

HP Database and Middleware Automation

For Linux, AIX, and Solaris

Software Version: 10.20

Sybase ASE Database Refresh User Guide

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 2012-2013 Hewlett-Packard Development Company, L.P.

Trademark Notices

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

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

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:

<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.

The following table indicates changes made to this document since the last released major edition.

Document Changes

Chapter	Version	Changes
Title Page Legal Notices	10.01	Updated version number, software release date, document release date, and copyright date range.
Import the Solution Pack	10.01	Updated the HP DMA web user interface URL to reference the correct default port (8443) for SSL communication.
Title Page Legal Notices	10.10	Updated version number, software release date, document release date, and copyright date range.
Title Page Legal Notices	10.20	Updated version number, software release date, document release date, and copyright date range.
Quick Start Tutorial Workflow Details	10.20	Removed the Quick Start chapter. In the "How to Run this Workflow" sections, pointed to the <i>HP DMA Quick Start Tutorial</i> .

Support

Visit the HP Software Support Online web site at:

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

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

Contents

Contents	5
About HP DMA Solution Packs	7
Audience	8
Document Map	9
Important Terms	10
The Sybase ASE Database Refresh Solution	11
Supported Products and Platforms	12
Prerequisites for this Solution	13
Workflow Details	15
Dump Sybase Database	16
Prerequisites for this Workflow	17
How this Workflow Works	19
How to Run this Workflow	22
Sample Scenarios	24
Parameters for Dump Sybase Database	28
Load Sybase Database Dump	31
Prerequisites for this Workflow	32
How this Workflow Works	34
How to Run this Workflow	37
Sample Scenarios	39
Parameters for Load Sybase Database Dump	43
Dump And Load Sybase Database	46
Prerequisites for this Workflow	48
How this Workflow Works	50
How to Run this Workflow	53
Sample Scenarios	56
Parameters for Dump and Load Sybase Database	60

Reference Information	64
Sybase ASE Product Documentation	64
HP DMA Documentation	64
Tips and Best Practices	65
How this Solution is Organized	66
How to Expose Additional Workflow Parameters	71
How to Use a Policy to Specify Parameter Values	72
Create a Policy	72
Extract a Policy	73
Reference the Policy in the Deployment	73
How to Import a File into the Software Repository	75
Troubleshooting	76
Target Type	76
User Permissions and Related Requirements	76
Discovery in HP DMA	77
Glossary	78

About HP DMA Solution Packs

HP Database and Middleware Automation (HP DMA) software automates administrative tasks like provisioning and configuration, compliance, patching, and release management for databases and application servers. When performed manually, these day-to-day operations are error-prone, time consuming, and difficult to scale.

HP DMA automates these daily, mundane, and repetitive administration tasks that take up 60-70% of a database or application server administrator's day. Automating these tasks enables greater efficiency and faster change delivery with higher quality and better predictability.

HP DMA provides role-based access to automation content. This enables you to better utilize resources at every level:

- End-users can deliver routine, yet complex, DBA and middleware tasks.
- Operators can execute expert level tasks across multiple servers including provisioning, patching, configuration, and compliance checking.
- Subject matter experts can define, enforce, and audit full stack automation across network, storage, server, database, and middleware.

An HP DMA workflow performs a specific automated task—such as provisioning database or application servers, patching database or application servers, or checking a database or application server for compliance with a specific standard. You specify environment-specific information that the workflow requires by configuring its parameters.

Related HP DMA workflows are grouped together in solution packs. When you purchase or upgrade HP DMA content, you are granted access to download specific solution packs.

Audience

This solution is designed for people who are responsible for the following types of tasks:

- Performing routine database backup and restore operations
- Restoring a database from a database dump file
- Copying production database data into a Dev/Test/Staging database environment

Minimal knowledge of Sybase Adaptive Server Enterprise (Sybase ASE) is required to run the workflows in this solution pack using the default settings.

To customize this solution, however, you should be familiar with Sybase ASE database administration, migration, and backup/restore. You should also have hands-on experience upgrading or downgrading a large database (see [Sybase ASE Product Documentation](#) on page 64).

Document Map

The following table shows you how to navigate this guide:

Topic	Description
The Sybase ASE Database Refresh Solution	General information about this solution, including what it contains and what it does.
Workflow Details	Information about the Sybase ASE database refresh workflows included in this solution, including: prerequisites, how the workflows work, how to run them, sample scenarios, and a list of input parameters.
Reference Information	Links to current Sybase ASE product documentation and additional HP DMA documentation.
Tips and Best Practices	Simple procedures that you can use to accomplish a variety of common HP DMA tasks.
Troubleshooting	Tips for solving common problems.

Important Terms

Here are a few basic HP DMA terms that you will need to know:

- In HP DMA, a **workflow** executes a process —such as installing a software product or checking a database instance for compliance with a specific security benchmark.
- A workflow consists of a sequence of **steps**. Each step performs a very specific task. Steps can be shared among workflows.
- Steps can have input and output **parameters**, whose values will be unique to your environment.

If you provide correct values for the input parameters that each scenario requires, the workflow will be able to accomplish its objective. Output parameters from one step often serve as input parameters to another step.

- A **solution pack** contains a collection of related workflows and the steps, functions, and policies that implement each workflow.

More precisely, solution packs contain **workflow templates**. These are read-only versions of the workflows that cannot be deployed. To run a workflow included in a solution pack, you must first create a deployable copy of the workflow template and then customize that copy for your environment.

- A **deployment** associates a workflow with the targets (servers, instances, or databases) where the workflow will run. To run a workflow, you execute a specific deployment. A deployment is associated with one workflow; a workflow can have many deployments, each with its own targets and parameter settings.
- The umbrella term **automation items** is used to refer to those items to which role-based permissions can be assigned. Automation items include workflows, deployments, steps, and policies.

Organizations also have role-based permissions. Servers, instances, and databases inherit their role-based permissions from the organization in which the server resides.

- The **software repository** contains any files that a workflow might need to carry out its purpose (for example, software binaries or patch archives). If the files that a workflow requires are not in the software repository, they must be stored locally on each target server.

When you are using HP DMA with HP Server Automation (HP SA), the software repository is the HP SA Software Library.

- An **organization** is a logical grouping of servers. You can use organizations to separate development, staging, and production resources—or to separate logical business units. Because user security for running workflows is defined at the organization level, organizations should be composed with user security in mind.

Additional terms are defined in the [Glossary](#) on page 78.

Chapter 1

The Sybase ASE Database Refresh Solution

The HP HP DMA Database Refresh solution contains three Sybase ASE workflows:

Workflow Name	Purpose
Dump Sybase Database	Dumps the contents of a Sybase ASE database into a file.
Load Sybase Database Dump	Loads a specific database in the target Sybase ASE instance from an existing dump file that you provide.
Dump And Load Sybase Database	Dumps the contents of a Sybase ASE database (the source) into a file (the database dump file), and loads the contents of that file into an existing Sybase ASE database (the destination).

Although minimal Sybase ASE knowledge is required to run these workflows using the default settings, the workflows are highly customizable and can support complex environment-specific deployment scenarios.

The remaining topics in this chapter provide the following contextual information about these workflows:

- [Supported Products and Platforms](#) on the next page
- [Prerequisites for this Solution](#) on page 13

Supported Products and Platforms

The Sybase ASE Database Refresh workflows support the following database refresh scenarios on Linux and Solaris platforms:

Dump File Version	Target Instance Version
Sybase ASE15.0.3 or 15.5	Sybase ASE15.0.3 or 15.5
Sybase ASE12.5.4	Sybase ASE15.0.3

Operating Systems

For specific target operating system versions supported by each workflow, see the *HP Database and Middleware Automation Support Matrix* available on the HP Software Product Manuals web site:

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

Hardware Requirements

For HP DMA server hardware requirements, see the *HP DMA Installation Guide* and the *HP DMA Release Notes*.

For Sybase ASE hardware and software requirements, see the [Sybase ASE Product Documentation](#) on page 64

HP Software Requirements

This solution requires HP DMA version 10.20 (or later).

Prerequisites for this Solution

The following prerequisites must be satisfied before you can run the Sybase ASE workflows in this solution pack:

1. The source database must exist before the [Dump Sybase Database](#) workflow or the [Dump And Load Sybase Database](#) workflow runs.
2. The destination database must exist before the [Load Sybase Database Dump](#) workflow or the [Dump And Load Sybase Database](#) workflow runs.
3. The source database cannot be mounted on the master device.
4. The destination database cannot be mounted on the master device.
5. Both the Adaptive Server instance and the local Backup Server instance must be running, and they must be able to communicate with each other.
6. The master database system table (syssservers) must contain an entry that assigns the local Backup Server instance to SYB_BACKUP.
7. The user specified in the Sybase OS User Name parameter (sybase by default) must own the Sybase ASE installation directory and be a member of the "Sybase" group.
8. You must specify an operating system file in the Dump File List parameter (for example: `/var/temp/mydbdump.dmp`). You cannot specify a dump device.

The database dump file must be accessible from the server where the workflow is executed. The file must be available on the local machine or via a Network File System (NFS) mount.

The workflows currently do not support writing or reading the database dump file from tape devices.

9. The Dump File Password parameter is required when the database dump file is encrypted with a password.
10. Adequate disk space must be available on the pertinent database servers.
11. On Linux and Solaris platforms, the `sudo` package must be installed on the target servers.

Additional Considerations

- If the dump file is password-protected, you cannot perform a cross-platform refresh.
- If database transactions occur on the source database after you create the dump file, you should apply the latest transaction log dump to the target database after you run the workflow. Otherwise, these transactions will be missing from the target database.

It is good practice to run basic database consistency checks (DBCCs) on the source database before running this workflow. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If database transactions occur on the source database after the dump file is created, you should apply the latest transaction log dump to the destination database after you run the [Load Sybase Database Dump](#) workflow. Otherwise, these transactions will be missing from the destination database.

- Be sure to set the appropriate database level configuration parameters (for example: trunc log on chkpt, abort tran on log full, no chkpt on recovery) prior to running these workflows. The workflows do not modify these settings.

Caution: Performing frequent database dump or load operations can hamper system performance and impede target database usage.

Note: There are also specific prerequisites for each workflow.

Chapter 2

Workflow Details

Each workflow included in this solution pack has a set of input parameters whose values will be unique to your environment. If you provide correct values for the parameters that each scenario requires, the workflow will be able to accomplish its objective.

Tip: Input parameters are described in the "Parameters" topic for each workflow.

There are two steps required to customize this solution:

1. Ensure that all required parameters are visible. You do this by using the workflow editor.
To perform a simple database refresh, you can use the default values for most parameters. To use more advanced features of this solution, you will need to expose additional parameters.
2. Specify the values for those parameters. You do this when you create a deployment.

