

HP Business Service Management

For the Windows and Linux operating systems

Software Version: 9.23

RTSM Best Practices

Document Release Date: December 2013

Software Release Date: December 2013



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 2005 - 2013 Hewlett-Packard Development Company, L.P.

Trademark Notices

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

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

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

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

iPod is a trademark of Apple Computer, Inc.

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

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

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

UNIX® is a registered trademark of The Open Group.

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

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

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

UNIX® is a registered trademark of The Open Group.

Acknowledgements

This product includes software developed by the Apache Software Foundation (www.apache.org).

This product includes software developed by the JDOM Project (www.jdom.org).

This product includes software developed by the MX4J project (mx4j.sourceforge.net).

Documentation Updates

The title page of this document contains the following identifying information:

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

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

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: <http://h20229.www2.hp.com/passport-registration.html>

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

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

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.

Support

Visit the HP Software Support Online web site at: <http://www.hp.com/go/hpsupportsupport>

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

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

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

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

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

To find more information about access levels, go to:

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

HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is <http://h20230.www2.hp.com/sc/solutions/index.jsp>

Contents

Contents	4
Chapter 1: Introduction	6
Chapter 2: The Role of the RTSM	7
RTSM Capabilities	8
Enriching RTSM Data	8
Manual, Semi-automated, and Automated Enriching Methods	9
Chapter 3: RTSM Data Feeding Methods	11
Chapter 4: RTSM Integration Strategy	12
Global ID	12
Determining which RTSM/UCMDB is Defined as CMS	13
Preventing Multiple CMSs	13
Initial Synchronization	13
Chapter 5: BSM Hierarchical Deployment	15
Global ID in Hierarchical Deployments	16
Topology Discovery in a Manager of Managers Hierarchy	17
Chapter 6: BSM-UCMDB Topology Synchronization	19
Population, Data Push, and Federation	19
When to Use UCMDB-BSM Synchronization	20
Integration Prerequisites	21
Data Flow Probe Installation and Deployment	22
Handling Class Model Differences	24
Effective Methods for Aligning the Class Model during Synchronization	25
Deletion Policy	25
Aging Settings	26
Chapter 7: Reconciliation	27
Auto Complete Reconciliation	27
Reconciliation Priorities	27
Chapter 8: Checking the Logs	32

Log Configuration 32

Population Synchronization Troubleshooting 32

Appendix A: Acronym Table 33

Appendix B: CITs 34

We appreciate your feedback! 45

Chapter 1: Introduction

The goal of this document is to provide best practices for integrating the HP BSM Run-time Service Model (RTSM) with other products, or with other BSM 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 HP BSM 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.

BSM 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 HP's CMDB product, the HP Universal CMDB (UCMDB), embedded in BSM. 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 BSM and is an integral part of the BSM 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 BSM. 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 BSM 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 BSM.

RTSM also serves a specific set of use cases supporting the requirements of consolidated event management, provided by BSM'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 BSM; BSM 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 ["BSM-UCMDB Topology Synchronization" on page 19](#).

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 BSM operational solutions, such as BPM, RUM, Diagnostics, TransactionVision, SiteScope, HPOM SPIs, NNMi, and so on.
- **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 BSM operational solutions.

Enriching RTSM Data

The various monitoring products included in BSM 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 UCMDB complement each other and should be considered, along with Discovery technologies, to be elements of a CMS.

Within the CMS, a UCMDB instance acts as the central CMDB (this is the CMDB that generates the global IDs – for more information, see ["Global ID" on page 12](#)). Other UCMDB 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 BSM, separate from the rest of the CMS, allows for independent management of downtime, and the upgrade of BSM deployments, without affecting the UCMDB and the rest of the HP products that are integrated into the CMS. The RTSM class model is an extension of the UCMDB class model, and allows for the management of the BSM private class model extensions that are required for BSM operational use cases.

Other use cases, which revolve around traditional CMDB services, such as configuration and change management, may require Discovery and Dependency Mapping (DDMA/UD) services. (UD is the next generation of DDMA.) When an end-to-end solution is deployed, DDMA services

are typically used to populate a central UCMDB, and the RTSM can leverage richer data from the UCMDB where applicable.

The RTSM is a highly optimized service model, created for the high performance needs of the BSM use cases, and as such, includes the CIs that are required for complete and accurate monitoring. Other CIs which are not crucial for the BSM/APM use cases should be stored in the central UCMDB 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 BSM 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 BSM 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.

Note: Similar modeling can be achieved in the CMS as well, but we recommend exporting and importing pattern-based models and enrichments, instead of synchronizing high level and business level CIs, which can result in over populating RTSM due to the auto complete mechanism that synchronizes irrelevant topology (see ["Auto Complete Reconciliation" on page 27](#)).

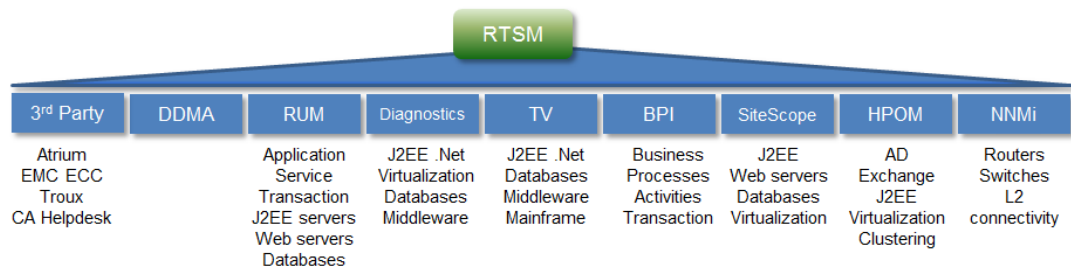
For details on pattern/instance-based modeling and topology enrichment methods, see the Modeling Guide and RTSM Best Practices.

