HP Database and Middleware Automation Solution Packs

For the Linux, Solaris, AIX, HP-UX, and Windows operating systems

Software Version: 9.14

Database Refresh - Sybase ASE Workflows

Document Release Date: June 2012

Software Release Date: June 2012

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

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

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

UNIX® is a registered trademark of The Open Group.

Documentation Updates

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

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

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

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

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

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

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

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

Support

Visit the HP Software Support Online web site at:

http://www.hp.com/go/hpsoftwaresupport

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

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

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

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

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

To find more information about access levels, go to:

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

Contents

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, & 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.

Chapter 1

Quick Start Tutorial

This topic shows you how to install your solution pack and run a workflow. There are five basic steps:

1. Install the Solution Pack below
2. Create a Deployable Workflow on next page
3. Create a Deployment on page 11
4. Run Your Workflow on page 12
5. View the Results on page 12

This tutorial provides a simplified demonstration using the Dump Sybase Database workflow. Default values are supplied for most input parameters. Before executing these procedures, make sure that these default values are suitable for your environment.

Note: See the Reference Information included in this guide for descriptions of all available input parameters for this workflow, including default values.

The information presented in this tutorial assumes the following:

- HP DMA is installed and operational.
- At least one valid target is available.

Note: For information about other automation scenarios, see How To Use This Solution.

Install the Solution Pack

The following instructions assume that you have purchased the Database Refresh solution pack.

To install the solution pack:

1. Go to HP Live Network to view a list of the latest available HP DMA solution packs.
2. Download the pertinent solution pack file from HP Software Support Online.
3. Extract the ZIP file that contains your solution pack (for example: DBRefresh.zip).
4. On the system where you downloaded the solution pack, open a web browser, and log in to the HP DMA server using an account with administrative privileges.

   For instructions, see “Getting Started” in the *User Guide: Database and Middleware Automation*. This guide is included in the HP Server Automation documentation library (SA version 9.10 and later).

5. On the Solutions > Installed tab, click the **Browse** button in the lower right corner. The Choose File dialog opens.
6. Locate and select the ZIP file that you extracted in step 3, and click **Open**.
7. Click **Import solution pack**.

Create a Deployable Workflow

The workflow templates provided by HP in your solution pack are read-only and cannot be deployed. When you are viewing a read-only item in the HP DMA web UI, you will see the lock icon in the lower right corner:

Read-only workflows are not deployable. You can create a deployable workflow by making a copy of a workflow template.

To create a deployable copy of the workflow template:

1. In the HP DMA web interface, go to Automation > Workflows.
2. From the list of workflows, select the Dump Sybase Database workflow template.
3. Click the **Copy** button in the lower left corner.
4. On the Documentation tab, specify the following:
   - Name – Name that will appear in the list of available workflows
   - Tags – Keywords that you can use later to search for this workflow (optional)
   - Type – Must be Sybase
   - Target level – Must be an Instance
5. On the Roles tab, grant Read access to at least one user or group and Write access to at least one user or group.
6. Click **Save**.

   Your new workflow now appears in the list of available workflows, and the following message is displayed:

   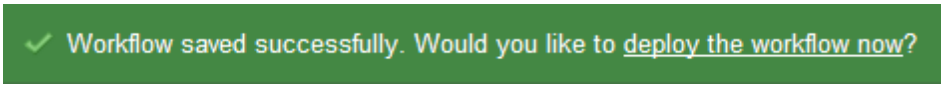

7. Click the **deploy the workflow now** link in the green message bar.

For more information about creating and working with workflows, see "Workflows" in the *User Guide: Database and Middleware Automation*. This guide is included in the HP Server Automation documentation library (SA version 9.10 and later).

Create a Deployment

Before you can run your new workflow, you must create a deployment. A deployment associates a workflow with one or more specific targets (in this case, an Instance).

To create a deployment:

1. If you do not see the green message bar—for example, if you navigated to another page after you created your copy of the workflow template—follow these steps:
   a. Go to the Automation > Deployments page.
   b. In the lower right corner, click **New deployment**.
2. Specify the following:
   - Name – Name that will appear in the list of available deployments.
   - Workflow – From the drop-down list, select the workflow that you just created.
   - Schedule – Frequency or date when the workflow will run. If you select None, the workflow will run only once when you explicitly tell it to run.
3. From the list of AVAILABLE servers on the left side of the Targets area, click the **ADD** link for the target (or targets) where the workflow will run.

   Note: If you are running a bridged execution workflow, the targets that you select on the Deployment page will be included in the lists of available targets that you can choose from on the Run page.

   For more information about bridged execution workflows, see the *User Guide: Database and Middleware Automation*. This guide is included in the HP Server Automation documentation library (SA version 9.10 and later).
4. On the Parameters tab, specify values for the input parameters listed there.

   These are a subset of the required parameters for this workflow. Parameters that are not visible in the deployment will have default values.

   Note: See the Reference Information included in this guide for descriptions of all available input parameters for this workflow, including default values.
5. If you do not want to explicitly enter the values here, you can create a policy that stores the values and then reference that policy in your deployment (see Using a Policy to Specify Parameter Values on page 112).
6. Click **Save**.

   Your new deployment now appears in the list of available workflows, and the following message is displayed:

   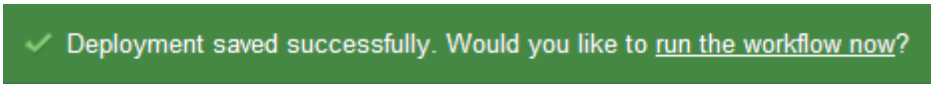

7. Click the **run the workflow now** link in the green message bar.

Run Your Workflow

Now you are ready to run your workflow against the server that you selected.

To run the workflow:

1. If you do not see the green message bar—for example, if you navigated to another page after you created your deployment—follow these steps:
   a. Go to the Automation > Run area.
   b. In the list of WORKFLOWS on the left side, select the workflow that you created.
   c. In the list of DEPLOYMENTS in the center, double-click the deployment that you just created.
2. If you are running a single-target workflow, select the check box for each target where you want to run the workflow.

   If you are running a bridged execution workflow, click the **SELECT** link to specify each target. The targets that are available to choose from here are the targets that you selected on the Deployment page.

   For more information about bridged execution workflows, see the *User Guide: Database and Middleware Automation*. This guide is included in the HP Server Automation documentation library (SA version 9.10 and later).
3. Click the **Run workflow** button.
4. The following message is displayed:

   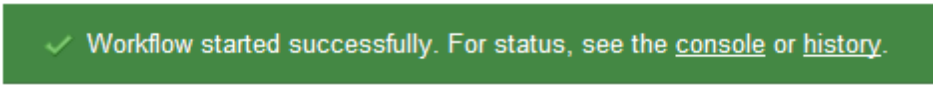

5. To view the progress of your deployment, click the **console** link in the green message bar.

View the Results

While your workflow is running, you can watch its progress on the Automation > Console page.