Tip: Each of these steps is explained in greater detail in the "How to Use this Workflow" topic for each workflow.

The information presented here assumes the following:

- HP DMA is installed and operational.
- At least one suitable target server is available (see [Supported Products and Platforms](#) on page 12).
- You are logged in to the HP DMA web interface.
- You have permission to create, edit, and deploy copies of the workflows included in this solution pack.

Dump Sybase Database

This workflow enables you to dump the contents of a Sybase Adaptive Server Enterprise (ASE) database (the source database) into a file (the database dump file).

The workflow performs extensive validation checks prior to and immediately after the dump operation to ensure that the dump file is valid.

This workflow can create and load database dump files that are striped, compressed (at any level 1-9), encrypted, or any combination thereof.

If any source database objects are bound to a specific (non-default) cache, the workflow will create a cache dump file—provided that you specify a valid value for the Cache Dump File parameter. The cache dump file contains details about the specific caches used by the source database and any objects that are bound to each cache. This file is in data-readable format.

To use this workflow in your environment, see the following information:

Topic	Information Included
Prerequisites for this Workflow	List of prerequisites that must be satisfied before you can run this workflow
How this Workflow Works	Information about what the workflow does, including validation checks performed, steps executed, and a high-level process flow
How to Run this Workflow	Instructions for running this workflow in your environment
Sample Scenarios	Examples of typical parameter values for this workflow
Parameters for Dump Sybase Database	List of input parameters for this workflow

The process of deploying and running this workflow is the same for all scenarios, but the parameters required will differ depending on the specific scenario that you are implementing.

The workflow provides default values for most parameters. These default values are usually sufficient for a "typical" database refresh. You can override the defaults by specifying parameter values in the deployment. You can also expose additional parameters in the workflow, if necessary, to accomplish more advanced scenarios.

Any parameters not explicitly specified in the deployment will have the default values listed in [Parameters for Dump Sybase Database](#) on page 28.

Note: To view detailed information about the steps included in this workflow, see the [Steps in this Workflow](#).

Prerequisites for this Workflow

Caution: You cannot perform a database dump if there are dependencies between the source database and another database in the same or a different Sybase ASE instance. You must remove or disable object referencing (for example: triggers, views, stored procedures, etc.) before you run this workflow.

The following prerequisites must be satisfied before you can run the Dump Sybase Database workflow:

1. The source database must exist before the workflow runs.
2. The source database must NOT be mounted on the master device.
3. Both the Adaptive Server instance that executes the `dump` command and the local Backup Server instance must be running, and they must be able to communicate with each other.
4. The master database system table (`syssservers`) must contain an entry that assigns the local Backup Server instance to `SYB_BACKUP`.
5. By default, the workflow will create the database dump file with the following format:

```
dump_file_path/<databaseName>_<dateTime>.dmp
```

For example: `/var/tmp/mytestdb_2012111283762.dmp`

If you specify a non-default file name (or names) in the Dump File List parameter, the path to each specified file must exist.

6. Adequate disk space must be available to store the database dump file, whether it is stored locally or in a shared NFS location.
7. On Linux and Solaris platforms, the `sudo` package must be installed on the server that hosts the source database.
8. The Dump File Password parameter is required if a password was used to encrypt the source database dump file.
9. The workflow assumes the following:
 - The Adaptive Server component is installed in the `/home/Sybase/ASE_15` directory.
 - The Adaptive Server instance name is `NY_DS`.
 - The database name is `mytestdb`.
 - The database dump file is stored in the `/var/tmp` directory.
 - `/var/tmp` is an NFS mount point.
 - The Sybase ASE user specified in the ASE SysAdmin Username parameter is permitted to access the `/var/tmp` directory.
 - The user specified in the Sybase OS User Name parameter (sybase by default) must own the installation directory and be a member of the “sybase” group.

Note: The workflow currently does not support reading the database dump file from tape devices.

Note: This workflow does not support dump file password encryption for cross-platform database refresh (for example: the database dump file was created on a Linux server, and you are loading it onto a Solaris server).

Additional Considerations

It is good practice to run basic database consistency checks (DBCCs) on the source database before running this workflow. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If database transactions occur on the source database after the dump file is created, you should apply the latest transaction log dump to the destination database after you run the [Load Sybase Database Dump](#) workflow. Otherwise, these transactions will be missing from the destination database.

For information about prerequisites for Sybase ASE, refer to the [Sybase ASE Product Documentation](#) on page 64.

How this Workflow Works

This topic contains information about the [Dump Sybase Database](#) workflow:

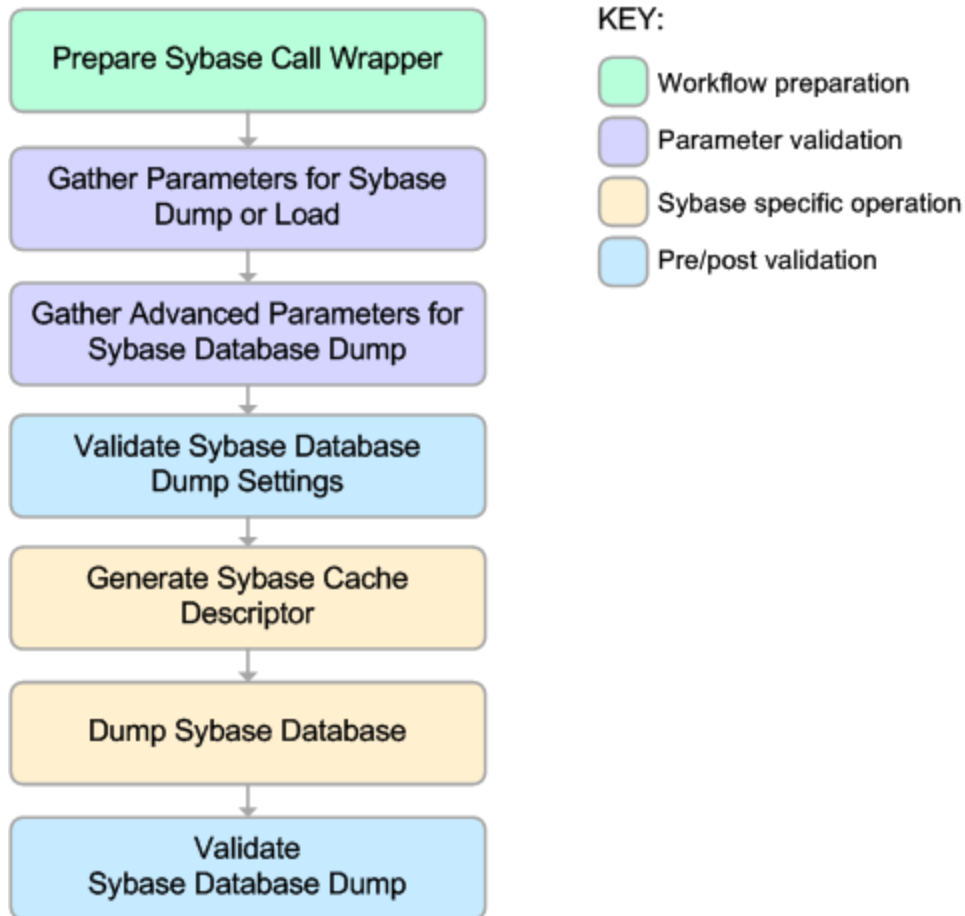
Validation Checks Performed

The workflow checks the following things prior to dumping the database. If any of these checks fails, the workflow fails.

1. All required parameters have values. If any required parameter does not have a value—either a value that you specify or a default value—the workflow fails in the Sybase - Validate Database Dump Settings step.
2. The Sybase ASE software is installed.
3. The source database exists in the specified Sybase ASE instance.
4. The source Adaptive Server and Backup Server components are running and able to communicate with each other so that they can perform the database dump.
5. The source database is online.
6. Adequate disk space is available to store the database dump file.

Steps Executed

The [Dump Sybase Database](#) workflow includes the following steps. Each step must complete successfully before the next step can start. If a step fails, the workflow reports a failure, and all subsequent steps are skipped.



Process Flow

This workflow performs the following tasks:

1. Creates the Instance Wrapper and Server Wrapper.
2. Verifies that the HP DMAserver is able to communicate with the server where the workflow is running.
3. Performs the preliminary [validation checks](#) described above.
4. Generates the cache descriptor file for the source database. This is used to replicate the cache objects on the destination server.
5. Performs the database `dump` operation to create the database dump file.
6. Performs post-dump validation checks to ensure that all required parameters had valid values.

Tips and Best Practices

It is good practice to run basic database consistency checks (DBCCs) on the source database before running this workflow to ensure that there are no internal errors in the database. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If you find errors in the source database, be sure to fix them before running this workflow. The workflow does not have the ability to diagnose or remediate problems in the database prior to performing the database dump.

How to Run this Workflow

This topic explains how to customize and run the [Dump Sybase Database](#) workflow in your environment.

Note: Prior to running this workflow, review the [Prerequisites for this Workflow](#), and ensure that all requirements are satisfied.

Tip: For detailed instructions to run HP DMA workflows—using the Run Oracle Compliance Audit workflow as an example—see *HP DMA Quick Start Tutorial*.

To customize and run the Dump Sybase Database workflow:

1. Create a deployable copy of the workflow (see "Create a Deployable Workflow" in *HP DMA Quick Start Tutorial*).
2. Determine the values that you will specify for the following parameters. This is the minimum set of parameters required to run this workflow.

Parameter Name	Default Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Cache File	no default	Database cache file associated with this database dump. This is a single filename (with absolute path—path must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.
Dump File Compression Level	7	Compression level (1-9) to apply to the dump file (or files) that will be created.
Dump File Password	no default	Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted). Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).

Note: This is the minimum set of parameters required to run this workflow. You may need to expose additional parameters depending on your provisioning objectives.

See [Parameters for Dump Sybase Database](#) on page 28 for detailed descriptions of all input parameters for this workflow, including default values.

Tip: To avoid having to re-enter passwords whenever they change, you can create a policy to provide them to the workflow (see [How to Use a Policy to Specify Parameter Values](#) on page 72).

3. In the workflow editor, expose any additional parameters that you need (see [How to Expose Additional Workflow Parameters](#) on page 71). You will specify values for those parameters when you create the deployment.
4. Save the changes to the workflow (click **Save** in the lower right corner).
5. Create a new deployment (see "Create a Deployment" in *HP DMA Quick Start Tutorial* for instructions).
6. On the Parameters tab, specify values for the required parameters listed in step 2 and any additional parameters that you have exposed. You do not need to specify values for those parameters whose default values are appropriate for your environment.
7. On the Targets tab, specify one or more targets for this deployment.
8. Save the deployment (click **Save** in the lower right corner).
9. Run the workflow using this deployment (see "Run Your Workflow" in *HP DMA Quick Start Tutorial* for instructions).

