Acronym | Meaning |
|----------------|---|
| APM | Application Performance Management |
| BPM | Business Process Monitor |
| APM | Application Performance Management |
| CI | Configuration Item |
| CIT | Configuration Item Type |
| CLIP | Closed Loop Incident Process |
| CMDB | Configuration Management Database |
| CMS | Configuration Management System |
| MDW | Monitor Deployment Wizard |
| MoM | Manager of Managers |
| RTSM | Run-time Service Model |
| RUM | Real User Monitor |
| SLA | Service Level Agreement |
| TQL | Topology Query Language |
| UCMDB | Universal Configuration Management Database |
| UD | Universal Discovery |

Appendix B: CITs

The following is a list of CI types that are composed of the topology reported by APM data collectors.

| Class Name | Description | Superclass | Producers | Consumers |
|-------------------------|--|----------------------|------------------|-----------------------|
| Agent | virtual class to all agents | RunningSoftware | SiS | SiS |
| BusinessApplication | A collection of software components that can be managed as an independent unit that supports a particular business function. An application is a logical composition of the functionality required to manipulate the data and provide the functional requirements of underlying business processes. An application has a set of supporting infrastructure entities. Example: Call center application, trade application etc. | BusinessElement | SiS, Diagnostics | RUM, SiS, Diagnostics |
| BusinessService | Represents business or IT service. A business service is a service that a business provides to another business (B2B) or that one organization provides to another within a business (e.g. payment processing). An IT service is a business service that the IT organization provides to support business services or IT's own operations. | Service | SiS | SiS |
| BusinessTransaction | An end-to-end service provided to end-users or to other related applications (e.g. stock trade, account login, product purchase or billing inquiry), in the IT realm will be provided by an application. Each business transaction has a unique profile that contains a business context (i.e. activity type, user, time and location, etc.) and a technical profile. | BusinessElement | Diagnostics | Diagnostics |
| BusinessTransactionFlow | BusinessTransactionFlow | Collection | RUM, Diagnostics | RUM, Diagnostics |
| Computer | This class represents a general purpose machine which has an IP address, such as Windows, Unix, Mainframe | Node | SiS | SiS |
| ConfigurationDocument | A document that contains a block of arbitrary information, or resource for storing information, which is available to a computer program and is usually based on some kind of durable storage | Application Resource | SiS | SiS |
| Database | A system that manages a collection of records arranged in a predefined structure and format allowing an efficient retrieval and search of data usually by key data items. | RunningSoftware | Diagnostics | Diagnostics |
| Datacenter | Represents a Data center entity | BusinessElement | SiS | SiS |
| DB2 | A database from IBM that serve a number of different operating system platforms. | Database | SiS | SiS |
| Diagnostics Probe | A Diagnostics probe deployed on a host | Monitor | Diagnostics | Diagnostics |
| Diagnostics Probe Group | A logical group of Diagnostics Probes | Collection | Diagnostics | Diagnostics |
| EMS Group | EMS group | Group | SiS | SiS |