- To view the progress of the workflow as the deployment proceeds, click the workflow name in the upper box on the Console page.
- To view the outcome of a specific step, select that step in the left box in the Output area. Informational messages are displayed in the right box, and the values of any output parameters are listed.

While the workflow is running, its status indicator on the Console says RUNNING. After the workflow finishes, its status indicator changes to SUCCESS, FAILURE, or FINISHED.

After the workflow has finished running, you can view a summary of your deployment on the History page. This page lists all the deployments that have run on this HP DMA server during the time period specified in the Filter box.

While the workflow is running, the History page shows nothing in the status column. A workflow that results in the FINISHED state also shows nothing in the status column on this page.

To view step-by-step results, select the row in the table that corresponds to your deployment. The tabs below the table show you information about each step in the workflow. This includes the start and end time for each step, the exit code, and the following information:

- Output tab – any informational messages that were produced
- Errors tab – any errors that were reported
- Header tab – values assigned to any output parameters

Chapter 2

About this Solution

The HP Database and Middleware Automation Solution Packs Database Refresh solution contains three Sybase ASE workflows:

Workflow Name	Purpose
Dump Sybase Database on page 26	Dumps the contents of a Sybase ASE database into a file.
Load Sybase Database Dump on page 36	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 chapterprovide the following contextual information about these workflows:

- Audience below
- Supported Products and Platforms on next page
- Prerequisites for this Solution on page 16
- How this Solution is Organized on page 18
- Additional Resources

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 Sybase ASE knowledge is required to run the workflows in this solution pack using the default settings.

To customize this solution, however, you should be familiar with the following Sybase ASEprocesses:

- Database administration
- Database migration
- Database backup and restore

You should also have hands-on experience upgrading or downgrading a large database (see Additional Resources on page 23).

Supported Products and Platforms

The Sybase – Database Refresh workflow supports the following database refresh scenarios on Linux and Solaris platforms:

Dump File Version	Target Instance Version
Sybase ASE15.0.3, 15.5, or 15.7	Sybase ASE15.0.3, 15.5, or 15.7
Sybase ASE12.5.4	Sybase ASE15.0.3

Operating Systems

For specific operating system versions supported, see the *HP Database and Middleware Automation Solution Packs version 9.14 Support Matrix* available in the HP Software product manuals library located here: http://h20230.www2.hp.com/selfsolve/manuals

Hardware Requirements

- See the *HP Server Automation Quick Reference: SA Installation Requirements* or the *HP Server Automation Standard/Advanced Installation Guide*.
- For Sybase ASE hardware and software requirements, see the Sybase ASE Product Documentation on page 110

HP Software Requirements

This solution can be used with the following HP products:

- HP Server Automation version 9.11 (or later)
- HP Database and Middleware Automation Web Server version 6.0.17 (or later)

Bridged execution workflows can only be used with HP Server Automation version 9.11 (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 able to communicate with each other.
6. The master database system table (sysservers) 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.

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.

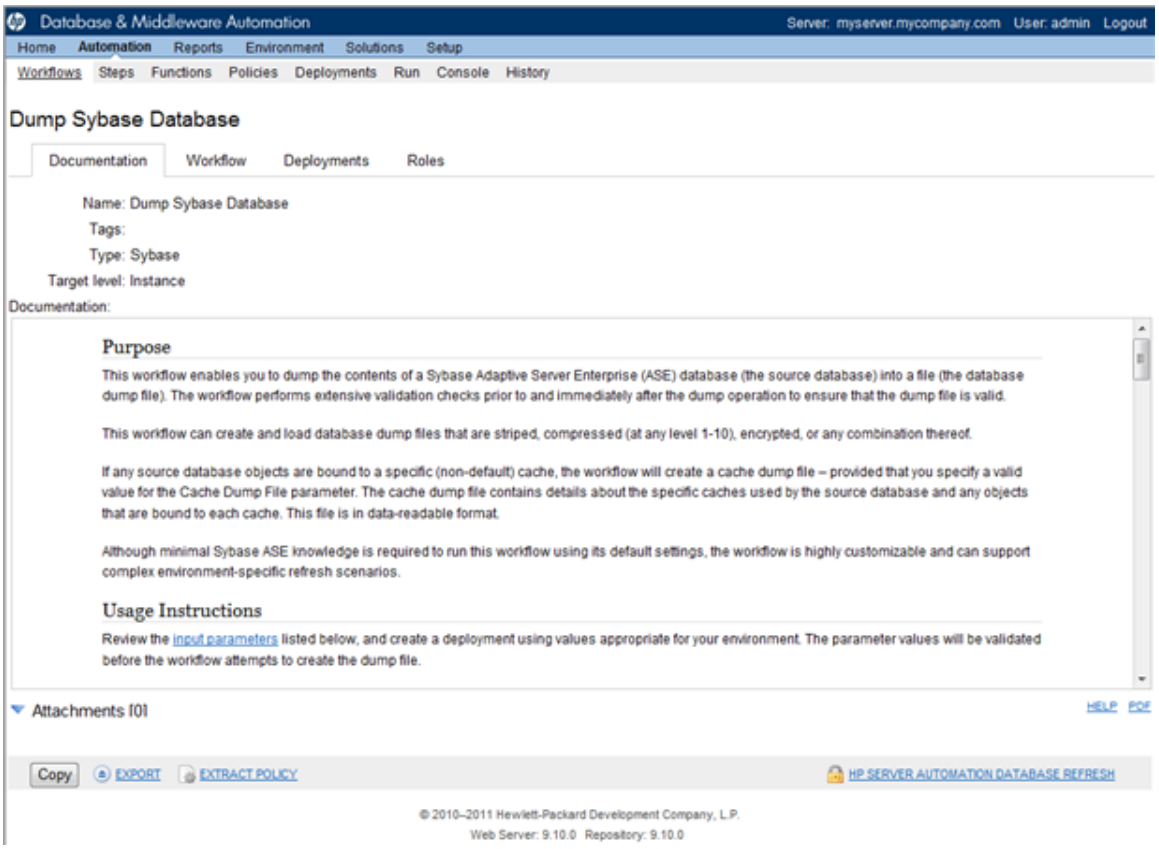

A workflow consist of a sequence of steps. Each step performs a very specific task. Each step includes a documentation panel that briefly describes its function.