The workflow will complete and report "Success" on the Console if it has run successfully. If an invalid parameter value is specified, an error is logged, and the workflow terminates in the "Failure" state.

Sample Scenarios

This topic shows you how to use various parameters to achieve the following database refresh scenarios in your environment using the [Dump Sybase Database](#) workflow:

Scenario 1: Perform a Database Refresh Using a Database Dump File that is Not Encrypted or Striped

This is the simplest Sybase ASE database dump scenario.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Dump Sybase Database](#) on page 28).

Scenario 2: Perform a Database Refresh Using a Database Dump File that is Encrypted and Compressed

This scenario requires you to specify the encryption password and compression level for the database dump file.

Parameter Name	Example Value	
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File Compression Level	8	Compression level (1-9) to apply to the dump file (or files) that will be created.
Dump File Password	MyPassword1@#	<p>Password that will be used to encrypt the database dump file.</p> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p>

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Dump Sybase Database](#) on page 28).

Scenario 3: Perform a Database Refresh Using a Database Dump File that is Striped and Encrypted

In this scenario, the database dump file will be striped across multiple files. You must specify all the individual stripe files in the Dump File List parameter (separate them with commas).

If you want the stripe files to be encrypted, you must also specify the Dump File Password parameter.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	MyPassword1@#	Password that will be used to encrypt the database dump file. Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Dump Sybase Database](#) on page 28).

Scenario 4: Perform a Database Refresh Using a Cache Dump File

In this scenario, the database dump file has an associated cache dump file. You must specify the name of the cache dump file by using the Cache File parameter. The workflow will use the cache dump file to rebuild and bind the cache after the database dump file is loaded into the destination database.

Parameter Name	Example Value	
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Cache File	<code>/var/tmp/runcache_mytestdb.txt</code>	File where the Sybase ASE database cache configuration data for the source database will be written. This is a single filename (with absolute path—path must exist).

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Parameters for Dump Sybase Database

The following tables describe the required and optional input parameters for this workflow. Most of these parameters are not initially visible in a deployment (see [How to Expose Additional Workflow Parameters](#) on page 71). For most parameters, if you do not specify a value for a parameter, a default value is assigned in the following steps:

- Gather Parameters for Sybase Dump or Load
- Gather Advanced Parameters for Sybase Database Dump

Input Parameters for the Dump Sybase Database Workflow

Parameter Name	Default Value	Required	Description
ASE SysAdmin Password	password	required	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
ASE SysAdmin Username	sa	required	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database load operation.
Cache File	no default	optional	Database cache file associated with this database dump. This is a single filename (with absolute path—the path and file must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.
Call Wrapper	jython	required	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Dump Device Name	n/a	n/a	Not used in this release.
Dump File Compression Level	7	optional	Compression level (1-9) to apply to the dump file (or files) that will be created.

Input Parameters for the Dump Sybase Database Workflow (continued)

Parameter Name	Default Value	Required	Description
Dump File List	<p><code>/var/tmp/<dbname>_<datetime>.dmp</code></p> <p>For example: <code>/var/tmp/mytestdb_2012111283762.dmp</code></p> <p>If you specify a non-default file name (or names) in the Dump File List parameter, the path to each specified file must exist.</p>	required	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	no default	optional	<p>Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted).</p> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p>
Dump File Path	n/a	n/a	Not used in this release.
Local Backup Instance Name	n/a	n/a	Not used in this release.
Local Backup Instance Port	n/a	n/a	Not used in this release.
Remote Backup Instance Name	n/a	n/a	Not used in this release.

Input Parameters for the Dump Sybase Database Workflow (continued)

Parameter Name	Default Value	Required	Description
Remote Backup Instance Port	n/a	n/a	Not used in this release.
Role Password SQL Statement	no default	required	Not used in this release.
Source Database Instance Name	NY_DS	required	Name of the Adaptive Server instance where the dump file (or files) will be created. You specify the value of this parameter in the deployment.
Source Database Name	mytestdb	required	Name of database from which the dump file (or files) will be created.
Sybase ASE Home Directory	/home/sybase/ASE_15	required	Sybase ASE installation home directory, where the source database resides. Sybase will examine the interface file that exists in this directory to determine how to create the specified database dump file (or files).
Sybase OS User Name	sybase	required	OS user (typically, sybase) who owns the Sybase ASE installation directory.

Load Sybase Database Dump

This workflow enables you to load the contents of a previously created Sybase ASE database dump file (the source data) into an existing Sybase ASE database (the destination database).

The workflow performs extensive validation checks prior to and immediately after loading the source data into the destination database to ensure that the schema and data have been loaded successfully. The workflow restores any existing database users after the source data is loaded into the destination database.

This workflow can perform a cross-platform database refresh (load) if necessary. After it performs a cross-platform load operation, the workflow rebuilds the indexes (clustered or non-clustered indexes on APL/DOL tables) to avoid any page linkage or index corruption issues.

The source database dump file (or files) can be striped, compressed (at any level 1-9), encrypted, or any combination thereof.

To use this workflow in your environment, see the following information:

Topic	Information Included
Prerequisites for this Workflow	List of prerequisites that must be satisfied before you can run this workflow
How this Workflow Works	Information about what the workflow does, including validation checks performed, steps executed, and a high-level process flow
How to Run this Workflow	Instructions for running this workflow in your environment
Sample Scenarios	Examples of typical parameter values for this workflow
Parameters for Load Sybase Database Dump	List of input parameters for this workflow

The process of deploying and running this workflow is the same for all scenarios, but the parameters required will differ depending on the specific scenario that you are implementing.

The workflow provides default values for most parameters. These default values are usually sufficient for a "typical" database refresh. You can override the defaults by specifying parameter values in the deployment. You can also expose additional parameters in the workflow, if necessary, to accomplish more advanced scenarios.

Any parameters not explicitly specified in the deployment will have the default values listed in [Parameters for Load Sybase Database Dump](#) on page 43.

Note: To view detailed information about the steps included in this workflow, see [Steps in this Workflow](#).

Prerequisites for this Workflow

The following prerequisites must be satisfied before you can run the Load Sybase Database Dump workflow:

1. The destination database must exist before the workflow runs.
2. The destination database must NOT be mounted on the master device.
3. The source and destination database servers must use the same page size.
4. No database users may be logged in to the destination database server when this workflow runs.
5. Both the Adaptive Server instance that executes the `load` command and the local Backup Server instance must be running, and they must be able to communicate with each other.
6. The master database system table (`syssservers`) must contain an entry that assigns the local Backup Server instance to `SYB_BACKUP`.
7. On Linux and Solaris platforms, the `sudo` package must be installed on the server that hosts the destination database.
8. You must specify an operating system file in the Dump File List parameter (for example: `/var/temp/mydbdump.dmp`). You cannot specify a dump device.

The database dump file must be accessible from the server where the workflow is executed. The file must be available on the local machine or via a Network File System (NFS) mount.

The workflows currently do not support writing or reading the database dump file from tape devices.

9. The file (or files) specified in the Dump File List parameter must exist in the specified location.
10. The Dump File Password parameter is required if a password was used to encrypt the source database dump file.
11. The workflow assumes the following:
 - The Adaptive Server component is installed in the `/home/sybase/ASE_15` directory.
 - The Adaptive Server instance name is `NY_DS`.
 - The database name is `mytestdb`.
 - The database dump file is stored in the `/var/tmp` directory.
 - `/var/tmp` is an NFS mount point.
 - The Sybase ASE user specified in the ASE SysAdmin Username parameter is permitted to access the `/var/tmp` directory.
 - The user specified in the Sybase OS User Name parameter (sybase by default) must own the installation directory and be a member of the “sybase” group.

Note: The workflow currently does not support reading the database dump file from tape devices.

Note: This workflow does not support dump file password encryption for cross-platform database refresh (for example: the database dump file was created on a Linux server, and you are loading it onto a Solaris server).

Additional Considerations

It is good practice to run basic database consistency checks (DBCCs) on the source database before running this workflow. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If database transactions occur on the source database after the dump file is created, you should apply the latest transaction log dump to the destination database after you run the [Load Sybase Database Dump](#) workflow. Otherwise, these transactions will be missing from the destination database.

For information about prerequisites for Sybase ASE, refer to the [Sybase ASE Product Documentation](#) on page 64.

How this Workflow Works

This topic contains information about the [Load Sybase Database Dump](#) workflow:

Caution: You cannot refresh the target database (load the database dump) if there are dependencies between the target database and another database in the same or a different Sybase ASE instance. You must remove or disable object referencing (for example: triggers, views, stored procedures, etc.) before you run this workflow.

Validation Checks Performed

The workflow checks the following things prior to refreshing the database. If any of these checks fails, the workflow fails.

1. All required parameters have values. If any required parameter does not have a value—either a value that you specify or a default value—the workflow fails in the Sybase - Validate Database Refresh Settings step.
2. The Sybase ASE software is installed.
3. The target database and the Backup Server are running and able to communicate with each other.
4. The dump file server page size matches the target database server page size.
5. The Sybase ASE version in the database dump file header is compatible with the target Sybase ASE instance version. The following versions are compatible:

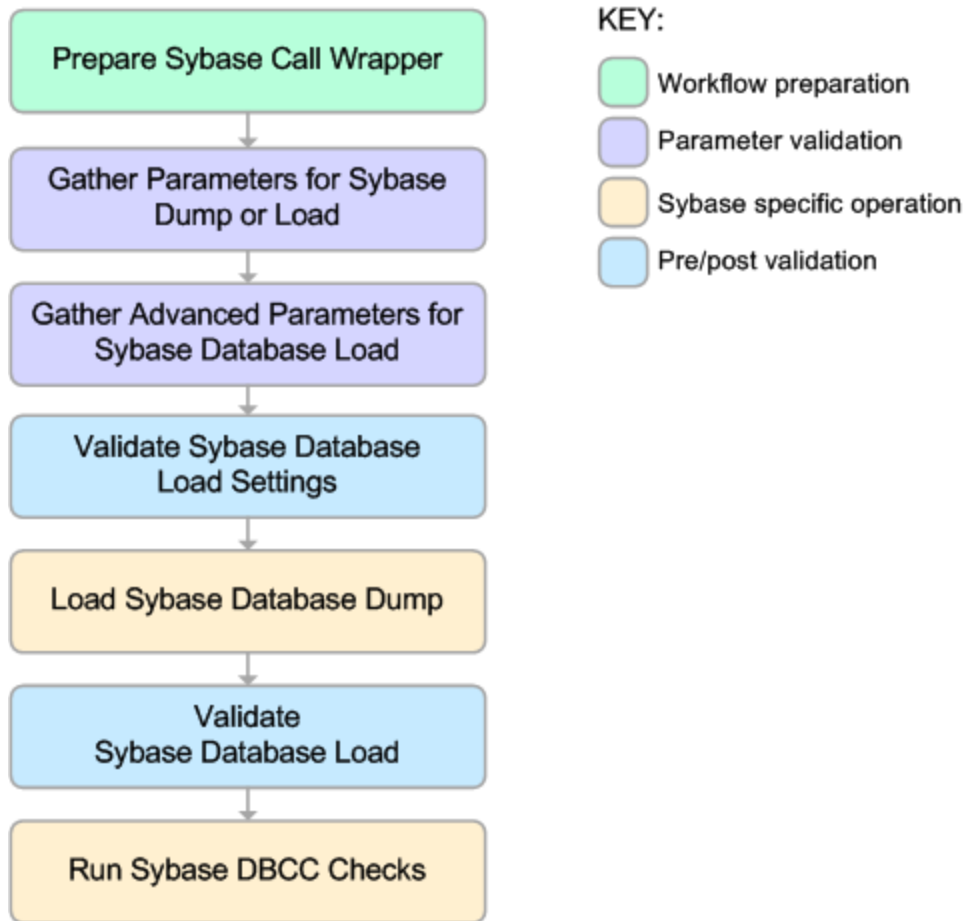
Dump File Version	Target Instance Version
Sybase ASE15.0.3 or 15.5	Sybase ASE15.0.3 or 15.5
Sybase ASE12.5.4	Sybase ASE15.0.3

6. The target database exists.
7. The size of the target database is sufficient to load the database dump file.

The workflow then determines whether the target database server is currently in use by Sybase ASE users. If the database is in use, the workflow creates a backup of the target database users and groups in `tempdb` before it refreshes the database. It restores the users after the database is refreshed.

Steps Executed

The [Load Sybase Database Dump](#) workflow includes the following steps. Each step must complete successfully before the next step can start. If a step fails, the workflow reports a failure, and all subsequent steps are skipped.



Process Flow

This workflow performs the following tasks:

1. Creates the Instance Wrapper and Server Wrapper.
2. Verifies that the HP DMAserver is able to communicate with the server where the workflow is running.
3. Performs the pre-refresh [validation checks](#) described above.
4. Reads the header of the database dump file, and validates that the source Sybase ASE page size matches the target Sybase ASE page size. The workflow fails if the page sizes do not match.
5. Backs up any existing database users and groups.
6. Determines whether the source (the database dump file) and target servers have different byte architectures (big-endian versus little-endian).
7. Enables the database level “`dbo use only`” option to ensure that no users except the database owner are accessing the destination database.
8. Loads the database dump file on the target database server, and refreshes the destination database.
9. Brings the destination database online and performs the post-refresh checks.
10. Rebuild the indexes using the Sybase ASE recommended `sp_post_xpload` system stored procedure.
11. Runs the specified database consistency checker (DBCC) checks to ensure that no database tables or objects have become corrupted. The output of these checks is printed in the step log and stored in files under the specified directory.
12. Builds the specified cache (if specified), and binds the database object to either the default data cache or the specified cache.

Tips and Best Practices

It is good practice to run basic database consistency checks (DBCCs) on the source database before you create the dump file (or files) to ensure that there are no internal errors in the database. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If you find errors in the source database, be sure to fix them before you create the dump file. The workflow does not have the ability to diagnose or remediate problems in the database.

How to Run this Workflow

This topic explains how to customize and run the [Load Sybase Database Dump](#) workflow in your environment.

Note: Prior to running this workflow, review the [Prerequisites for this Workflow](#), and ensure that all requirements are satisfied.

Tip: For detailed instructions to run HP DMA workflows—using the Run Oracle Compliance Audit workflow as an example—see *HP DMA Quick Start Tutorial*.

To customize and run the Load Sybase Database Dump workflow:

1. Create a deployable copy of the workflow (see "Create a Deployable Workflow" in *HP DMA Quick Start Tutorial*).
2. Determine the values that you will specify for the following parameters. This is the minimum set of parameters required to run this workflow.

Parameter Name	Default Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/dump.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	no default	Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted). Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).

Note: This is the minimum set of parameters required to run this workflow. You may need to expose additional parameters depending on your provisioning objectives.

See [Parameters for Load Sybase Database Dump](#) on page 43 for detailed descriptions of all input parameters for this workflow, including default values.

Tip: To avoid having to re-enter passwords whenever they change, you can create a policy to provide them to the workflow (see [How to Use a Policy to Specify Parameter Values](#) on page 72).

3. In the workflow editor, expose any additional parameters that you need (see [How to Expose Additional Workflow Parameters](#) on page 71). You will specify values for those parameters when you create the deployment.
4. Save the changes to the workflow (click **Save** in the lower right corner).
5. Create a new deployment (see "Create a Deployment" in *HP DMA Quick Start Tutorial* for instructions).
6. On the Parameters tab, specify values for the required parameters listed in step 2 and any additional parameters that you have exposed. You do not need to specify values for those parameters whose default values are appropriate for your environment.
7. On the Targets tab, specify one or more targets for this deployment.
8. Save the deployment (click **Save** in the lower right corner).
9. Run the workflow using this deployment (see "Run Your Workflow" in *HP DMA Quick Start Tutorial* for instructions).

The workflow will complete and report "Success" on the Console if it has run successfully. If an invalid parameter value is specified, an error is logged, and the workflow terminates in the "Failure" state.

Sample Scenarios

This topic shows you how to use various parameters to achieve the following database refresh scenarios in your environment using the [Load Sybase Database Dump](#) workflow:

Scenario 1: Database Dump File is Not Encrypted or Striped

This is the simplest Sybase ASE database refresh scenario. It does not matter whether the database dump file is compressed – if decompression is required, it is handled automatically by the workflow prior to the refresh.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Scenario 2: Database Dump File is Encrypted

This scenario requires you to specify the encryption password for the database dump file. It does not matter whether the database dump file is compressed – if decompression is required, it is handled automatically by the workflow prior to the refresh.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	/var/tmp/mytestdb.dmp	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	MyPassword1@#	<p>Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted).</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p> </div>

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Scenario 3: Database Dump File is Striped

In this scenario, the database dump file has been striped across multiple files. You must specify all the individual stripe files in the Dump File List parameter (separate them with commas).

If the stripe files are encrypted, you must specify the Dump File Password parameter.

It does not matter whether the database dump file is compressed – if decompression is required, it is handled automatically by the workflow prior to the refresh.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	MyPassword1@#	Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted). Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Scenario 4: Using a Cache Dump File

In this scenario, the database dump file has an associated cache dump file. You must specify the name of the cache dump file by using the Cache Dump File parameter. The workflow will rebuild and bind the cache after the database dump file is loaded into the target database.

If the database dump file is encrypted, you must specify the Dump File Password parameter.

If the cache dump file is encrypted, you must specify the Cache Dump File Password parameter.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	MyPassword1@#	Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted). Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).
Cache Dump File	<code>/var/tmp/runcache_mytestdb.txt</code>	Database cache file associated with this database dump. This is a single filename (with absolute path—the path and file must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on the next page).

Parameters for Load Sybase Database Dump

The following tables describe the required and optional input parameters for this workflow. Most of these parameters are not initially visible in a deployment (see [How to Expose Additional Workflow Parameters](#) on page 71). For most parameters, if you do not specify a value for a parameter, a default value is assigned in the following steps:

- Gather Parameters for Sybase Dump or Load
- Gather Advanced Parameters for Sybase Database Load

Input Parameters for the Load Sybase Database Dump Workflow

Parameter Name	Default Value	Required	Description
ASE SysAdmin Password	password	required	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
ASE SysAdmin Username	sa	required	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database load operation.
Cache Dump File	no default	optional	Database cache file associated with this database dump. This is a single filename (with absolute path—the path and file must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.
Call Wrapper	jython	required	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
DBCC Checks	checkdb, checkalloc, checkcatalog	optional	List of database consistency checker (DBCC) checks that you want to run to ensure that there are no problems with the database after the dump file is loaded.
DBCC Error Directory	/var/tmp	optional	The directory (with absolute path) where you want to store the DBCC results (output files) for post-load checks. This directory must exist.

Input Parameters for the Load Sybase Database Dump Workflow (continued)

Parameter Name	Default Value	Required	Description
Database Instance Name	NY_DS	required	The name of the Sybase ASE instance where the database will be loaded from the dump file (or files).
Dump Device Name	n/a	n/a	Not used in this release.
Dump File List	/var/tmp/dump.dmp	required	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	no default	optional	<p>Password required to decrypt a password-protected encrypted database dump file (required if the dump file is encrypted).</p> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p>
Dump File Path	n/a	n/a	Not used in this release.
Local Backup Instance Name	n/a	n/a	Not used in this release.
Local Backup Instance Port	n/a	n/a	Not used in this release.
Remote Backup Instance Name	n/a	n/a	Not used in this release.
Remote Backup Instance Port	n/a	n/a	Not used in this release.

Input Parameters for the Load Sybase Database Dump Workflow (continued)

Parameter Name	Default Value	Required	Description
Sybase ASE Home Directory	/home/sybase/ASE_15	required	Sybase ASE installation home directory, where the destination database resides. Sybase ASE will examine the interface file that exists in this directory to determine where to load the specified database dump file (or files).
Sybase OS User Name	sybase	required	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Target Database Instance Name	NY_DS	required	Name of the Sybase ASE instance where the dump file (or files) will be loaded.
Target Database Name	mytestdb	required	Name of the database where the dump file (or files) will be loaded.
Target Database Page Size	4 KB	optional	Page size of the target database server (in kilobytes).

Dump And Load Sybase Database

This workflow enables you to dump the contents of a Sybase ASE database (the source) into a file (the database dump file) and load the contents of that file into an existing Sybase ASE database (the destination).

The workflow performs extensive validation checks prior to and immediately after the dump operation at the source to ensure that the dump file is valid. It also performs validation checks prior to and immediately after the load operation at the destination to ensure that the data was successfully loaded.

This workflow can perform a cross-platform database refresh (dump and load) when necessary. After it performs a cross-platform load operation, the workflow rebuilds the indexes (clustered or non-clustered indexes on APL/DOL tables) to avoid page linkage or index corruption issues. Password protected dump cannot be supported for cross platform dump and load.

If any source database objects are bound to a specific (non-default) cache, the workflow will create a cache dump file—provided that you specify a valid value for the Cache Dump File parameter. The cache dump file contains details about the specific caches used by the source database and any objects that are bound to each cache. This file is in data-readable format. The workflow uses the cache dump file to refresh the destination database cache (provided that ample cache space is available). The workflow cannot, however, configure or enable cache buffering.