Chapter 3: RTSM Data Feeding Methods

Note: RTSM is designed to support BSM 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:

- BSM 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 12](#).
- 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 19](#).


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 BSM and other HP centers (for example, HP Service Manager), we recommend that you install a central UCMDB instance as part of a CMS, rather than integrate the other HP center directly with the RTSM. We also recommend that you connect DDMA to the central UCMDB instance, and then synchronize the topology relevant for BSM operational use cases with BSM, as necessary.

BSM provides a set of predefined data synchronization Topology Query Languages (TQLs) for pulling the topology from the central UCMDB to BSM. 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.

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

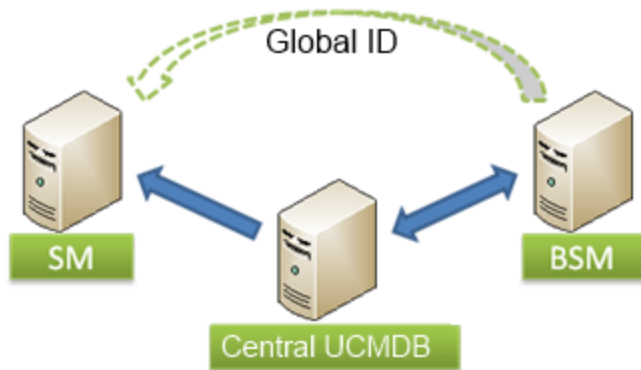


Direct integration with configuration management products (such as HP Release Control) is not supported by the RTSM and requires the installation of a CMS or dedicated UCMDB instance.

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 HP centers (for example, a new incident is opened in HP 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 a BSM hierarchical deployment with no CMS, where the top instance of the RTSM is required as 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 BSM 9.10 and above.

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

Chapter 5: BSM Hierarchical Deployment

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

- **Event related scaling.** When a hierarchy of BSM 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 BSM server is planned to reach BSM’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 BSM 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 BSM instance. The consolidated picture is achieved either by two-way synchronization between BSM instances, or by defining one central MoM instance.
- **Functional structure.** When the IT Department chooses to manage applications and infrastructure separately, by creating separate BSM instances for application owners and for infrastructure owners. In this deployment scenario, there can be several BSM instances, each one operated by domain experts, for example, APM (performed by BSM), NNMi, HPOM.



For details on how to set up the NNMi-BSM integration, refer to “NNMi-BSM Topology Integration - Best Practices” in the BSM Help.

- **Organizational mergers and acquisitions.** When there are several BSM 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 BSM instances under MoM deployment.
- Push synchronization to synchronize between CMS into BSM instances, according to the support matrix.

Global ID in Hierarchical Deployments

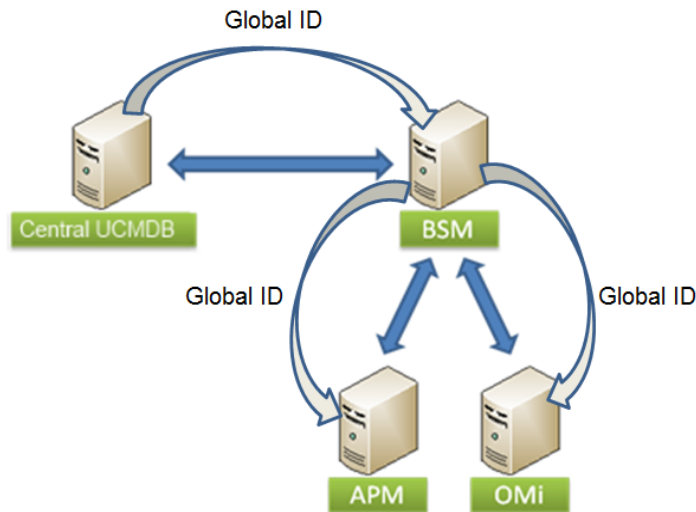
One of the main motivations for synchronizing the topology in hierarchical deployments is to provide a topological context for the forwarded events. To resolve the incoming event, and find the associated CI, the information available in the event is used. Part of this information is the global CI ID that provides the shared topological context between two BSM 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 BSM instance as a global CI ID generator and use pushback of IDs towards each of the child BSM 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 BSM instance, and between the top BSM instance and its child BSMs.



- **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 BSM instance, and then from the top BSM instance to its child BSMs. This entails a full synchronization between the central UCMDB and the top BSM instance, and again between the top BSM instance and its child BSMs.

Note: The central UCMDB is by default the global CI ID generator. The UCMDB cannot generate global IDs for BSM 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 BSM instance.

Topology Discovery in a Manager of Managers Hierarchy

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

1. The global BSM Probe (MoM) pulls the topology from the local (child) BSMs, or alternatively, the local BSM pushes the topology into the global BSM. We recommend using the Push method.

2. The central CMDB pulls the topology from the global BSM, and generates global IDs, or alternatively, the global BSM pushes topology into the CMS. We recommend using the Push method.
3. The global BSM pulls the topology from the central CMDB, or alternatively, the CMS pushes the topology into the global BSM. Push is the recommended method when using the latest versions.
4. The local BSMs pull topology from the global BSM, or alternatively, the global BSM pushes the topology to the local BSM.
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 BSM 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: BSM-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 BSM, see the BSM-CMS Synchronization Integration Matrix in the BSM 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.

Note: Although we recommend using Data Push for topology synchronization from CMS to BSM, it is supported only when integrating data from UCMDB/CMS 10.01 CUP5 and up to BSM 9.22 and above.

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

Population integration should generally be used for topology integrations between a central UCMDB and the RTSM, or between RTSM instances (see ["BSM Hierarchical Deployment" on page 15](#)).