| Class Name | Description | Superclass | Producers | Consumers |
|-----------------------|--|------------------------|-----------------------|------------------------|
| EMS Measurement | EMS event | System Monitor | SiS | SiS |
| EMS Monitor | EMS Monitor | System Monitor | SiS | SiS |
| End User Group | Logic group which separates content from the monitors | BusinessElement | RUM | RUM |
| End User Subgroup | End User Subgroup | End User Group | RUM | RUM |
| HP Diagnostics Agent | A Diagnostics probe deployed on a host | Agent | Diagnostics | Diagnostics |
| IBMWebSphereMQ | Represents IBM WebSphere MQSeries Software | MessagingServer | Diagnostics | Diagnostics |
| IIS Application Pool | IIS Application Pool | IIS Resource | Diagnostics | Diagnostics |
| IIS Virtual Dir | IIS Virtual Dir | IIS Web Dir | Diagnostics | Diagnostics |
| IIS Web Dir | IIS Web Dir | IIS Resource | Diagnostics | Diagnostics |
| IIS Web Server | This class represents an instance of IIS web-server (Microsoft Internet Information Server). | WebServer | SiS, Diagnostics | SiS, Diagnostics |
| IIS Web Service | IIS Web Service | IIS Service | Diagnostics | Diagnostics |
| IIS Web Site | IIS Web Site | IIS Resource | Diagnostics | Diagnostics |
| InfrastructureService | An IT Service that is not directly used by the Business, but is required by the IT Service Provider so they can provide other IT Services. For example Directory Services, naming services, or communication services. Infrastructure service is a synonym for technical service in ITILv3. | Service | SiS | SiS |
| Interface | Describes a logical interface (Note: This is NOT a physical interface card) that supports various data link layer and other higher layer protocols such as tunnels, mpls, frame, atm etc. | NodeElement | Diagnostics | Diagnostics |
| IpAddress | This class represents the logical network identifier of a node on the network. An IP Address can represent either an IPv4 or an IPv6 address. | NetworkEntity | RUM, SiS, Diagnostics | RUM, SiS, Diagnostics, |
| IpServiceEndpoint | Represents any kind of a network service end point based on a variety of protocols such as UDP or TCP or more specific ones such as SMTP, SNMP etc. The class describes the IP address and network port a service end point is bound to and the type of protocol it uses to expose an end point. | CommunicationEndpoint | RUM, SiS, Diagnostics | RUM, SiS, Diagnostics |
| J2EE Application | J2EE application - is a deployable unit of J2EE functionality Key Attributes: 1. Name (data_name) - A unique name that identifies application in the J2EE domain. 2. Container (root_container) - The relevant J2EE Domain | J2EE Deployed Object | Diagnostics | Diagnostics |
| J2EE Domain | Management domain of Java EE resources (i.e. servers, clusters, application components, etc.) | ApplicationSystem | Diagnostics | Diagnostics |
| J2EE Server | The J2EE server identifies the server core of one instance of a J2EE platform product as described in the Java 2 Enterprise Edition Platform specification. | ApplicationServer | Diagnostics | Diagnostics |
| JBoss AS | The Jboss Application Server | J2EE Server | Diagnostics | Diagnostics |
| MQ Queue | A Queue is a container of messages in the MQ Infrastructure | Message Queue Resource | Diagnostics | Diagnostics |

| Class Name | Description | Superclass | Producers | Consumers |
|-----------------------------|--|-----------------------|------------------|------------------|
| MSSQL Database | SQL Data Base from Microsoft | Database Schema | Diagnostics | Diagnostics |
| Node | The Node class represents a general purpose machine (i.e., computer). This also is a base class from which all the other classes representing devices such as virtual machines will inherit. Machines are typically reachable via the network but there are times when users model a machine even when it is not reachable via the network. Examples of machines include UNIX/Windows systems, switches, routers, firewalls etc. | InfrastructureElement | RUM, SiS | RUM, SiS |
| NTCMD | xcmd or pstools | Shell | SiS | SiS |
| Oracle | Oracle database | Database | SiS, Diagnostics | SiS, Diagnostics |
| Oracle iAS | The Oracle IAS Application Server | J2EE Server | Diagnostics | Diagnostics |
| Process | An instance of a program. | NodeElement | Diagnostics | Diagnostics |
| RunningSoftware | This class represents the runtime aspects of a software system that is currently running or intended to be running on a Node. An instance of the RunningSoftware class is the place for runtime overview information such as last-startup-time and application-instance-name. The rather static characteristics of an installed software are represented by the InstalledSoftware class. | InfrastructureElement | RUM, SiS | RUM, SiS |
| SAP ABAP Application Server | SAP's integrated software solution for client/server and distributed open systems. | SapApplicationServer | SiS, Diagnostics | SiS, Diagnostics |
| SAP J2EE Application Server | Represents an instance of the J2EE Application server | SapApplicationServer | Diagnostics | Diagnostics |
| SAP System | A logical unit, grouping together SAP-related entities (and possibly other entities as well) into one homogenous SAP deployment | ApplicationSystem | SiS | SiS |
| SAP Work Process | A logical single-instance representation of a certain type of work process | SAP Resource | SiS | SiS |
| Siebel Application Server | An application running the business logic tier | ApplicationServer | SiS | SiS |
| Siebel Component | A process on the Siebel Application Server encapsulating some Siebel application functionality | Application Resource | SiS | SiS |
| Siebel Component Group | An application running on the Siebel application server. | Application Resource | SiS | SiS |
| Siebel Enterprise | A logical unit, which describes a group of servers that function together to build a full-blown Siebel toolset experience | ApplicationSystem | SiS | SiS |
| Siebel Gateway | A coordinating server, which routes requests to the correct component | ApplicationServer | SiS | SiS |
| Siebel Web Application | A Siebel application's location on the web server | Application Resource | SiS | SiS |
| Siebel Web Server Extension | A server running the Web tier of Siebel | Application Resource | SiS | SiS |
| SiteScope Group | SiteScope Logical Group | Group | SiS | SiS |
| SiteScope Measurement | Represents a SiteScope measurement. | System Monitor | SiS | SiS |
| SiteScope Measurement Group | A group of measurements of the same type | Group | SiS | SiS |