This workflow can create and load database dump files that are striped, compressed (at any level 1-9), encrypted, or any combination thereof.

To use this workflow in your environment, see the following information:

Topic	Information Included
Prerequisites for this Workflow	List of prerequisites that must be satisfied before you can run this workflow
How this Workflow Works	Information about what the workflow does, including validation checks performed, steps executed, and a high-level process flow
How to Run this Workflow	Instructions for running this workflow in your environment
Sample Scenarios	Examples of typical parameter values for this workflow
Parameters for Dump and Load Sybase Database	List of input parameters for this workflow

The process of deploying and running this workflow is the same for all scenarios, but the parameters required will differ depending on the specific scenario that you are implementing.

The workflow provides default values for most parameters. These default values are usually sufficient for a "typical" database refresh. You can override the defaults by specifying parameter values in the deployment. You can also expose additional parameters in the workflow, if necessary, to accomplish more advanced scenarios.

Any parameters not explicitly specified in the deployment will have the default values listed in [Parameters for Dump and Load Sybase Database](#) on page 60.

Note: To view detailed information about the steps included in this workflow, see [Steps in this Workflow](#).

Prerequisites for this Workflow

Caution: You cannot perform a database dump if there are dependencies between the source database and another database in the same or a different Sybase ASE instance. You must remove or disable object referencing (for example: triggers, views, stored procedures, etc.) before you run this workflow.

The following prerequisites must be satisfied before you can run the Dump and Load Sybase Database workflow:

1. The source and destination databases must exist before the workflow runs.
2. The source and destination databases must NOT be mounted on the master device.
3. The source and destination database servers must use the same page size.
4. No database users may be logged in to the destination database server when this workflow runs.
5. The source Adaptive Server instance that executes the `dump` command and the local source Backup Server instance must both be running, and they must be able to communicate with each other.
6. The destination Adaptive Server instance that executes the `load` command and the local destination Backup Server instance must both be running, and they must be able to communicate with each other.
7. The master database system table (`sys.servers`) for both source and destination must contain an entry that assigns the local Backup Server instance to `SYB_BACKUP`.
8. By default, the workflow will create the database dump file with the following file name format:

`dump_file_path/databasename_datetime.dmp`

For example: `/var/tmp/mytestdb_2012111283762.dmp`

9. If you specify a non-default file name (or names) in the Dump File List parameter, the path to each specified file must exist.

You must specify an operating system file in the Dump File List parameter (for example: `/var/temp/mydbdump.dmp`). You cannot specify a dump device.

The database dump file must be accessible from the server where the workflow is executed. The file must be available on the local machine or via a Network File System (NFS) mount.

The workflows currently do not support writing or reading the database dump file from tape devices.

10. Adequate disk space must be available on the shared NFS location where the dump file will be stored.
11. On Linux and Solaris platforms, the `sudo` package must be installed on the target servers.
12. The workflow assumes the following for both the source and destination:
 - The Adaptive Server component is installed under `/home/Sybase/ASE_15`
 - The Adaptive Server instance name is `NY_DS`

- The database name is mytestdb
- The database dump file is stored in the `/var/tmp` directory
- `/var/tmp` is an NFS mount point.
- The Sybase ASE user specified in the ASE SysAdmin Username parameter is permitted to access the `/var/tmp` directory.
- The user specified in the Sybase OS User Name parameter (sybase by default) must own the installation directory and be a member of the “sybase” group.

Note: The workflow currently does not support reading the database dump file from tape devices.

Note: This workflow does not support dump file password encryption for cross-platform database refresh (for example: the database dump file was created on a Linux server, and you are loading it onto a Solaris server).

Additional Considerations

It is good practice to run basic database consistency checks (DBCCs) on the source database before running this workflow. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If database transactions occur on the source database after the dump file is created, you should apply the latest transaction log dump to the destination database after you run the [Load Sybase Database Dump](#) workflow. Otherwise, these transactions will be missing from the destination database.

For information about prerequisites for Sybase ASE, refer to the [Sybase ASE Product Documentation](#) on page 64.

How this Workflow Works

This topic contains information about the [Dump And Load Sybase Database](#) workflow:

Caution: You cannot perform a database dump if there are dependencies between the source database and another database in the same or a different Sybase ASE instance. You must remove or disable object referencing (for example: triggers, views, stored procedures, etc.) before you run this workflow.

Validation Checks Performed

The workflow checks the following things prior to refreshing the database. If any of these checks fails, the workflow fails.

1. All required parameters have values. If any required parameter does not have a value—either a value that you specify or a default value—the workflow fails in the Sybase - Validate Database Dump Settings step.
2. The Sybase ASE software is installed.
3. The source database exists in the specified Sybase ASE instance and is online.
4. The source Adaptive Server and Backup Server components are running and able to communicate with each other so that they can perform the database dump.
5. Adequate disk space is available to store the database dump file.
6. The destination database exists in the specified Sybase ASE instance and is online.
7. The destination Adaptive Server and Backup Server components are running and able to communicate with each other.
8. The destination database server page size matches the source database server page size.
9. The Sybase ASE version of the source database is compatible with that of the destination database. The following versions are compatible:


Dump File Version	Target Instance Version
Sybase ASE15.0.3 or 15.5	Sybase ASE15.0.3 or 15.5
Sybase ASE12.5.4	Sybase ASE15.0.3

10. The size of the destination database is sufficient to load the database dump file created from the source database.

The workflow then determines whether the destination database is currently in use by Sybase ASE users. If the database is in use, the workflow creates a backup of the destination database users and groups in `tempdb` before it loads the contents of the dump file. It restores the users after the database is refreshed.

Steps Executed

The [Dump And Load Sybase Database](#) workflow includes the following steps. Each step must complete successfully before the next step can start. If a step fails, the workflow reports a failure, and all subsequent steps are skipped.



KEY:

- Workflow preparation
- Parameter validation
- Sybase specific operation
- Pre/post validation

Process Flow

This workflow performs the following tasks:

1. Creates the Instance Wrapper and Server Wrapper.
2. Verifies that the HP DMA server is able to communicate with the server where the workflow is running.
3. Performs the pre-dump [validation checks](#) described above.
4. Generates the cache descriptor file for the source database. This is used to replicate the cache objects on the destination server.
5. Performs the database `dump` operation to create the database dump file.
6. Performs post-dump validation checks to ensure that all required parameters had valid values.
7. Reads the header of the database dump file, and validates that the source Sybase ASE page size matches the target Sybase ASE page size. The workflow fails if the page sizes do not match.
8. Backs up any existing database users and groups.
9. Determines whether the source and destination database servers have different byte architectures (big-endian versus little-endian).
10. Enables the database level “`dbo use only`” option to ensure that no users except the database owner are accessing the destination database.
11. Loads the database dump file on the destination database server, and refreshes the destination database.
12. Brings the destination database online and performs the post-refresh checks.
13. Rebuild the indexes using the Sybase ASE recommended `sp_post_xpload` system stored procedure.
14. Runs the specified database consistency checker (DBCC) checks to ensure that no database tables or objects have become corrupted. The output of these checks is printed in the step log and stored in files under the specified directory.
15. Builds the specified cache (if specified), and binds the database object to either the default data cache or the specified cache.

Tips and Best Practices

It is good practice to run basic database consistency checks (DBCCs) on the source database before running this workflow to ensure that there are no internal errors in the database. You can do this by creating a simple workflow that includes the Run Sybase DBCC Checks step included in this solution pack.

If you find errors in the source database, be sure to fix them before running this workflow. The workflow does not have the ability to diagnose or remediate problems in the database prior to performing the database dump.

How to Run this Workflow

This topic explains how to customize and run the [Dump And Load Sybase Database](#) workflow in your environment.

Note: Prior to running this workflow, review the [Prerequisites for this Workflow](#), and ensure that all requirements are satisfied.

Tip: For detailed instructions to run HP DMA workflows—using the Run Oracle Compliance Audit workflow as an example—see *HP DMA Quick Start Tutorial*.

To customize and run the Dump and Load Sybase Database workflow:

1. Create a deployable copy of the workflow (see "Create a Deployable Workflow" in *HP DMA Quick Start Tutorial*).
2. Determine the values that you will specify for the following parameters. This is the minimum set of parameters required to run this workflow.

Parameter Name	Default Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Cache File	no default	Database cache file associated with this database dump. This is a single filename (with absolute path—path must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.
Dump File Compression Level	7	Compression level (1-9) to apply to the dump file (or files) that will be created.

Parameter Name	Default Value	Description
Dump File Password	no default	<p>Password required to encrypt and decrypt the database dump file.</p> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p>
Sybase ASE Home Directory	/home/sybase/ASE_15	<p>Sybase ASE installation home directory, where the destination database resides. Sybase ASE will examine the interface file that exists in this directory to determine where to load the specified database dump file (or files).</p> <p>If the Sybase ASE installation home directory is the same on the source and the destination servers, you do not need to specify this parameter. The default is assumed for the source—if you want to specify a different home directory for the source, you will need to expose the Sybase ASE Home Directory parameter in the Gather Advanced Parameters for Sybase Database Dump step (see step 3).</p>

Note: This is the minimum set of parameters required to run this workflow. You may need to expose additional parameters depending on your provisioning objectives.

See [Parameters for Load Sybase Database Dump](#) on page 43 for detailed descriptions of all input parameters for this workflow, including default values.

Tip: To avoid having to re-enter passwords whenever they change, you can create a policy to provide them to the workflow (see [How to Use a Policy to Specify Parameter Values](#) on page 72).

3. In the workflow editor, expose any additional parameters that you need (see [How to Expose Additional Workflow Parameters](#) on page 71). You will specify values for these parameters when you create the deployment.
4. Save the changes to the workflow (click **Save** in the lower right corner).
5. Create a new deployment (see "Create a Deployment" in *HP DMA Quick Start Tutorial* for instructions).
 - a. On the Targets tab, select all the target servers—both source and destination—that will participate in this database refresh. The targets that you select here will be available in the Target Parameters drop-down menus on the Run page (see [step 7](#)).

- b. On the Parameters tab, specify values for the required parameters listed in [step 2](#) and any additional parameters that you exposed in [step 3](#). You do not need to specify values for those parameters whose default values are appropriate for your environment.
6. Save the deployment (click **Save** in the lower right corner).
7. Run the workflow using this deployment (see "Run Your Workflow" in *HP DMA Quick Start Tutorial* for instructions).

On the Run page, select the following targets from the respective drop-down menus:

Parameter Name	Default	Description
Source Instance	no default	<p>The Adaptive Server instance where the dump file will be created. You specify this when you run the workflow.</p> <p>Note: The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.</p>
Target Instance	no default	<p>The Adaptive Server instance where the destination database will be loaded from the dump file (or files). You specify this when you run the workflow.</p> <p>Note: The Target Instance that you specify at run time must match the Target Database Instance Name that you specify in the deployment.</p>

The workflow will complete and report "Success" on the Console if it has run successfully. If an invalid parameter value is specified, an error is logged, and the workflow terminates in the "Failure" state.