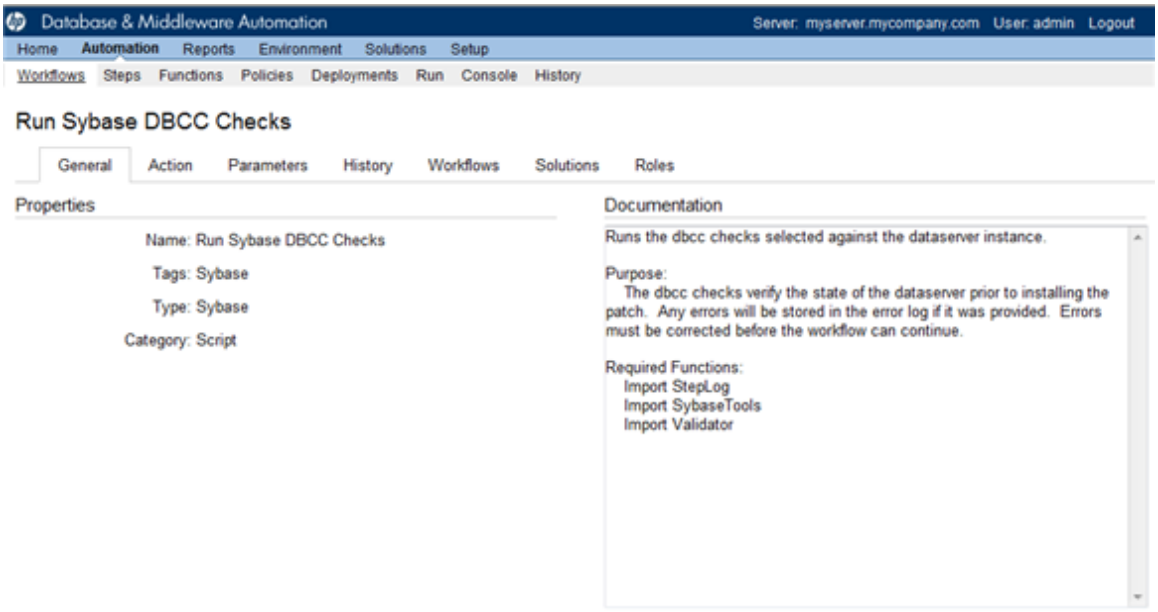

Steps can have input and output parameters. Output parameters from one step often serve as input parameters to another step. Steps can be shared among workflows.

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

On the Workflow tab for each workflow.

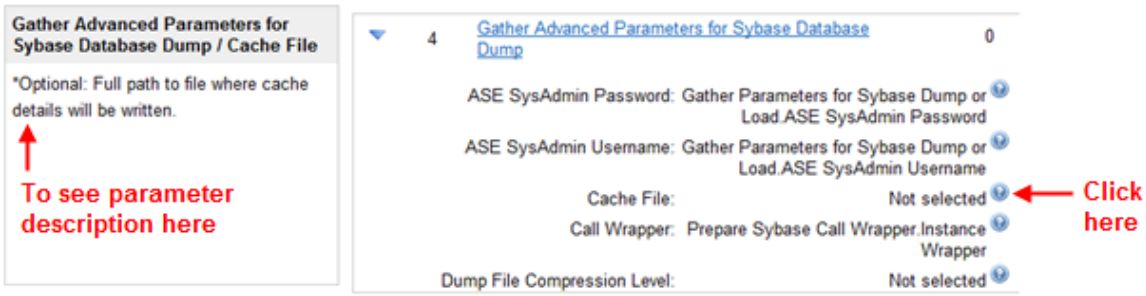

On the Parameters tab for each step in the workflow

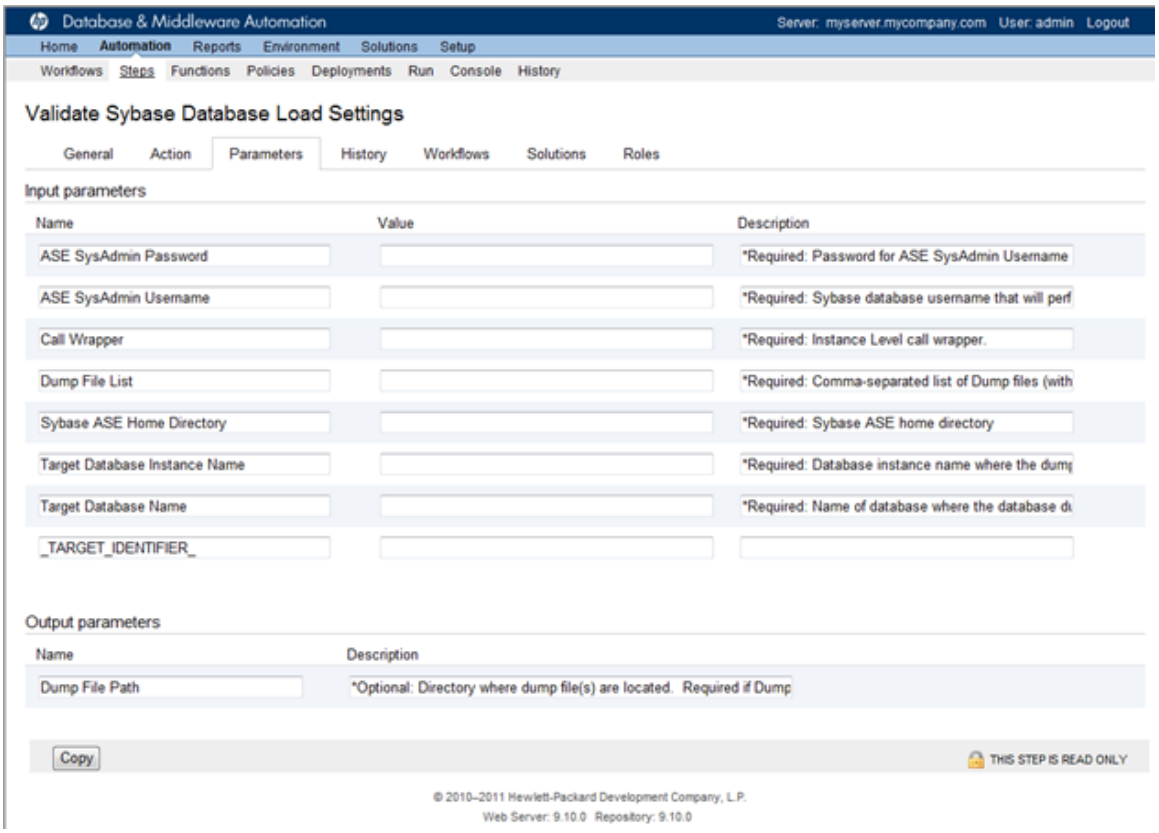

On the Parameters tab in the deployment (organized by step)