A Federation integration type is used when BSM needs to retrieve information from other HP centers that are integrated with the central UCMDB, without replicating the data; for example, incidents from HP Service Manager, or planned changes from HP Release Control. For these use cases, there is no reason to replicate the data as BSM 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 HP offerings for BSM, SM, UCMDB, and OO. For details, see [Closed Loop Incident Process 9.30](#). 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. BSM provides an out-of-the-box integration point that includes all required configuration for out-of-the-box BSM 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 BSM" in the RTSM Data Flow Management Guide.

When to Use UCMDB-BSM Synchronization

The following sections describe use cases for topology synchronization.

Synchronizing the Topology from UCMDB to BSM

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

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

- BSM flows that are integrated with topology reported by DDMA:
 - **Vertical Solutions:** Application Management for SAP, Application Management for Siebel, Application Management for SOA.
 - **Monitor Deployment Wizard (MDW) in SiteScope:** The MDW uses the topology discovered by DDMA when configuring SiteScope monitors. The MDW retrieves relevant information from CIs reported by DDMA, and use it to help configure the monitors and assign the monitors to the relevant CIs.
 - **Consume DDMA credentials:** Once DDMA 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 BSM flows: MDW, integration with Operations Orchestration, monitor assignment for vertical solutions.
- Enrich the topology in BSM:

Some BSM data collectors report "shallow" topology. Synchronizing topology discovered by DDMA, or other UCMDB topology reporting sources, enriches the CIs reported by BSM. For example, BSM data collectors can report the nodes (hosts) using IP. Once those same nodes are discovered by DDMA, additional information is available, such as DNS name, OS, OS version, and so forth. This can help BSM 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 BSM to UCMDB

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

- BSM data collectors (such as RUM, Diagnostics, TransactionVision) 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.
- BSM 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 DDMA; BSM discovers a more “shallow” topology that is used as a trigger for DDMA 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 DDMA).

Setting Up UCMDB-BSM 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 BSM side, see “Set Up Integrations Between CMS and BSM” in the Data Flow Management Guide.
- For instructions on populating the data from BSM 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 content packs on both synchronization sides.
- Install a Data Flow Probe according to the instructions in ["Data Flow Probe Installation and Deployment" on the next page](#).

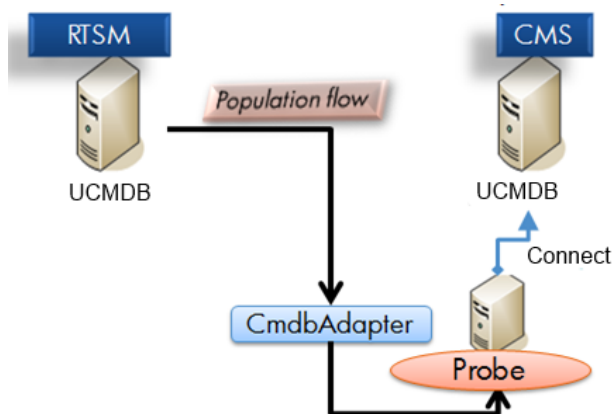
- Connect the Data Flow Probe to the side initiating the synchronization according to the instructions in ["Data Flow Probe Installation and Deployment"](#) below.
- For data population, prepare an integration TQL or use one of the pre-defined TQLs. Synchronization TQL should be defined on the source (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 24 and define the CITs and attributes to be synchronized accordingly.

Data Flow Probe Installation and Deployment

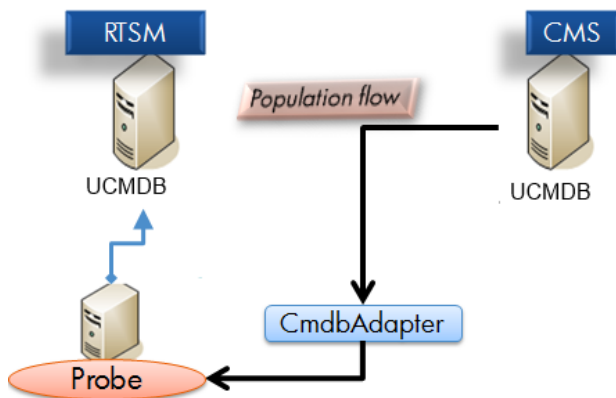
Topology synchronization requires installation of the HP Data Flow Probe. A single Data Flow Probe can be used to pull the topology from multiple UCMDB sources. For example, when using a hierarchical BSM deployment, the same Data Flow Probe can be used to pull the topology from several child BSM instances.

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 BSM versions.

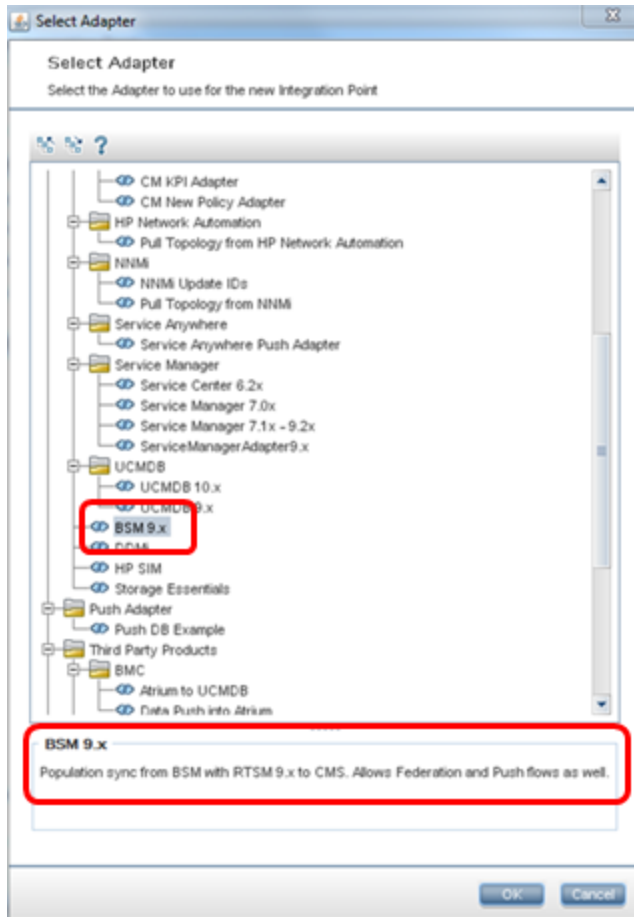
Although technically, the same Data Flow Probe can be used for running both data synchronization and Discovery jobs, we recommend having a dedicated Data Flow Probe for data synchronization flow. The Data Flow Probe that is used for data synchronization can be installed on one of the BSM server machines. If a customer has a two-server BSM 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 BSM Downloads page (Select **Admin > Platform > Setup and Maintenance > Downloads**), and does not require an additional license unless you are running discovery.