Sample Scenarios

This topic shows you how to use various parameters to achieve the following database refresh scenarios in your environment using the [Dump And Load Sybase Database](#) workflow:

Scenario 1: Perform a Database Refresh Using a Dump File is Not Encrypted or Striped

This is the simplest Sybase ASE database refresh scenario.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Scenario 2: Perform a Database Refresh Using a Dump File that is Compressed and Encrypted

This scenario requires you to specify the encryption password and compression level for the database dump file.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Compression Level	8	Compression level (1-9) to apply to the dump file (or files) that will be created.
Dump File Password	MyPassword1@#	<p>Password required to encrypt and decrypt the database dump file.</p> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p>

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Scenario 3: Perform a Database Refresh Using a Dump File that is Striped

In this scenario, the database dump file will be striped across multiple files. You must specify all the individual stripe files in the Dump File List parameter (separate them with commas). If the stripe files are encrypted, you must also specify the Dump File Password parameter.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	MyPassword1@#	<p>Password required to encrypt and decrypt the database dump file.</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p>Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).</p> </div>

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Scenario 4: Perform a Database Refresh Using a Cache Dump File

In this scenario, the database dump file has an associated cache dump file. You must specify the name of the cache dump file by using the Cache File parameter. The workflow will rebuild and bind the cache after the database dump file is loaded into the target database.

If the database dump file is encrypted, you must specify the Dump File Password parameter.

If the cache dump file is encrypted, you must specify the Cache Dump File Password parameter.

Parameter Name	Example Value	Description
Call Wrapper	jython	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Dump File List	<code>/var/tmp/mytestdb.dmp</code>	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	MyPassword1@#	Password required to encrypt and decrypt the database dump file. Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris).
Cache File	<code>/var/tmp/runcache_mytestdb.txt</code>	Database cache file associated with this database dump. This is a single filename (with absolute path—path must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.

Be sure that the default values for all remaining parameters are appropriate for your environment (see [Parameters for Load Sybase Database Dump](#) on page 43).

Parameters for Dump and Load Sybase Database

The following tables describe the required and optional input parameters for this workflow. Most of these parameters are not initially visible in a deployment (see [How to Expose Additional Workflow Parameters](#) on page 71). For most parameters, if you do not specify a value for a parameter, a default value is assigned in the following steps:

- Gather Parameters for Sybase Dump or Load
- Gather Advanced Parameters for Sybase Database Dump

Input Parameters for the Dump Sybase Database Workflow

Parameter Name	Default Value	Required	Description
ASE SysAdmin Password	password	required	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
ASE SysAdmin Username	sa	required	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database dump and load operations.
Cache Dump File	no default	optional	Database cache file associated with this database dump. This is a single filename (with absolute path—path must exist). The file contains detailed information about any specific (non-default) data caches used by the source database and any database objects bound to those caches.
Call Wrapper	jython	required	Command that will be used to construct the call wrapper. The workflow uses the call wrapper to execute subsequent steps as either the OS administrative user (for example: <code>sudo su - root /opt/hp/dma/client/bin/jython.sh</code>) or the owner of the Sybase ASE installation (for example: <code>sudo su - sybase /opt/hp/dma/client/bin/jython.sh</code>).
DBCC Checks	checkdb, checkalloc, checkcatalog	optional	List of database consistency checker (DBCC) checks that you want to run to ensure that there are no problems with the database after the dump file is loaded.

Input Parameters for the Dump Sybase Database Workflow (continued)

Parameter Name	Default Value	Required	Description
DBCC Error Directory	<code>/var/tmp</code>	optional	The directory (with absolute path) where you want to store the DBCC results (output files) for post-load checks. This directory must exist.
Dump Device Name	n/a	n/a	Not used in this release.
Dump File Compression Level	7	optional	Compression level (1-9) to apply to the dump file (or files) that will be created.
Dump File List	<code>/var/tmp/< databasename>_<datetime>.dmp</code> For example: <code>/var/tmp/mytestdb_2012111283762.dmp</code> If you specify a non-default file name (or names) in the Dump File List parameter, the path to each specified file must exist.	required	Comma-separated list of database dump files (with absolute paths—all specified paths must exist). For a single dump file, no comma is necessary.
Dump File Password	no default	optional	Password required to encrypt and decrypt the database dump file. <div style="background-color: #f0f0f0; padding: 5px;"> Note: You cannot use an encrypted dump file to perform a cross-platform refresh when an architectural endian difference exists (for example: create dump on Linux, load dump on Solaris). </div>
Dump File Path	n/a	n/a	Not used in this release.
Local Backup Instance Name	n/a	n/a	Not used in this release.

Input Parameters for the Dump Sybase Database Workflow (continued)

Parameter Name	Default Value	Required	Description
Local Backup Instance Port	n/a	n/a	Not used in this release.
Remote Backup Instance Name	n/a	n/a	Not used in this release.
Remote Backup Instance Port	n/a	n/a	Not used in this release.
Role Password SQL Statement	no default	required	Not used in this release.
Source Database Instance Name	NY_DS	required	<p>Name of the Adaptive Server instance where the dump file (or files) will be created. You specify the value of this parameter in the deployment.</p> <p>Note: The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.</p>
Source Database Name	mytestdb	required	Name of database from which the dump file (or files) will be created.
Source Instance	no default	required	<p>The Adaptive Server instance where the dump file will be created. You specify this when you run the workflow.</p> <p>Note: The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.</p>

Input Parameters for the Dump Sybase Database Workflow (continued)

Parameter Name	Default Value	Required	Description
Sybase ASE Home Directory	/home/sybase/ASE_15	required	Sybase ASE installation home directory, where the source and destination databases each reside. Sybase ASE will examine the interface file that exists in this directory to determine where to first create and then load the specified database dump file (or files).
Sybase OS User Name	sybase	required	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Target Database Instance Name	NY_DS	required	Name of the Adaptive Server instance where the destination database will be loaded from the dump file (or files). You specify the value of this parameter in the deployment. Note: The Target Instance that you specify at run time must match the Target Database Instance Name that you specify in the deployment.
Target Database Name	mytestdb	required	Name of the database where the dump file (or files) will be loaded.
Target Instance	no default	required	The Adaptive Server instance where the destination database will be loaded from the dump file (or files). You specify this when you run the workflow. Note: The Target Instance that you specify at run time must match the Target Database Instance Name that you specify in the deployment.

Chapter 3

Reference Information

This page shows you where to find additional information about the Sybase ASE products discussed in this guide and about HP DMA.

Sybase ASE Product Documentation

SAP provides an extensive documentation library for Sybase ASE at this location:

<http://infocenter.sybase.com/help/index.jsp>

For information about Adaptive Server specifications—including database requirements based on page size—see this document:

[Adaptive Server Specifications](#)

HP DMA Documentation

For information about using the HP DMA web interface, see the *HP DMA User Guide*, the *HP DMA Administrator Guide*, and the *HP DMA Quick Start Tutorial*.

These documents are part of the HP DMA documentation library, which is available on the HP Software Product Manuals web site:

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

Chapter 4

Tips and Best Practices

This portion of the document contains a collection of tips and best practices that will enable you to use HP DMA more effectively. It contains the following topics:

[How this Solution is Organized](#) on the next page

[How to Expose Additional Workflow Parameters](#) on page 71

[How to Use a Policy to Specify Parameter Values](#) on page 72

[How to Import a File into the Software Repository](#) on page 75

How this Solution is Organized

In HP DMA, a workflow executes a process—such as installing a software product or creating a database.

A solution pack contains one or more related workflow templates. This solution contains the following workflow templates:

Dump Sybase Database

Use this workflow to dump the contents of a Sybase Adaptive Server Enterprise (Sybase ASE version 15) database (the source database) into a file (the database dump file).

- It performs extensive validation checks prior to and immediately after the database `dump` operation to ensure that the refresh is successful.
- It can create a database dump file that is striped, compressed, encrypted or any combination thereof.
- The dump file can be used to subsequently perform a cross-platform database refresh, if necessary.

Load Sybase Database Dump

Use this workflow to load (restore) a specific database in the target instance from an existing dump (backup) file that you provide.

In addition to quickly performing a database load, the workflow provides the following benefits:

- It performs extensive validation checks prior to and immediately after the database `load` operation to ensure that the refresh is successful.
- It can perform a cross-platform database refresh when necessary. After it performs a cross-platform load operation, the workflow rebuilds the indexes (clustered or non-clustered indexes on APL/DOL tables) to avoid any page linkage or index corruption.
- It can utilize a database dump file that is striped, compressed, encrypted or any combination thereof.
- It restores any existing database users after the refresh is completed.

The workflow will perform a "typical" database refresh using default parameter values. You can override these default values by specifying parameter values in the deployment. The input parameter values that you specify are validated before provisioning begins.

Dump And Load Sybase Database

Use this workflow to dump the contents of a Sybase ASE database (the source) into a file (the database dump file) and load the contents of that file into an existing Sybase ASE database (the destination).

The workflow performs extensive validation checks prior to and immediately after the dump operation at the source to ensure that the dump file is valid. It also performs validation checks prior to and immediately after the load operation at the destination to ensure that the data was successfully loaded.

This workflow can perform a cross-platform database refresh (dump and load) when necessary. After it performs a cross-platform load operation, the workflow rebuilds the indexes (clustered or non-clustered indexes on APL/DOL tables) to avoid page linkage or index corruption issues. Password protected dump cannot be supported for cross platform dump and load.

If any source database objects are bound to a specific (non-default) cache, the workflow will create a cache dump file—provided that you specify a valid value for the Cache Dump File parameter. The cache dump file contains details about the specific caches used by the source database and any objects that are bound to each cache. This file is in data-readable format. The workflow uses the cache dump file to refresh the destination database cache (provided that ample cache space is available). The workflow cannot, however, configure or enable cache buffering.

This workflow can create and load database dump files that are striped, compressed (at any level 1-9), encrypted, or any combination thereof.

What's Inside

Each **workflow** template has a Documentation tab that provides detailed information about that workflow.