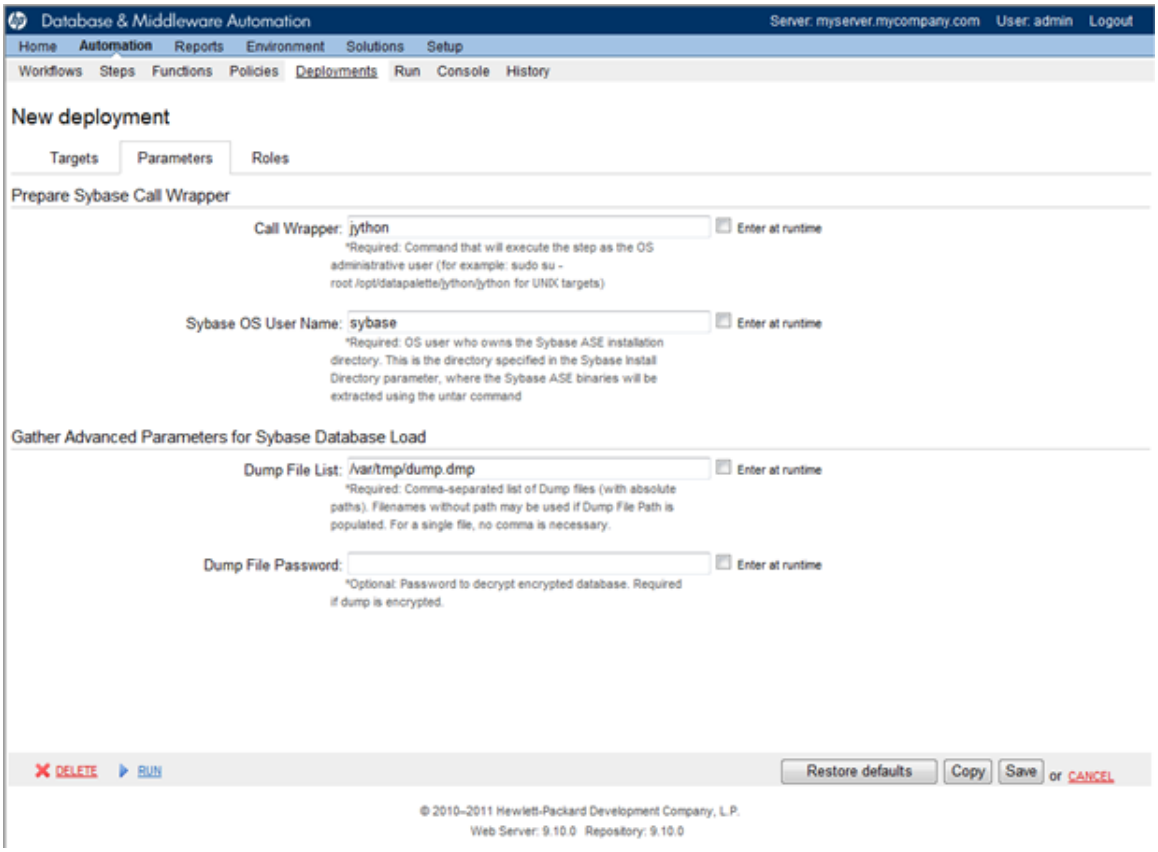

All parameters used by the workflows in this solution pack are also described in the Reference Information for this solution pack.

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 on page 10

Additional Resources

If you are using HP Server Automation version 9.10 (or later), see these documents:

- *HP Server Automation User Guide: Application Deployment Manager*
- *HP Server Automation User Guide: Database and Middleware Automation*

If you are using HP Server Automation version 9.0x, see these documents:

- *HP Database and Middleware Automation User Guide*
- *HP Server Automation Integration Guide*

If you are using HP Database and Middleware Automation version 1.00, see these documents:

- *HP Database and Middleware Automation Installation Guide*
- *HP Database and Middleware Automation User Guide*

How to Use this Solution

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.

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.

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 15).
- 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.

Note: All parameters used by each workflow in this solution are described in the Reference Information on page 61.

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 (see Create a Deployable Workflow on page 10).
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:

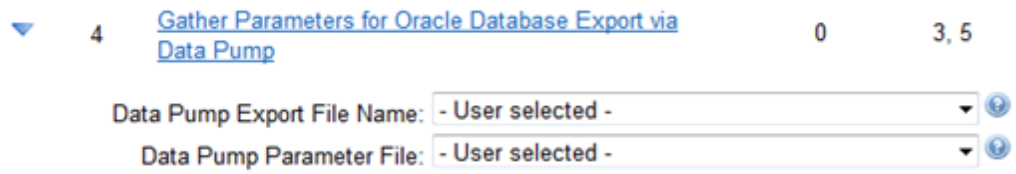

6. Repeat steps 4 and 5 for all the parameters that you would like to specify in the deployment.
7. Click **Save** in the lower right corner.

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 on next page	List of prerequisites that must be satisfied before you can run this workflow
How this Workflow Works on page 29	Information about what the workflow does, including validation checks performed, steps executed, and a high-level process flow
How to Run this Workflow on page 32	Instructions for running this workflow in your environment
Sample Scenarios on page 34	Examples of typical parameter values 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 63 .

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 able to communicate with each other.
4. The master database system table (sysservers) 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 110.

How this Workflow Works

This topic contains the following information about the Dump Sybase Database workflow:

Overview

This workflow dumps the contents of a Sybase ASE database (the source database) into a file (the database dump file). It has the following capabilities:

- 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.

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 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.