Linux DDMA 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.

For push synchronization, open the Integration Studio from the CMDB server, create a new integration point, and select the **BSM9.x** adapter. When configuring the integration point, select the **Data Push** tab.



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 BSM 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.

Note: Currently it is not possible to export content (i.e., CI Types) from UCMDB v10.x and import it into RTSM. Therefore, define the new content on the RTSM side and export and then import the content on the UCMDB 10 side.

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 BSM 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 CmdbAdapter 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.



Chapter 7: Reconciliation

Reconciliation process is designed to update the database model and avoid duplicate CIs in the BSM 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, Real User Monitor, TransactionVision, and Business Process Insight).

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 BSM 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
OMi/OM using topo-sync	BSM
*NNMi (using NNMi topology push)	According to the created integration user

Data Collector	Integration point name (data store origin)
NNMi (using DDMA based topology pull)	Integration point is name is not significant. Select Population from NNMi adapter in the integration adapters list
**SiteScope/BSM Connector	DiscoverySiteScope

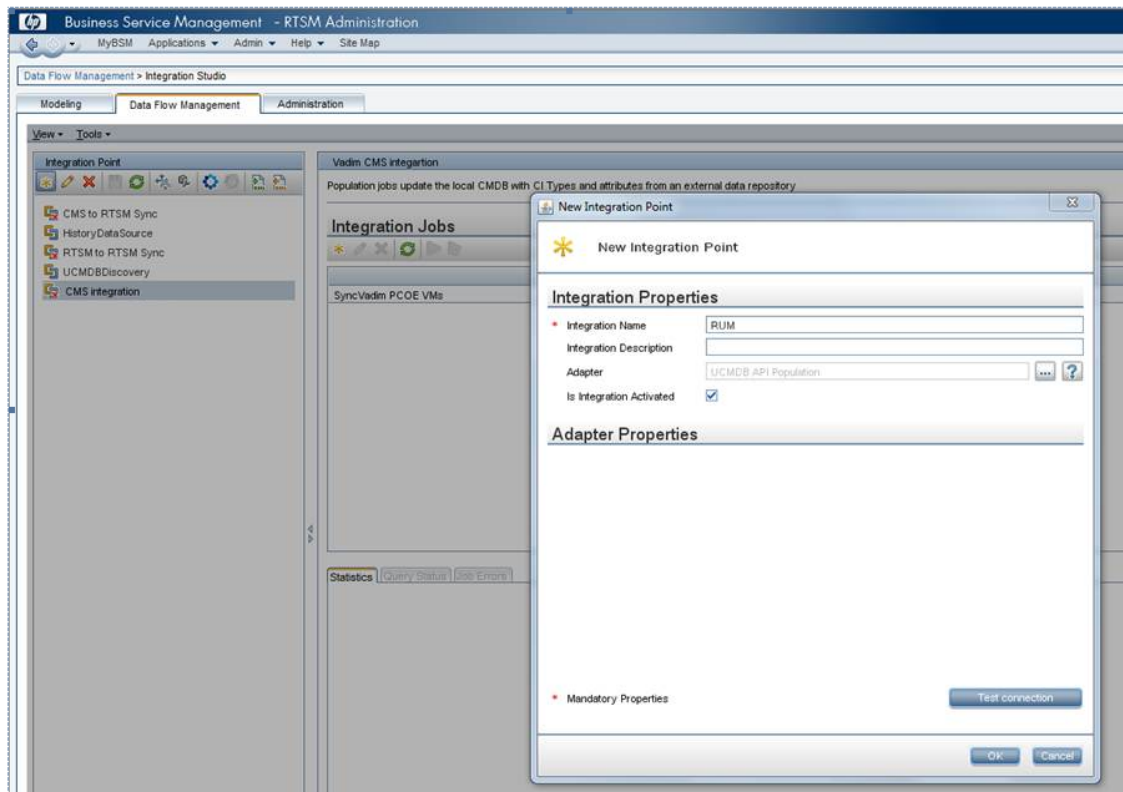
*As part of the integration with NNMi, the integration needs to be defined specifically for the NNMi integration. For details, see “NNMi–BSM Topology Integration - Best Practices” in the BSM Help.

****Limitation:** If you configure a monitor in SiteScope which integrates with one of the integrations under **Administration > Integrations > EMS Integrations Administration**, the “reconciliation prioritization” feature for SiteScope will not work, meaning SiteScope cannot be prioritized over other data collectors.

In order to define a priority for a certain BSM 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**.