The screenshot shows the HP Database & Middleware Automation interface. The top navigation bar includes 'Home', 'Automation', 'Reports', 'Environment', 'Solutions', and 'Setup'. Below this, a secondary navigation bar lists 'Workflows', 'Steps', 'Functions', 'Policies', 'Deployments', 'Run', 'Console', and 'History'. The main content area is titled 'My Copy of Dump And Load Sybase Database' and has tabs for 'Documentation', 'Workflow', 'Deployments', and 'Roles'. The 'Documentation' tab is active, showing the following details:

- Name: My Copy of Dump And Load Sybase Database
- Tags: BEW, Bridged, Execution, Workflow
- Type: Sybase
- Target level: Instance

The documentation text is as follows:

Purpose

This workflow enables you to dump the contents of a Sybase Adaptive Server Enterprise (ASE) database (the source) into a file (the database dump file) and load the contents of that file into an existing Sybase ASE database (the destination).

The workflow performs extensive validation checks prior to and immediately after the dump operation at the source to ensure that the dump file is valid. It also performs validation checks prior to and immediately after the load operation at the destination to ensure that the data was successfully loaded.

This workflow can perform a cross-platform database refresh (dump and load) when necessary. After it performs a cross-platform load operation, the workflow rebuilds the indexes (clustered or non-clustered indexes on APL/DOL tables) to avoid page linkage or index corruption issues. Password protected dump cannot be supported for cross platform dump and load.

If any source database objects are bound to a specific (non-default) cache, the workflow will create a cache dump file if you specify a valid value for the Cache Dump File parameter. The cache dump file contains details about the specific caches used by the source database and any objects that are bound to each cache. This file is in data-readable format. The workflow uses the cache dump file to refresh the

At the bottom of the documentation area, there are links for [HELP](#), [PDF](#), and [EDIT](#).

The footer of the interface contains several action buttons: **DELETE**, **EXPORT**, **EXTRACT POLICY**, **DEPLOY**, **Copy**, **Save**, and **CANCEL**.

A workflow consist of a sequence of [steps](#). Each step performs a very specific task. Each step includes a documentation panel that briefly describes its function.

hp Database & Middleware Automation

Home Automation Reports Environment Solutions Setup

Workflows Steps Functions Policies Deployments Run Console History

Run Sybase DBCC Checks

General Action Parameters History Workflows Solutions Roles

Properties

Name: Run Sybase DBCC Checks
Tags: Sybase
Type: Sybase
Category: Script
Targetable:

Documentation

Runs the specified database consistency checks ([DBCC](#)) against the [dataserver](#) instance.

Purpose:
The [DBCC](#) checks verify the state of the [dataserver](#) prior to installing the patch. Any errors will be stored in the error log if it was provided. Errors must be corrected before the workflow can continue.

Required Functions:
Import [StepLog](#)
Import [SybaseTools](#)
Import [Validator](#)

Copy

THIS STEP IS READ ONLY

Steps can have input and output [parameters](#). Output parameters from one step often serve as input parameters to another step. Steps can be shared among workflows.

Parameter descriptions are also displayed in the following locations in the HP DMA web interface:

On the Workflow tab for each workflow.

Gather Advanced Parameters for Sybase Database Dump / Cache File

*Optional: Full path to file where cache details will be written.

To see parameter description here

4 Gather Advanced Parameters for Sybase Database Dump 0

ASE SysAdmin Password: Gather Parameters for Sybase Dump or Load.ASE SysAdmin Password

ASE SysAdmin Username: Gather Parameters for Sybase Dump or Load.ASE SysAdmin Username

Cache File: Not selected

Call Wrapper: Prepare Sybase Call Wrapper.Instance Wrapper

Dump File Compression Level: Not selected

Click here

On the Parameters tab for each step in the workflow

HP Database & Middleware Automation

Home Automation Reports Environment Solutions Setup

Workflows Steps Functions Policies Deployments Run Console History

Validate Sybase Database Load Settings

General Action Parameters History Workflows Solutions Roles

Input parameters

Name	Value	Description
ASE SysAdmin Password	<input type="text"/>	*Required: Password for ASE SysAdmin Username
ASE SysAdmin Username	<input type="text"/>	*Required: Sybase database username that will perf
Call Wrapper	<input type="text"/>	*Required: Instance Level call wrapper.
Dump File List	<input type="text"/>	*Required: Comma-separated list of Dump files (with
Sybase ASE Home Directory	<input type="text"/>	*Required: Sybase ASE home directory
Target Database Instance Name	<input type="text"/>	*Required: Database instance name where the dump
Target Database Name	<input type="text"/>	*Required: Name of database where the database d

Output parameters

Name	Description
Dump File Path	*Optional: Directory where dump file(s) are located. Required if Dump

THIS STEP IS READ ONLY

On the Parameters tab in the deployment (organized by step)

The screenshot shows the HP Database & Middleware Automation interface. The top navigation bar includes 'Home', 'Automation', 'Reports', 'Environment', 'Solutions', and 'Setup'. Below this, a secondary navigation bar shows 'Workflows', 'Steps', 'Functions', 'Policies', 'Deployments', 'Run', 'Console', and 'History'. The main content area is titled 'New deployment' and has three tabs: 'Targets', 'Parameters', and 'Roles'. The 'Parameters' tab is active, showing two sections:

- Prepare Sybase Call Wrapper:**
 - Call Wrapper:** A text input field containing 'jython' and a dropdown menu set to 'Text'. Below it, a note states: '*Required: Command that will execute the step as the OS administrative user (for example: sudo su - root /opt/datapalette/jython/jython for UNIX targets)'
 - Sybase OS User Name:** A text input field containing 'sybase' and a dropdown menu set to 'Text'. Below it, a note states: '*Required: OS user who owns the Sybase ASE installation directory. This is the directory specified in the Sybase Install Directory parameter, where the Sybase ASE binaries will be extracted using the untar command'
- Gather Advanced Parameters for Sybase Database Load:**
 - Dump File List:** A text input field containing '/var/tmp/dump.dmp' and a dropdown menu set to 'Text'. Below it, a note states: '*Required: Comma-separated list of Dump files (with absolute paths). Filenames without path may be used if Dump File Path is populated. For a single file, no comma is necessary.'
 - Dump File Password:** A text input field containing '*****' and a dropdown menu set to 'Text'. Below it, a note states: '*Optional: Password to decrypt encrypted database. Required if dump is encrypted.'

At the bottom right of the configuration area, there are three buttons: 'Restore defaults', 'Save', and 'CANCEL'.

All parameters used by the workflows in this solution pack are also described in the [Reference Information](#) for this solution pack.

Note: The workflow templates included in this solution pack are read-only and cannot be deployed. To use a workflow template, you must first create a copy of the template and then customize that copy for your environment (see "Create a Deployable Workflow" in *HP DMA Quick Start Tutorial*).

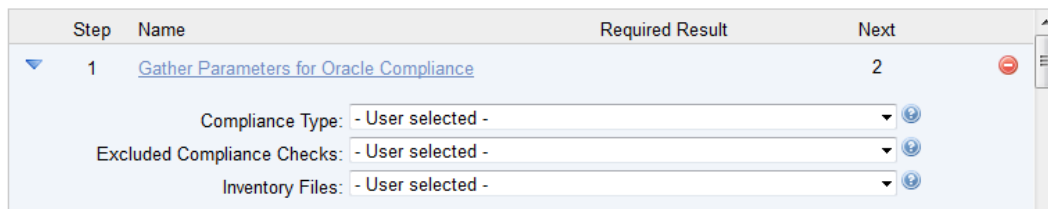
How to Expose Additional Workflow Parameters

Each workflow in this solution pack has a set of input parameters. Some are required and some are optional. To run a workflow in your environment, you must specify values for a subset of these parameters when you create a deployment.

By default, only a few of the input parameters for each workflow are visible on the Deployment page, and the rest are hidden. In order to specify a value for a parameter that is currently hidden, you must first expose that parameter by changing its mapping in the workflow editor.

To expose a hidden workflow parameter:

1. In the HP DMA web interface, go to Automation > Workflows.
2. From the list of workflows, select a deployable workflow.
3. Go to the Workflow tab.
4. In the list of steps below the workflow diagram, click the ▶ (blue arrow) to the immediate left of the pertinent step name. This expands the list of input parameters for this step.
5. For the parameter that you want to expose, select - User Selected - from the drop-down list.
For example:



Step	Name	Required Result	Next
1	Gather Parameters for Oracle Compliance		2

Compliance Type: - User selected -

Excluded Compliance Checks: - User selected -

Inventory Files: - User selected -

6. Repeat steps 4 and 5 for all the parameters that you would like to specify in the deployment.
7. Click **Save** in the lower right corner.

How to Use a Policy to Specify Parameter Values

It is sometimes advantageous to provide parameter values by using a policy rather than explicitly specifying the values in a deployment. This approach has the following advantages:

- The policy can be used in any deployment.
- It is faster and less error-prone than specifying parameter values manually.
- For parameter values that change frequently—for example, passwords that must be changed regularly—you only need to update them in one place.

To establish a policy, you can either [Create a Policy](#) or [Extract a Policy](#) from a workflow.

After you establish the policy, you must [Reference the Policy in the Deployment](#).

For more information, see the *HP DMA User Guide*. This document is available on the HP Software Product Manuals web site: <http://h20230.www2.hp.com/selfsolve/manuals>

Create a Policy

The first step in this approach is to create a policy that provides parameter values. There are two ways to do this: (1) create a new policy, and define all attributes manually (as shown here) or (2) extract a policy from a workflow (see [Extract a Policy](#) on the next page).

To create a policy that provides parameter values:

1. In the HP DMA web UI, go to Automation > Policies.
2. Click **New Policy**.
3. In the **Name** box, specify the name of the policy
4. For each parameter value that you want to provide using this policy, perform the following actions on the Attributes tab:
 - a. From the drop-down list, select the type of attribute:
 - A Text attribute contains simple text that users can view while deploying and running workflows.
 - A List attribute contains a comma-separated list of values (or a large amount of text not suitable for a Text attribute).
 - A Password attribute contains simple text, but the characters are masked so that users cannot see the text.
 - b. In the text box to the left of the Add button, specify the name of the attribute.
For your convenience, this name should be similar to the parameter name used in the pertinent workflow (or workflows).
 - c. Click **Add**.
 - d. In the new text box to the right of the attribute's name, enter a value for this attribute.
To remove an attribute, click the **Remove** button.
5. On the Roles tab, grant Read and Write permission to any additional users and groups who will

be using this policy. By default, any groups to which you belong have Read and Write permission.

6. Click the **Save** button (lower right corner).

Extract a Policy

An alternative to creating your own policy one attribute at a time is to extract the policy. This automatically creates a reusable policy that provides values for all input parameters associated with a workflow. This is a convenient way to create a policy.

To extract a policy:

1. Go to Automation > Workflows.
2. Select the Workflow that you want to work with.
3. Click the Extract Policy link at the bottom of the screen.
4. Specify values for each attribute listed.
5. *Optional:* Remove any attributes that you do not want to use.
6. *Optional:* Add any new attributes that you want to use.
7. *Optional:* On the Roles tab, select the Read box for any users or user groups that you want to be able to use this policy to provide parameter values in a Deployment. Select the Write box for any users or groups that you want to be able to modify this Policy (add or remove attributes).
8. Click **Save**.