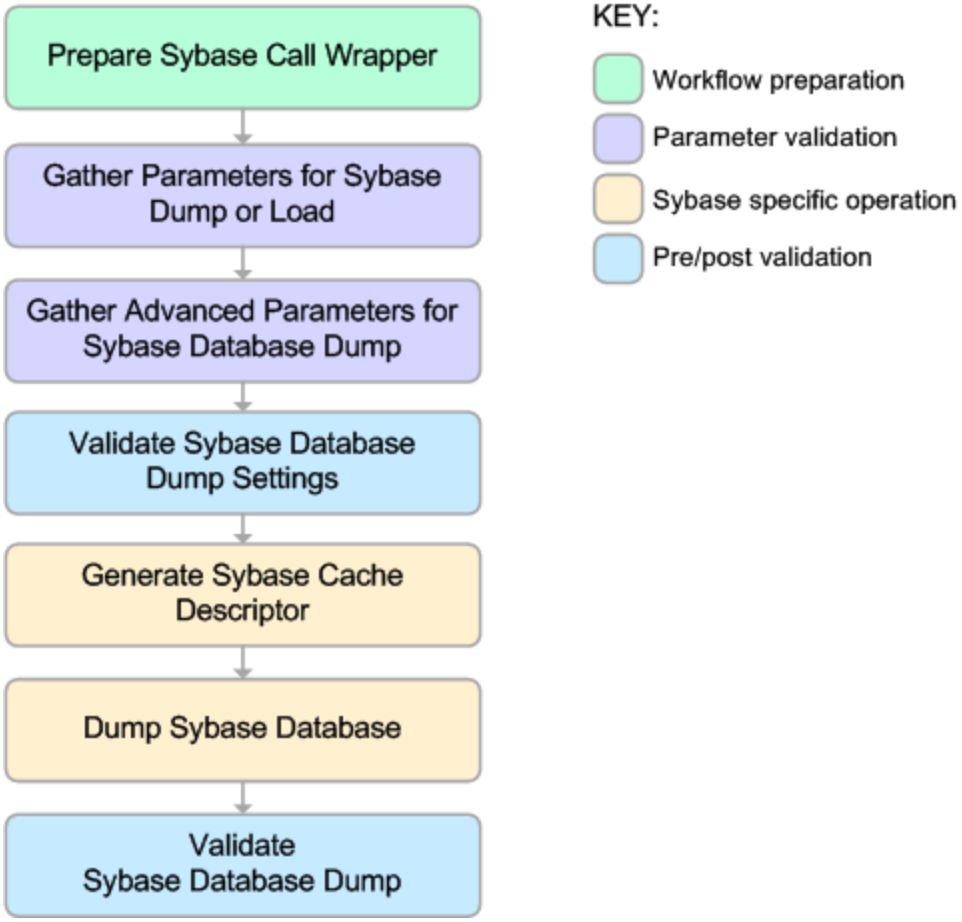

Process Flow

This workflow performs the following tasks:

1. Creates the Instance Wrapper and Server Wrapper (see Prepare Sybase Call Wrapper on page 79), and verifies that the HP Server Automation agent is able to communicate with the server where the workflow is running.
2. Performs the preliminary validation checks described above.
3. Generates the cache descriptor file for the source database. This is used to replicate the cache objects on the destination server.
4. Performs the database `dump` operation to create the database dump file.
5. 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: To learn the basic steps required to deploy and run any workflow, see the Quick Start Tutorial on page 9.

To customize and run the Dump Sybase Database workflow:

1. Create a deployable copy of the workflow (see Create a Deployable Workflow on page 10).
2. Determine the values that you will specify for the following parameters. This is the minimum set of parameters required to run this workflow.

Step Name	Parameter Name	Default Value	Description
Prepare Sybase Call Wrapper on page 79	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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
	Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Gather Advanced Parameters for Sybase Database Dump on page 88	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.

Step Name	Parameter Name	Default Value	Description
	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 63 for detailed descriptions of all input parameters for this workflow, including default values.

Note: To avoid entering passwords in clear text, see Using a Policy to Specify Parameter Values on page 112.

3. In the workflow editor, expose any additional parameters that you need (see How to Expose Additional Workflow Parameters on page 25). 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 on page 11 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 on page 12 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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Dump Sybase Database on page 63).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File Compression Level	8
	Dump File Password	MyPassword1@#

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Dump Sybase Database on page 63).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp`
	Dump File Password	MyPassword1@#

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Dump Sybase Database on page 63).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
	Dump File List	`/var/tmp/mytestdb.dmp`
Gather Advanced Parameters for Sybase Database Load on page 83	Cache File	`/var/tmp/runcache_ mytestdb.txt`

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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 on next page	List of prerequisites that must be satisfied before you can run this workflow
How this Workflow Works on page 39	Information about what the workflow does, including validation checks performed, steps executed, and a high-level process flow
How to Run this Workflow on page 43	Instructions for running this workflow in your environment
Sample Scenarios on page 45	Examples of typical parameter values 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 66.

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 able to communicate with each other.
6. The master database system table (sysservers) 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 110.

How this Workflow Works

This topic contains the following information about the Load Sybase Database Dump workflow:

Overview

This workflow provides the capability to refresh the specified target database from a database dump file that you provide. It has the following capabilities:

- 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.

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, 15.5, or 15.7	Sybase ASE15.0.3, 15.5, or 15.7
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 ASEusers. 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.

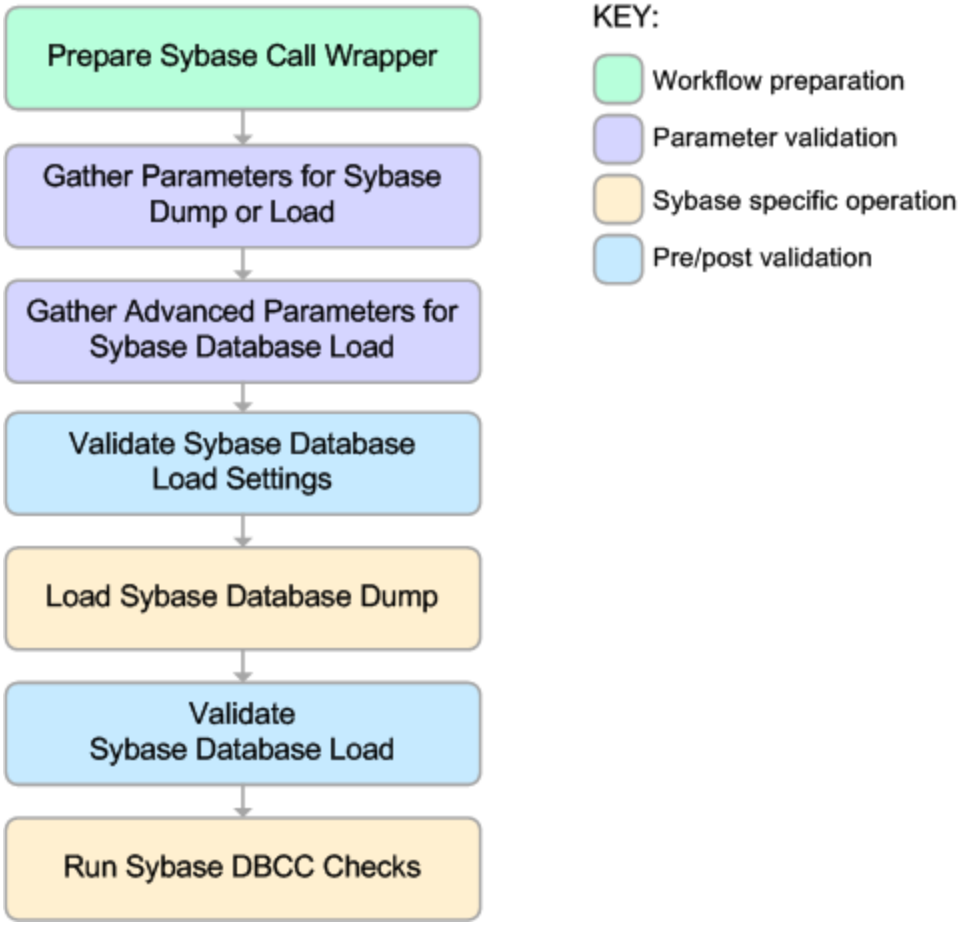

Process Flow

This workflow performs the following tasks:

1. Creates the Instance Wrapper and Server Wrapper (see Prepare Sybase Call Wrapper on page 79), and verifies that the HP Server Automation agent is able to communicate with the server where the workflow is running.
2. Performs the pre-refresh validation checks described above.
3. Reads the header of the database dump file, and validates that the sourceSybase ASE page size matches the target Sybase ASE page size. The workflow fails if the page sizes do not match.
4. Backs up any existing database users and groups.
5. Determines whether the source (the database dump file) and target servers have different byte architectures (big-endian versus little-endian).
6. Enables the database level "`dbo use only`" option to ensure that no users except the database owner are accessing the destination database.
7. Loads the database dump file on the target database server, and refreshes the destination database.
8. Brings the destination database online and performs the post-refresh checks.
9. Rebuild the indexes using theSybase ASE recommended `sp_post_xpload` system stored procedure.
10. 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 (see Run Sybase DBCC Checks on page 107).
11. 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: To learn the basic steps required to deploy and run any workflow, see the Quick Start Tutorial on page 9.

To customize and run the Load Sybase Database Dump workflow:

1. Create a deployable copy of the workflow (see Create a Deployable Workflow on page 10).
2. Determine the values that you will specify for the following parameters. This is the minimum set of parameters required to run this workflow.

Step Name	Parameter Name	Default Value	Description
Prepare Sybase Call Wrapper on page 79	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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
	Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/dump.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.

Step Name	Parameter Name	Default Value	Description
	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 66 for detailed descriptions of all input parameters for this workflow, including default values.

Note: To avoid entering passwords in clear text, see Using a Policy to Specify Parameter Values on page 112.

3. In the workflow editor, expose any additional parameters that you need (see How to Expose Additional Workflow Parameters on page 25). 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 on page 11 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 on page 12 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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb.dmp`

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb.dmp`
	Dump File Password	MyPassword1@#

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp`
	Dump File Password	MyPassword1@#

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb.dmp`
	Dump File Password	MyPassword1@#
	Cache Dump File	`/var/tmp/runcache_ mytestdb.txt`

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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 on page 50	List of prerequisites that must be satisfied before you can run this workflow
How this Workflow Works on page 52	Information about what the workflow does, including validation checks performed, steps executed, and a high-level process flow
How to Run this Workflow on page 56	Instructions for running this workflow in your environment
Sample Scenarios on page 59	Examples of typical parameter values 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 69.

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 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 able to communicate with each other.
7. The master database system table (sysservers) 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 110.

How this Workflow Works

This topic contains the following information about the Dump And Load Sybase Database workflow:

Overview

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.

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, 15.5, or 15.7	Sybase ASE15.0.3, 15.5, or 15.7
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 ASEusers. 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.