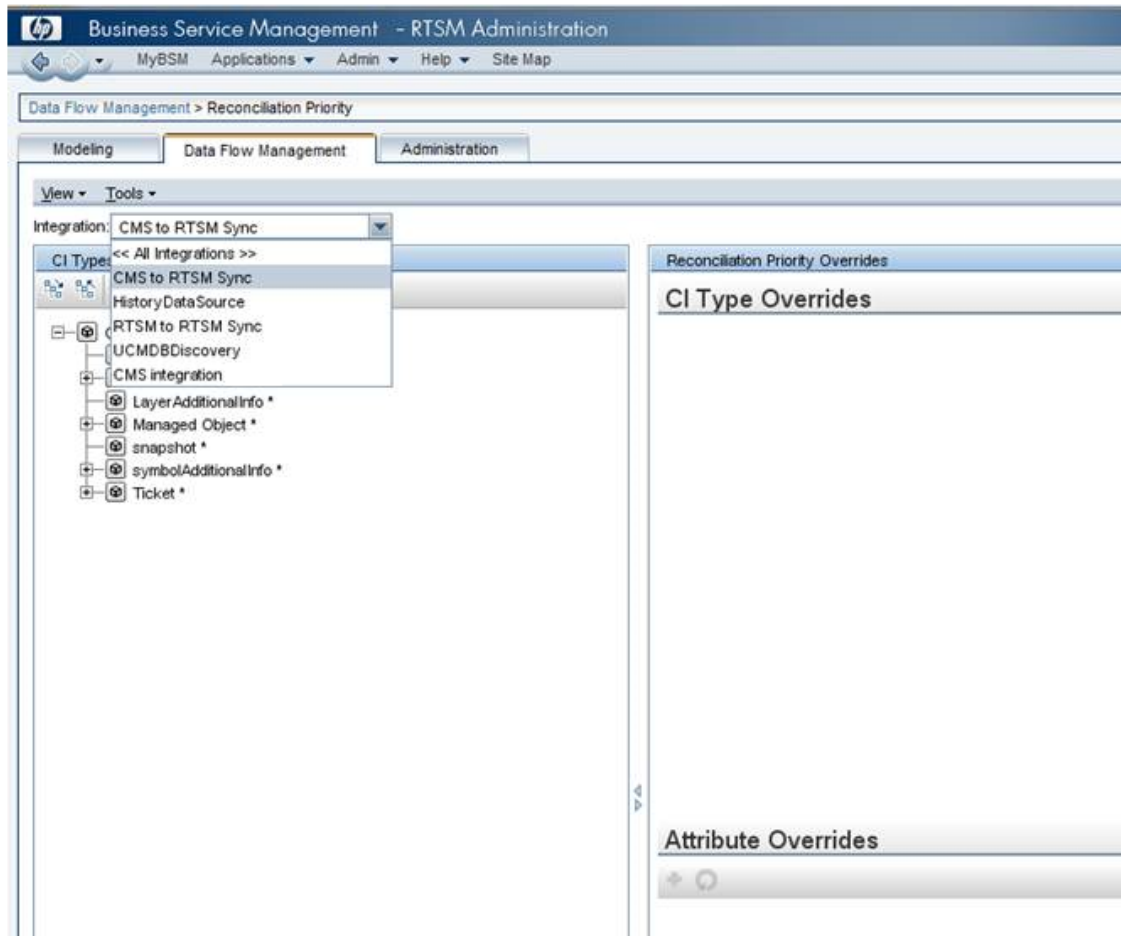


After creating the integration point, you can add it to the reconciliation priorities.

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 displays the HP Business Service Management - RTSM Administration interface. The main window is titled "Reconciliation Priority" and is divided into two main sections: "CI Type Overrides" and "Attribute Overrides".

CI Type Overrides

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

Attribute Overrides

Attribute	Integration Name	Priority	
Name	CMS to RTSM Sync	100	ConfigurationItem

The left sidebar shows a tree view of CI Types, with "ConfigurationItem" selected. The top navigation bar includes "MyBSM", "Applications", "Admin", "Help", and "Site Map". The "Data Studio" tab is active, and the "Reconciliation Priority" sub-tab is selected.

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\<UCMDBServer|DataFlowProbe>\conf\log\fcmdb.properties
- UCMDB\<UCMDBServer|DataFlowProbe>\conf\log\fcmdb.gdba.properties

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 **<BSM DPS install root>\odb\runtime\log**.

To change the log level go to:

<BSM 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
BPI	Business Process Insight
BPM	Business Process Monitor
BSM	Business Service Management
CI	Configuration Item
CIT	Configuration Item Type
CLIP	Closed Loop Incident Process
CMDB	Configuration Management Database
CMS	Configuration Management System
DDMA	Discovery and Dependency Mapping
HPOM	HP Operations Manager
HPOM SPIs	HP Operations Manager Smart Plug-ins
MDW	Monitor Deployment Wizard
MoM	Manager of Managers
NNMi	Network Node Manager i
OM	Operation Manager
RTSM	Run-time Service Model
RUM	Real User Monitor
SLA	Service Level Agreement
TQL	Topology Query Language
TV	TransactionVision
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 BSM data collectors.

Class Name	Description	Superclass	Producers	Consumers
.NetAppDomain	.NET AppDomain	Application Resource	Diagnostics	Diagnostics
Active Directory Domain	A domain is a partition in an Active Directory forest, and enables organizations to replicate data only to where it is needed. Since Active directory domains are organized in hierarchical structure Active Directory Domain" CIT can have one of the two types of container: "Active Directory Forest" or "Active Directory Domain"	Active Directory	OM	OM
Active Directory Forest	Active Directory Forest is a logical group of Active Directory Domains.	Active Directory	OM	OM
Active Directory Site	A Site object in Active Directory represents a physical geographic location that hosts networks. Sites contain objects called Subnets.	Active Directory	OM	OM
Active Directory System	Active Directory System collection of Forests, Domains and Domain Controllers.	Active Directory	OM	OM
ActiveDirectoryApplicationMode	Identifies the Active Directory Application Mode software	DirectoryServer	OM	OM
Agent	virtual class to all agents	RunningSoftware	SiS	SiS
Archiving Server	Archiving Server enables you to archive IM communications and meeting content for compliance reasons.	Lync Server Role	OM	OM
AV Conferencing Server	A/V Conferencing Server provides A/V conferencing functionality to your deployment. It can be colocated with Front End Server, or deployed separately as a single server or A/V Conferencing Server pool.	Lync Server Role	OM	OM
BridgeHeadServer	Bridgehead Servers are the contact point for the exchange of directory information between sites.	DomainControllerRole	OM	OM
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, OM, 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