Reference the Policy in the Deployment

After you create a policy, you can reference its attributes in a deployment.

To reference policy attributes in a deployment:

1. Create or access the deployment.
See “Deployments” in the *HP DMA User Guide* for details.
2. On the Parameters tab, perform the following steps for each parameter whose value you want to provide by referencing a policy attribute:
 - a. In the drop-down menu for that parameter, select **Policy Attribute**.
 - b. In the text box for that parameter, type any character. A drop-down list of policy attributes appears. For example:

Admin Password: Policy Attribute

- Discovery.Web Service Password
- DTE - Policy.Password
- MyParameterValues.MyAdminPassword**
- MyParameterValues.MyAdminUser
- MyParameterValues.MyDBUser
- MyParameterValues.MyDBUserPassword
- oracle software.oracle software

c. From the drop-down list, select the attribute that you want to reference. For example:

Admin Password: Policy Attribute

3. Click **Save** to save your changes to the deployment.

How to Import a File into the Software Repository

Many HP DMA workflows are capable of downloading files from the software repository on the HP DMA server to the target server (or servers) where the workflow is running. The following procedure shows you how to import a file into the software repository so that it can be downloaded and deployed by a workflow.

HP DMA uses the HP Server Automation (HP SA) Software Library as its software repository.

Tip: Be sure to use unique file names for all files that you import into the software repository.

To import a file into the HP SA Software Library:

1. Launch the HP SA Client from the Windows Start Menu.

By default, the HP SA Client is located in Start → All Programs → HP Software → HP Server Automation Client

If the HP SA Client is not installed locally, follow the instructions under “Download and Install the HP SA Client Launcher” in the *HP Server Automation Single-Host Installation Guide*.
2. In the navigation pane in the HP SA Client, select Library → By Folder.
3. Select (or create) the folder where you want to store the file.
4. From the Actions menu, select **Import Software**.
5. In the Import Software dialog, click the **Browse** button to the right of the File(s) box.
6. In the Open dialog:
 - a. Select the file (or files) to import.
 - b. Specify the character encoding to be used from the Encoding drop-down list. The default encoding is English ASCII.
 - c. Click **Open**. The Import Software dialog reappears.
7. From the Type drop-down list, select **Unknown**.
8. If the folder where you want to store the files does not appear in the Folder box, follow these steps:
 - a. Click the **Browse** button to the right of the Folder box.
 - b. In the Select Folder window, select the import destination location, and click **Select**. The Import Software dialog reappears.
9. From the Platform drop-down list, select all the operating systems listed.
10. Click **Import**.

If one of the files that you are importing already exists in the folder that you specified, you will be prompted regarding how to handle the duplicate file. Press F1 to view online help that explains the options.
11. Click **Close** after the import is completed.

Chapter 5

Troubleshooting

These topics can help you address problems that might occur when you install and run the workflows in this solution pack:

- [Target Type](#) below
- [User Permissions and Related Requirements](#) below
- [Discovery in HP DMA](#) on the next page

Target Type

In your deployment, make sure that you have specified the correct type of target. The workflow type and the target type must match. A workflow designed to run against an instance target, for example, cannot run against a server target.

User Permissions and Related Requirements

Roles define access permissions for organizations, workflows, steps, policies, and deployments. Users are assigned to roles, and they gain access to these automation items according to the permissions and capabilities defined for their roles.

Roles are assigned by your server management tool administrator. They are then registered in HP DMA by your HP DMA administrator.

Your HP DMA administrator will ensure that the users in your environment are assigned roles that grant them the permissions and capabilities they need to accomplish their tasks. For example:

- To create a workflow, your role must have Workflow Creator capability.
- To view a workflow, your role must have Read permission for that workflow.
- To edit a workflow, your role must have Write permission for that workflow.
- To view a deployment, your role must have Read permission for that deployment.
- To modify a deployment, your role must have Write permission for that deployment.
- To run a deployment, your role must have Execute permission for that deployment and Deploy permission for the organization where it will run.

Capabilities determine what features and functions are available and active in the HP DMA UI for each user role.

For more information, see the *HP DMA Administrator Guide*. This document is available on the HP Software Product Manuals web site: <http://h20230.www2.hp.com/selfsolve/manuals>

Discovery in HP DMA

HP DMA uses a process called “discovery” to find information about the servers, networks, and database instances on target machines in your managed environment.

You must explicitly initiate the process of discovery—it is not automatic. See the *HP DMA User Guide* for instructions. This document is available on the HP Software Product Manuals web site: <http://h20230.www2.hp.com/selfsolve/manuals>

Glossary

A

automation items

The umbrella term automation items is used to refer to those items to which role-based permissions can be assigned. Automation items include workflows, deployments, steps, and policies.

B

bridged execution

A bridged execution workflow includes some steps that run on certain targets and other steps that run on different targets. An example of a bridged execution workflow is Extract and Refresh Oracle Database via RMAN (in the Database Refresh solution pack). This workflow extracts the contents of a database on one target (the Source) and creates a new database with the same contents on another target (the Destination). This workflow is useful when you want to clone a database - for example, to move it from a traditional IT infrastructure location into a private cloud. Bridged execution workflows are supported on HP DMA version 9.11 (and later).

C

capability

Capabilities are collections of related privileges. There are three capabilities defined in HP DMA. Login Access capability enables a user to log in to the web interface. This capability does not guarantee that this user can view any

organizations or automation items—permissions are required to access those items. Workflow Creator capability enables a user to create new workflows and make copies of other workflows. Administrator capability enables a user to perform any action and view all organizations. If you have Administrator capability, you do not need Workflow Creator capability. The Administrator can assign any of these capabilities to one or more roles registered roles.

connector

HP DMA includes a Connector component that enables it to communicate with your server management tool. You must configure the Connector before you can run an workflow against a target.

cross-platform

Cross-platform database refresh involves converting the data from one type of byte ordering to another. This is necessary, for example, if you want to load a database dump file on a little-endian Linux target that was created on a big-endian Solaris server.

custom field

Custom Fields are used to customize workflows or show information about the environment. Custom Fields can be used in workflow steps to automatically supply information that is specific to an organization, server, instance, or database.

D

Data Pump

Oracle Data Pump is a utility that enables you to move data or metadata from one database to another. You can use Data Pump to move a complete database or a subset of a database.

deployment

Deployments associate a workflow with a target environment in which a workflow runs. You can customize a deployment by specifying values for any workflow parameters that are designated - User Selected - in the workflow. You must save a deployment before you can run the workflow. You can re-use a saved deployment as many times as you like.

destination

In a database refresh scenario, the contents of a database dump file are loaded into the destination database.

F

function

Functions are reusable pieces of code that can be included in automation steps. Any common routine or operation that multiple steps perform is a good candidate for a function. Functions can be tagged with keywords indicating the language in which they are written and the operating system with which they work. Functions are "injected" into the step code just prior to step execution.

I

input parameters

A workflow has a set of required parameters for which you must specify a value. The required parameters are a subset of all the parameters associated

with that workflow. The remaining parameters are considered optional. You can specify a value for an optional parameter by first exposing it using the workflow editor and then specifying the value when you create a deployment.

M

mapping

An input parameter is said to be "mapped" when its value is linked to an output parameter from a previous step in the workflow or to a metadata field. Mapped parameters are not visible on the Deployment page. You can "unmap" a parameter by specifying - User Selected - in the workflow editor. This parameter will then become visible on the Deployment page.

O

organization

An organization is a logical grouping of servers. You can use organizations to separate development, staging, and production resources - or to separate logical business units.

P

parameters

Parameters are pieces of information - such as a file system path or a user name - that a step requires to carry out its action. Values for parameters that are designated User Selected in the workflow can be specified in the deployment. Parameters that are marked Enter at Runtime in the deployment must be specified on the target system when the workflow runs.

policy

Policies are reusable sets of attributes that can be used as parameter values in

deployments. Deployments can reference policy attributes to change the automation behavior. Policies provide values for input parameters. They can contain fixed values or reference Custom Fields. Policies enable HP DMA to manage groups of hundreds or thousands of servers at a time without the need to configure each individual server.

R

raw devices

In Sybase ASE version 15, you can create and mount database devices on raw bound devices. This enables Sybase ASE to use direct memory access from your address space to the physical sectors on the disk. This can improve performance by reducing memory copy operations from the user address space to the operating system kernel buffers.

Recovery Manager (RMAN)

Oracle Recovery Manager (RMAN) is a backup and recovery tool included in Oracle Database Enterprise Edition (and related products). RMAN enables you to efficiently backup and restore data files, control files, server parameter files, and archived redo log files. It provides block-level corruption detection during both the backup and restore phases. It is optimized for performance and space consumption.

role

Each HP DMA user has one or more roles. Roles are used to grant users permission to log in to and to access specific automation items and organizations. Roles are defined in your server management tool. Before you can associate a role with an automation item or organization, however, you must register that role in HP DMA.

S

smart group

Smart Groups are dynamic groups of servers, instances, or databases defined by some criteria. They are used to specify targets for deployments. As information about an environment object changes, its membership in the groups is re-evaluated.

software repository

The software repository is where the workflow will look for any required files that are not found on the target server. If you are using HP DMA with HP Server Automation (SA), this repository is the SA Software Library.

solution pack

A solution pack contains one or more related workflow templates. These templates are read-only and cannot be deployed. To run one of the workflows included in a solution pack, you must first create a deployable copy of that template and then customize that copy for your environment. Solution packs are organized by function - for example: database patching or application server provisioning.

source

In a database refresh scenario, the contents of the SOURCE database are extracted and stored in a file (or multiple files).

source database

In the context of MS SQL database refresh, the "source database" is the database from which the backup file is created.

steps

Steps contains the actual code used to perform a unit of work detailed in a

workflow.

T

target instance

In the context of MS SQL database refresh, the term "target instance" refers to the SQL Server instance where the database that will be restored resides.

W

workflow

A workflow automates the process followed for an operational procedure. Workflows contain steps, which are linked together to form business logic for a common task. Workflows connect existing tasks in order to perform a new business process by building on existing best practices and processes.

workflow editor

The workflow editor is the tool that you use to assemble steps into workflows. You can map each input parameter to output parameters of previous steps or built-in metadata (such as the server name, instance name, or database name). You can also specify User Selected to expose a parameter in the deployment; this enables the person who creates the deployment to specify a value for that parameter.

workflow templates

A workflow template is a read-only workflow that cannot be deployed. To run one of the workflows included in a solution pack, you must first create a deployable copy of the workflow template and then customize that copy for your environment.