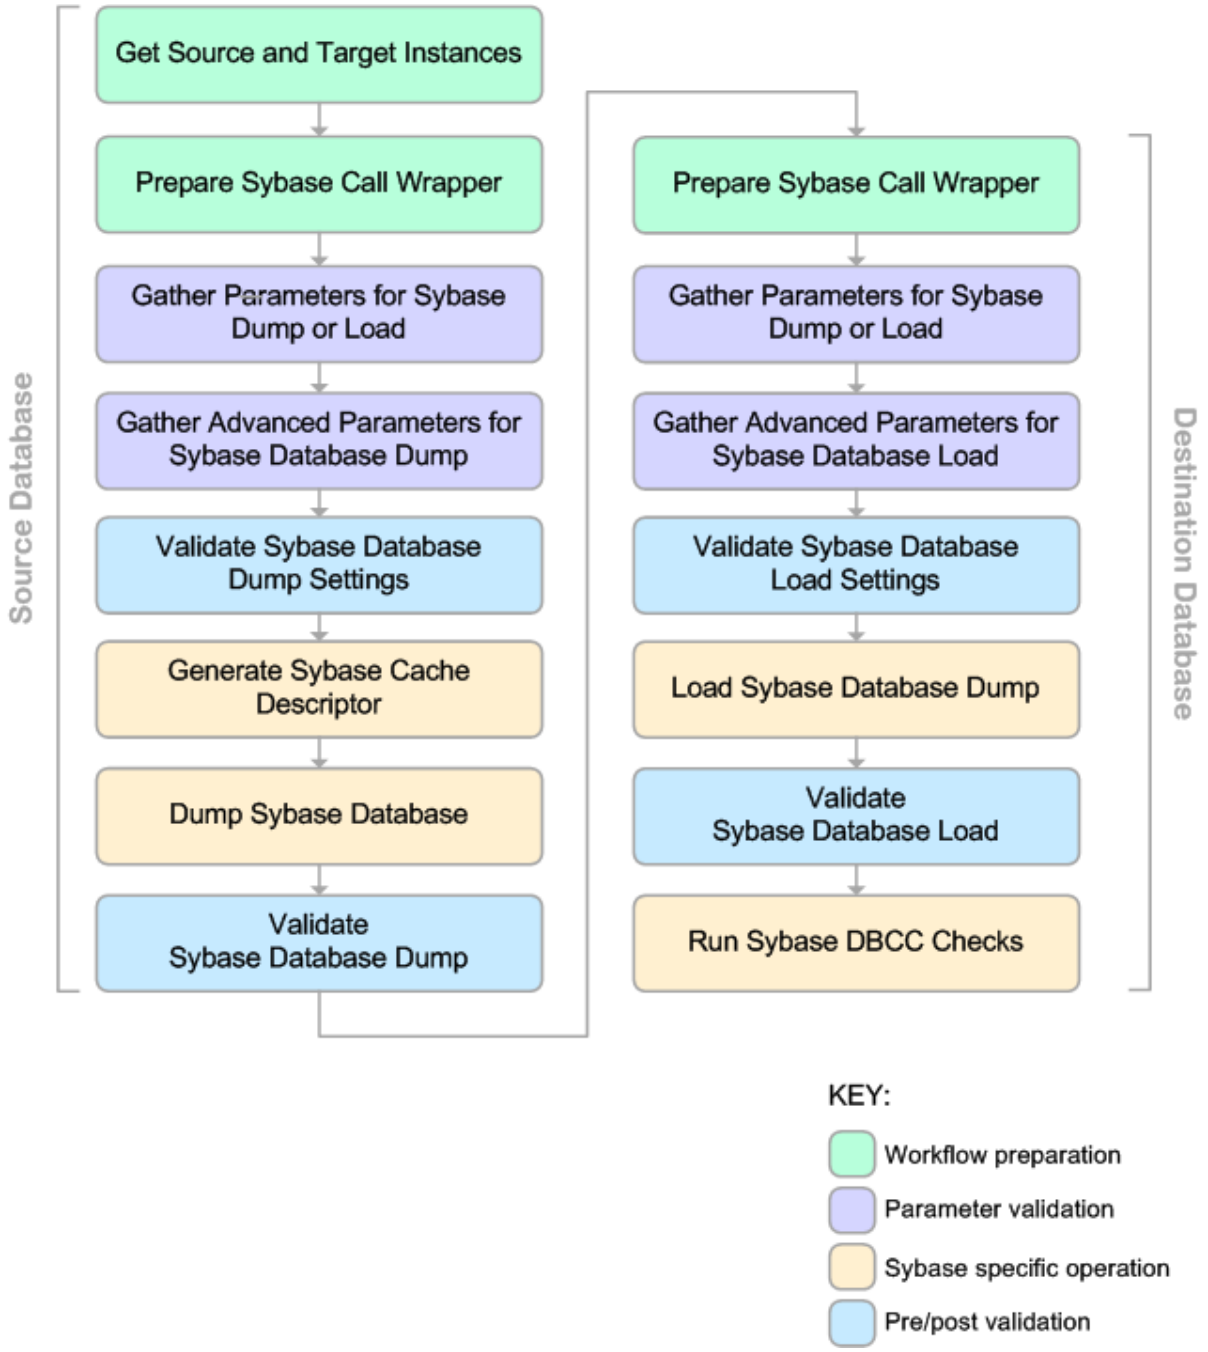

Process Flow

This workflow performs the following tasks:

1. Creates the Instance Wrapper and Server Wrapper (see Prepare Sybase Call Wrapper on page 79), and verifies that the HP Server Automation agent is able to communicate with the server where the workflow is running.
2. Performs the pre-dump validation checks described above.
3. Generates the cache descriptor file for the source database. This is used to replicate the cache objects on the destination server.
4. Performs the database `dump` operation to create the database dump file.
5. Performs post-dump validation checks to ensure that all required parameters had valid values.
6. Reads the header of the database dump file, and validates that the sourceSybase ASE page size matches the target Sybase ASE page size. The workflow fails if the page sizes do not match.
7. Backs up any existing database users and groups.
8. Determines whether the source and destination database servers have different byte architectures (big-endian versus little-endian).
9. Enables the database level "`dbo use only`" option to ensure that no users except the database owner are accessing the destination database.
10. Loads the database dump file on the destination database server, and refreshes the destination database.
11. Brings the destination database online and performs the post-refresh checks.
12. Rebuild the indexes using theSybase ASE recommended `sp_post_xpload` system stored procedure.
13. 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 (see Run Sybase DBCC Checks on page 107).
14. 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: To learn the basic steps required to deploy and run any workflow, see the Quick Start Tutorial on page 9.

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

1. Create a deployable copy of the workflow (see Create a Deployable Workflow on page 10).
2. Determine the values that you will specify for the following parameters. This is the minimum set of parameters required to run this workflow.

Step Name	Parameter Name	Default Value	Description
Prepare Sybase Call Wrapper on page 79	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: `sudo su - root /opt/op-sware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/op-sware/dma/jython.sh`).
	Sybase OS User Name	sybase	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Gather Advanced Parameters for Sybase Database Dump on page 88	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.

Step Name	Parameter Name	Default Value	Description
	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 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).
Gather Advanced Parameters for Sybase Database Load on page 83	Sybase ASE Home Directory	`/home/s-ybase/ASE_15`	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). 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 Gather Advanced Parameters for Sybase Database Dump on page 88 (see step 3).

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 66 for detailed descriptions of all input parameters for this workflow, including default values.

Note: To avoid entering passwords in clear text, see Using a Policy to Specify Parameter Values on page 112.

3. In the workflow editor, expose any additional parameters that you need (see How to Expose Additional Workflow Parameters on page 25). 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 on page 11 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 on page 12 for instructions).

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

Parameter Name	Default	Description
Source Instance	no default	The Adaptive Server instance where the dump file will be created. You specify this when you run the workflow. **Note:** The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.
Target Instance	no default	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.

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb.dmp`

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb.dmp`
	Dump File Compression Level	8
	Dump File Password	MyPassword1@#

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb1.dmp, /var/tmp/mytestdb2.dmp, /var/tmp/mytestdb3.dmp`
	Dump File Password	MyPassword1@#

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

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.

Step Name	Parameter Name	Example Value
Prepare Sybase Call Wrapper on page 79	Call Wrapper	jython
	Sybase OS User Name	sybase
Gather Advanced Parameters for Sybase Database Load on page 83	Dump File List	`/var/tmp/mytestdb.dmp`
	Dump File Password	MyPassword1@#
	Cache File	`/var/tmp/runcache_mytestdb.txt`

Be sure that the default values for all remaining parameters are appropriate for your environment (see Parameters for Load Sybase Database Dump on page 66).

Reference Information

This chapter contains the following information:

Workflow Parameters

Workflow Steps

Other Information

Chapter 3

Parameter Information

The following topics provide detailed information about the input parameters used by the workflows in this solution pack:

- Parameters for Dump Sybase Database on next page
- Parameters for Load Sybase Database Dump on page 66
- Parameters for Dump and Load Sybase Database on page 69

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 25). 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 on page 81
- Gather Advanced Parameters for Sybase Database Dump on page 88

For information about which steps use which input and output parameters, see How this Workflow Works on page 29.

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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
Dump Device 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
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	`/var/tmp/<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.	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 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).
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.

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 25). 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 on page 81
- Gather Advanced Parameters for Sybase Database Load on page 83

For information about which steps use which input and output parameters, see How this Workflow Works on page 39.

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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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 Load Sybase Database Dump Workflow (continued)

Parameter Name	Default Value	Required	Description
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.
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	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).
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 Load Sybase Database Dump Workflow (continued)

Parameter Name	Default Value	Required	Description
Remote Backup Instance Port	n/a	n/a	Not used in this release.
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).

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 25). 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 on page 81
- Gather Advanced Parameters for Sybase Database Dump on page 88

For information about which steps use which input and output parameters, see How this Workflow Works on page 29.

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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).