Class Name	Description	Superclass	Producers	Consumers
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, TV	Diagnostics, TV
BusinessTransactionFlow	BusinessTransactionFlow	CiCollection	RUM, Diagnostics, TV	RUM, Diagnostics, TV
Central Management Server	The Central Management Store is a centralized configuration database that is used to save configuration data for the entire Lync Server deployment.	Lync Server Role	OM	OM
ClusterResourceGroup	Represents a cluster resource group (Cluster Package" in MC/ServiceGuard terminology) on a failover cluster. The CRG provides a runtime environment that is similar to a virtual node (network connectivity	Node	OM	OM
ClusterSoftware	The software that provides failover cluster capabilities on a certain node (e.g. the runtime software of MC/ServiceGuard on an HP-UX cluster node)	RunningSoftware	OM	OM
Computer	This class represents a general purpose machine which has an IP address, such as Windows, Unix, Mainframe	Node	SiS, OM	SiS, OM
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
Cpu	A central processing unit (CPU) is an electronic circuit that can execute computer programs.	NodeElement	OM	OM
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	OM, Diagnostics
Datacenter	Represents a Data center entity	BusinessElement	SiS, OM	SiS, OM
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	CiCollection	Diagnostics	Diagnostics
Director Server	Directors can authenticate Lync Server user requests, but do not home user accounts, or provide presence or conferencing services. Directors are most useful in deployments that enable external user access, where the Director can authenticate requests before sending them on to internal servers.	Lync Server Role	OM	OM

Class Name	Description	Superclass	Producers	Consumers
DirectoryServer	A directory service is the software system that stores, organizes and provides access to information in a directory. A directory service is a shared information infrastructure for locating, managing, administering, and organizing common items and network resources, which can include volumes, folders, files, printers, users, groups, devices, telephone numbers and other objects.	RunningSoftware	OM	OM
DiskDevice	A DiskDevice is a peripheral device used to record and retrieve information. Main implementations are hard disks, floppy disks and optical discs. They are identified by device name on a given node e.g. /dev/sd0 in UNIX	NodeElement	OM	OM
DnsServer	A DNS or a Domain Name Server is a registry system for electronic services and resources. The DNS translates domain names which human understand into the appropriate network or application address.	RunningSoftware		OM
DomainController	A domain controller is a directory server that physically store the Active Directory information. All domain controllers that belong to the same domain replicate information to each other. A domain controller which is defined as a global catalog server replicate and stores objects from all the domains that belong to the forest.	DirectoryServer	OM	OM
DomainController Resource	DomainController Resourc	Application Resource	OM	OM
DomainControllerRole	DomainController Role	DomainController Resource	OM	OM
DomainNamingMaster	The domain naming master domain controller controls the addition or removal of domains in the forest.	DomainControllerRole	OM	OM
Edge Server	Edge Server enables your users to communicate and collaborate with users outside the organization's firewalls.	Lync Server Role	OM	OM
EMS Group	EMS group	Group	SiS	SiS
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
Exchange Client Access Server	The Client Access server is the server that users connect to with their mail client, mobile device, or web browser. The Client Access server handles all connections whether they come from an application such as Outlook, Outlook Express, or any other MAPI, POP3 or IMAP4 client. The Client Access server also handles connections made from mobile devices such as a Windows Mobile 5 Smartphone, or any other device using Exchange ActiveSync. This role also provides Outlook Web Access (OWA).	Exchange role	OM	OM

Class Name	Description	Superclass	Producers	Consumers
Exchange Database Availability Group	A database availability group (DAG) is a set of up to 16 Microsoft Exchange Server 2010 Mailbox servers that provide automatic database-level recovery from a database, server, or network failure. Mailbox servers in a DAG monitor each other for failures. When a Mailbox server is added to a DAG, it works with the other servers in the DAG to provide automatic, database-level recovery from database, server, and network failures.	Exchange	OM	OM
Exchange Edge Server	The Edge Transport role is installed on the edge of the network and therefore is installed on a standalone server that is not a member of the Active Directory domain. Active Directory Application Mode (ADAM) is used to sync AD with the Edge Transport server. ADAM and a component called EdgeSync are used to perform scheduled one-way synchronization of the configuration and recipient information from Active Directory. This allows the Edge Transport to perform recipient lookups and Spam filtering.	Exchange Transport Server	OM	OM
Exchange Hub Server	The Hub Transport role is responsible for all internal mail flow. This role is similar to the bridgehead server in an Exchange 2000/2003 organization. The Hub Transport server is installed on member server(s) in an Active Directory domain. Because it is a member of an AD domain, all its configuration information is stored in AD and any other Hub Transport servers will get their configuration from AD.	Exchange Transport Server	OM	OM
Exchange Mail Server	Mailbox role holds the Exchange databases within which the user mailboxes are contained. It is also home to the Public Folder databases if you enabled Public Folders.	Exchange role	OM	OM
Exchange Mailbox Database	Exchange storage group - a logical container for Exchange databases and their associated system and transaction log files.	Microsoft Exchange Resource	OM	OM
Exchange Organization	This class represents Exchange Organization.	Exchange	OM	OM
Exchange role	Role of exchange server in global scope. One exchange server can have several roles. Server roles allow an administrator to split the functions of an Exchange server and place each role, or a combination of roles, on different servers in the organization. This type is a common super-type for specific roles, which are: - Edge Transport - Hub Transport - Client Access - Mailbox - Unified Messaging	Microsoft Exchange Resource	OM	OM
Exchange Storage Group	Exchange storage group - a logical container for Exchange databases and their associated system and transaction log files.	Microsoft Exchange Resource	OM	OM
Exchange Transport Server	This CIT is a common super-type for two transport roles: Edge role and Hub role.	Exchange role	OM	OM