| Class Name | Description | Superclass | Producers | Consumers |
|-------------------------------|--|---------------------|------------------|------------------|
| SiteScope Monitor | SiteScope Monitor | System Monitor | SiS | SiS |
| SiteScope Profile | SiteScope Profile | Group | SiS | SiS |
| SiteScope Profile Monitor | SiteScope Profile Monitor (Internal) | System Monitor | SiS | SiS |
| SiteScope Web Service Monitor | SiteScope Web Service Monitor | SiteScope Monitor | SiS | SiS |
| SQL Server | SQL server | Database | SiS, Diagnostics | SiS, Diagnostics |
| SSH | ssh agent | Shell | SiS | SiS |
| Telnet | telnet | Shell | SiS | SiS |
| Unix | Host with Unix Operation system. | Computer | SiS, Diagnostics | SiS, Diagnostics |
| Virtualization Layer Software | Represents the thin software layer running on bare metal. It can represent a kernel of the ESX Server's hypervisor for Vmware virtualization solution or any other hypervisor for other vendors | Hypervisor | SiS, Diagnostics | SiS, |
| VMware Cluster | VMware Cluster is a collection of ESX Server hosts with shared resources and a shared management interface. When you add a host to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts | FailoverCluster | SiS | SiS |
| VMware ESX Server | VMware ESX Server is a computer that uses virtualization software, such as ESX Server 3.5 or ESX Server 3i, to run virtual machines. ESX Hosts provide CPU and memory resources, access to storage, and network connectivity to virtual machines that reside on them | Computer | SiS, Diagnostics | SiS |
| VMware Resource Pool | VMware Resource Pool represents a pool where Virtual machines execute in, and draw their resources from | Resource Pool | SiS | SiS |
| Web Service | Represent a web service | WebService Resource | SiS, Diagnostics | SiS, Diagnostics |
| Web Service Operation | Web Service Operation | WebService Resource | SiS, Diagnostics | SiS, Diagnostics |
| Weblogic AS | The Weblogic Application Server | J2EE Server | SiS, Diagnostics | SiS, Diagnostics |
| WebServer | A software that accepts http requests and delivers as a response web pages. WebServers allow web-clients to access the pages and resources of web-applications. Examples: 'Apache web Server' , 'IIS'. | RunningSoftware | SiS | SiS |
| Websphere AS | The Websphere Application Server | J2EE Server | SiS, Diagnostics | SiS, Diagnostics |
| Windows | Host with Microsoft Operation system (NT). | Computer | SiS, Diagnostics | SiS, Diagnostics |

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 Best Practices (Application Performance Management 9.40)

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 docteam@hpe.com.

We appreciate your feedback!