Class Name	Description	Superclass	Producers	Consumers
Exchange Unified Messaging Server	This server role enables Unified Messaging for an Exchange 2007 organization. Unified Messaging lets users access their Exchange 2007 mailbox over any telephone for e-mail, voice mail, fax messages, and calendaring and contact information. The Unified Messaging role is responsible for merging VOIP infrastructure with Exchange organization. It provides the capability to: - combined voice, fax, and mail in one inbox - access to voice, fax and mail via multiple interfaces	Exchange role	OM	OM
FailoverCluster	A FailoverCluster is a Cluster that is implemented primarily for the purpose of providing high availability of services which the cluster provides. They operate by having redundant computers or nodes which are then used to provide service when system components fail.	Cluster	OM	OM
FileSystem	A file system mounted on a Node. A file system provides the means for storing and organizing computer files and the data they contain to make it easy to find and access them. This class corresponds to the general notion of file system used by Unix, Windows, Linux, etc.	NodeElement	OM	OM
Front End Server	The Front End Server is the core server role, and runs many basic Lync Server functions. The Front End Server, along with the Back End Servers that provide the database, is the only server role required to be in any Lync Server Enterprise Edition deployment.	Lync Server Role	OM	OM
GlobalCatalogServer	The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain Active Directory forest.	DomainControllerRole	OM	OM
HP Diagnostics Agent	A Diagnostics probe deployed on a host	Agent	Diagnostics	Diagnostics
HP Operations Agent	Key Attributes: 1. Name (data_name) - Core ID of the Operations Agent 2. Container (root_container) - The container Host Expected attributes: 1. Application IP - The primary IP that is used to communicate with the Agent 2. Application Listening Port Number - The port number that is used to communicate with the Agent (default: 383)	Agent	SiS, OM	OM
Hypervisor	Base class for different kinds of hypervisors (software that allows to host virtual machines) Represents the virtualization service provided by the Hypervisor, its properties and health. The Hypervisor is the link object between the virtual machine (a Node) and the physical host (a Node).	RunningSoftware	OM	OM
IBM MQ Alias Queue	An IBM Alias Queue is simply an alias of another queue. It may be an alias of a local, remote, transmission or another alias queue. The alias queue and the queue for which it is an alias are within the same queue manager. Messages/commands issued on the alias queue are forwarded to the queue for which it is an alias	IBM MQ Queue	TV	TV

Class Name	Description	Superclass	Producers	Consumers
IBM MQ Cluster	An MQ Cluster provides a flexible approach to join multiple queue managers with minimal configuration. This allows multiple instances of the same service to be hosted through multiple queue managers which allows for higher performance, capacity and resiliency	FailoverCluster	TV	TV
IBM MQ Local Queue	A Local Queue is a basic message queue and container of messages. An application can place a message in it for delivery or request/retrieve a message from it	IBM MQ Queue	TV	TV
IBM MQ Queue	The IBM MQ Queue is a container of messages in the MQ Infrastructure and controls how messages are routed between queue managers. Queues may be setup in several configurations to control message ordering and delivery (F/LIFO, message priority, sequential delivery, guaranteed delivery, etc.) and are optimized to carry small amounts of information	MQ Queue	TV	TV
IBM MQ Queue Manager	A WebSphere MQ instance may have one or more queue managers. The queue manager is responsible for functions not directly related to data movement such as storage, timing, triggering, etc. Queue Managers use a proprietary IBM technology known as a bindings" connection to communicate with MQ objects it manages and with remote clients via a network Key Attributes: 1. Name (data_name) - The name of the MQ Queue Manager 2. Container (root_container) - The container IBM WebSphere MQ Software Element"	Message Queue Resource	TV	TV
IBM MQ Remote Queue	A Remote Queue is a remote or proxy instance of another queue. It may be a remote instance for a local, remote, transmission or another alias queue. The remote queue and the queue for which it is a remote may be on different queue managers	IBM MQ Queue	TV	TV
IBMWebSphereMQ	Represents IBM WebSphere MQSeries Software	MessagingServer	Diagnostics, TV	Diagnostics, TV
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
InfrastructureMaster	The infrastructure is responsible for updating references from objects in its domain to objects in other domains.	DomainControllerRole	OM	OM
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

Class Name	Description	Superclass	Producers	Consumers
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	OM, NNMi, Diagnostics, TV	OM, Diagnostics, TV
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, OM, NNMi, Diagnostics, TV	RUM, SiS, OM, Diagnostics, TV
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.	CommunicationEnd point	RUM, SiS, Diagnostics, TV	RUM, SiS, Diagnostics, TV
IpSubnet	This class represents an IP subnet in a network. The network could be an IPv4 or an IPv6 network. An instance of this class will exist for each IP subnet in a routing domain.	NetworkEntity		OM
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	OM, Diagnostics, TV	OM, Diagnostics, TV
J2EE Cluster	Cluster of Java EE Servers that provides high availability, scalability and fault tolerance services	Load Balancing Cluster	OM	OM
J2EE Domain	Management domain of Java EE resources (i.e. servers, clusters, application components, etc.)	ApplicationSystem	OM, Diagnostics, TV	OM, Diagnostics, TV
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, TV	OM, Diagnostics, TV
JBoss AS	The Jboss Application Server	J2EE Server	Diagnostics	Diagnostics
JDBC Data Source	Represents JDBC connector to data source. JDBC Data Source provides database connectivity through a pool of JDBC connections	Application Resource	OM	OM
JMS Destination	jms destination	JMS Resource	TV	TV
JMS Server	jms server	JMS Resource	TV	TV
JVM	Key Attributes: 1. Name (data_name) - JVM 2. Container (root_container) - The container Software Element	J2EE Managed Object	OM	OM
Layer2Connection	Represents a ISO Layer-2 connection between 2 or more interfaces.	NetworkEntity	NNMi	OM
Lync Pool	A Lync Server pool is a collection of Lync Servers. All Lync Servers in the pool run exactly the same services and one server in a pool of many can go down without detribalizing the pool. A pool is defined by a full qualified domain name.	Lync	OM	OM

Class Name	Description	Superclass	Producers	Consumers
Lync Server	Microsoft Lync Server (previously Microsoft Office Communications Server) is an enterprise real-time communications server, providing the infrastructure for enterprise instant messaging, presence, file transfer, peer-to-peer and multiparty voice and video calling, ad hoc and structured conferences (audio, video and web) and PSTN connectivity.	Communication Server	OM	OM
Lync Server Role	Lync Server 2010 allows for role based deployment. During Lync Server 2010 installation user can select for the components to be installed on a machine. The services installed on the Lync Server identify the role on the server.	Lync Resource	OM	OM
Lync Site	A Lync Server site can be either a central site or a branch site. A central site contains at least one Front End pool or one Standard Edition server. A branch site is associated with exactly one central site, and the users at the branch site get most of their Lync Server functionality from the servers at the associated central site.	Lync	OM	OM
Mainframe	Main computer which can serve many users	Computer	TV	TV
Mediation Server	Mediation Server is a necessary component for implementing Enterprise Voice and dial-in conferencing. Mediation Server translates signaling and, in some configurations, media between your internal Lync Server infrastructure and an Internet Protocol/Public Switched Telephone Network (IP-PSTN) gateway or a Session Initiation Protocol (SIP) trunk.	Lync Server Role	OM	OM
MicrosoftExchangeServer	This class represents Microsoft Exchange server software installed on some host.	MailServer	OM	OM
Monitoring Server	Monitoring Server collects data about the quality of your network media, in both Enterprise Voice calls and A/V conferences.	Lync Server Role	OM	OM
MQ Queue	A Queue is a container of messages in the MQ Infrastructure	Message Queue Resource	Diagnostics	Diagnostics
MS Cluster	Microsoft cluster server provides a clustering technology that keeps server-based applications available, regardless of individual component failures	FailoverCluster	OM	OM
MSSQL Database	SQL Data Base from Microsoft	Database Schema	Diagnostics	Diagnostics
Net Device	The Net Device class represents a specific purpose machines such as Routers, Switches, Printers	Node	OM	OM
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, NNMi, TV	RUM, SiS, OM, TV

Class Name	Description	Superclass	Producers	Consumers
NTCMD	xcmd or pstools	Shell	SiS	SiS
Oracle	Oracle database	Database	SiS, OM, Diagnostics	SiS, OM, Diagnostics
Oracle iAS	The Oracle IAS Application Server	J2EE Server	Diagnostics	Diagnostics
Oracle RAC	Oracle RAC allows multiple computers to run the Oracle RDBMS software simultaneously while accessing a single database	Load Balancing Cluster	OM	OM
PrimaryDomainController Master	The PDC emulator is a domain controller that advertises itself as the primary domain controller (PDC) to workstations, member servers, and domain controllers that are running earlier versions of Windows.	DomainControllerRole	OM	OM
Process	An instance of a program.	NodeElement	Diagnostics	Diagnostics
Registrar Server	A registrar is a service that accepts register requests from Session Initiation Protocol (SIP) endpoints. Registrars then save the registration information into a location database.		OM	OM
RelativeIDMaster	The RID master is responsible for processing RID pool requests from all domain controllers in a particular domain.	DomainControllerRole	OM	OM
Resource Pool	A generic Resource Pool	Application Resource	OM	OM
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, TV	RUM, SiS, TV
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
SchemaMaster	The schema master domain controller controls all updates and modifications to the schema.	DomainControllerRole	OM	OM
Service Guard Cluster	A Service Guard Cluster	FailoverCluster	OM	OM
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

Class Name	Description	Superclass	Producers	Consumers
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
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, OM, Diagnostics	SiS, OM, Diagnostics
SSH	ssh agent	Shell	SiS	SiS
Telnet	telnet	Shell	SiS	SiS
Unix	Host with Unix Operation system.	Computer	SiS, OM, Diagnostics, TV	SiS, OM, Diagnostics, TV
Veritas Cluster	Veritas cluster server provides a clustering technology that keeps server-based applications available, regardless of individual component failures	FailoverCluster	OM	OM
Virtual Management Software	Represents a Virtual Management Software	RunningSoftware	OM	OM
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, OM, Diagnostics	SiS, OM
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, OM	SiS, OM
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, OM, Diagnostics	SiS, OM
VMware Resource Pool	VMware Resource Pool represents a pool where Virtual machines execute in, and draw their resources from	Resource Pool	SiS, OM	SiS, OM
Web Service	Represent a web service	WebService Resource	SiS, Diagnostics	SiS, Diagnostics
Web Service Operation	Web Service Operation	WebService Resource	SiS, Diagnostics	SiS, Diagnostics

Class Name	Description	Superclass	Producers	Consumers
Weblogic AS	The Weblogic Application Server	J2EE Server	SiS, OM, Diagnostics, TV	SiS, OM, Diagnostics, TV
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, OM, Diagnostics, TV	SiS, OM, Diagnostics, TV
Windows	Host with Microsoft Operation system (NT).	Computer	SiS, OM, Diagnostics, TV	SiS, OM, Diagnostics, TV

We appreciate your 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 (Business Service Management 9.23)

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 Sw-doc@hp.com.