Input Parameters for the Dump Sybase Database Workflow (continued)

Parameter Name	Default Value	Required	Description
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.
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	`/var/tmp/<`*`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.	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. **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).
Dump File Path	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 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.
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. **Note:** The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.
Source Database Name	mytestdb	required	Name of database from which the dump file (or files) will be created.
Source Instance	no default	required	The Adaptive Server instance where the dump file will be created. You specify this when you run the workflow. **Note:** The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.

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 4

Step Information

The following topics provide detailed information about the steps used by the workflows in this solution pack:

Steps for Dump Sybase Database

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.

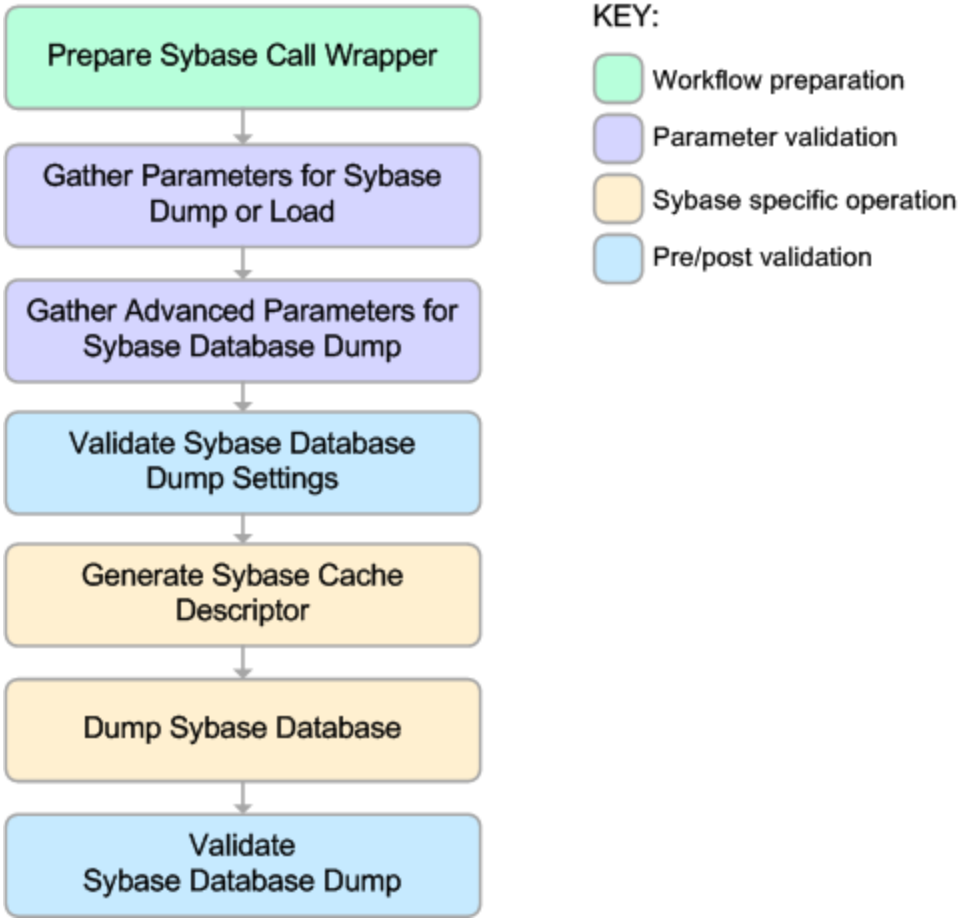

For parameter descriptions and defaults, see Parameters for Dump Sybase Database on page 63.

Steps for Load Sybase Database Dump

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.

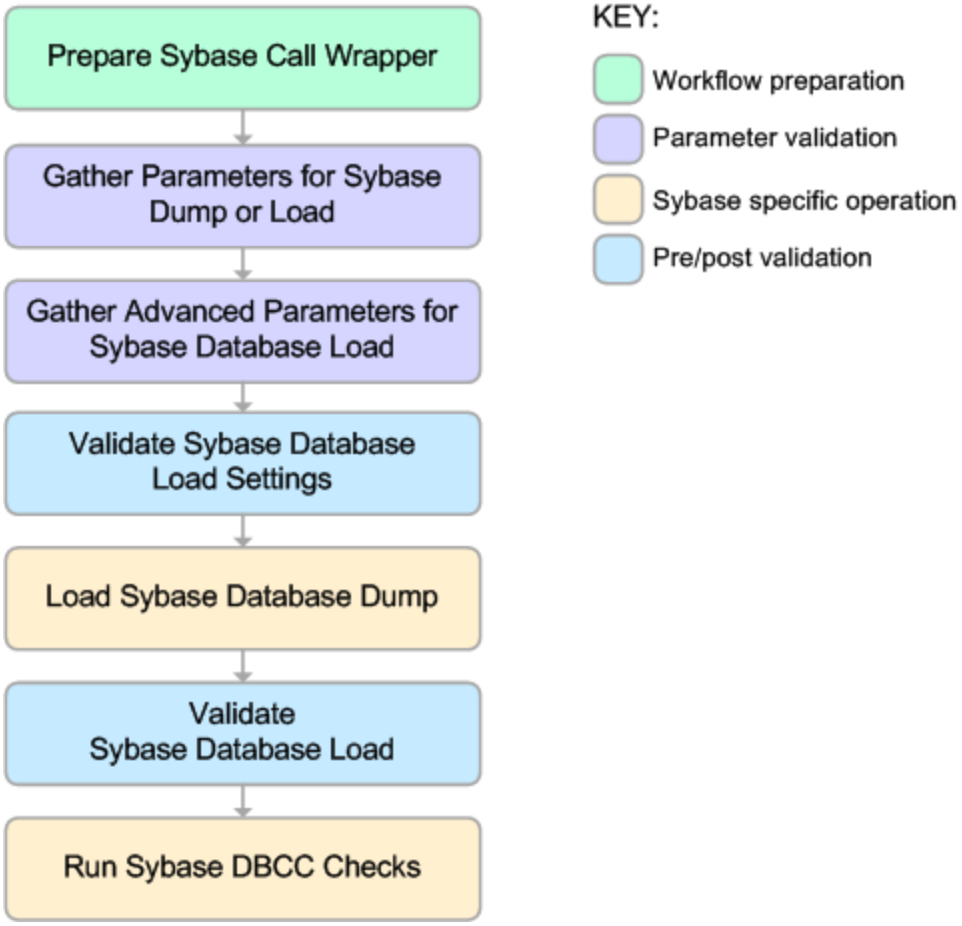

For parameter descriptions and defaults, see Parameters for Load Sybase Database Dump on page 66.

Steps for Dump and Load Sybase Database

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.

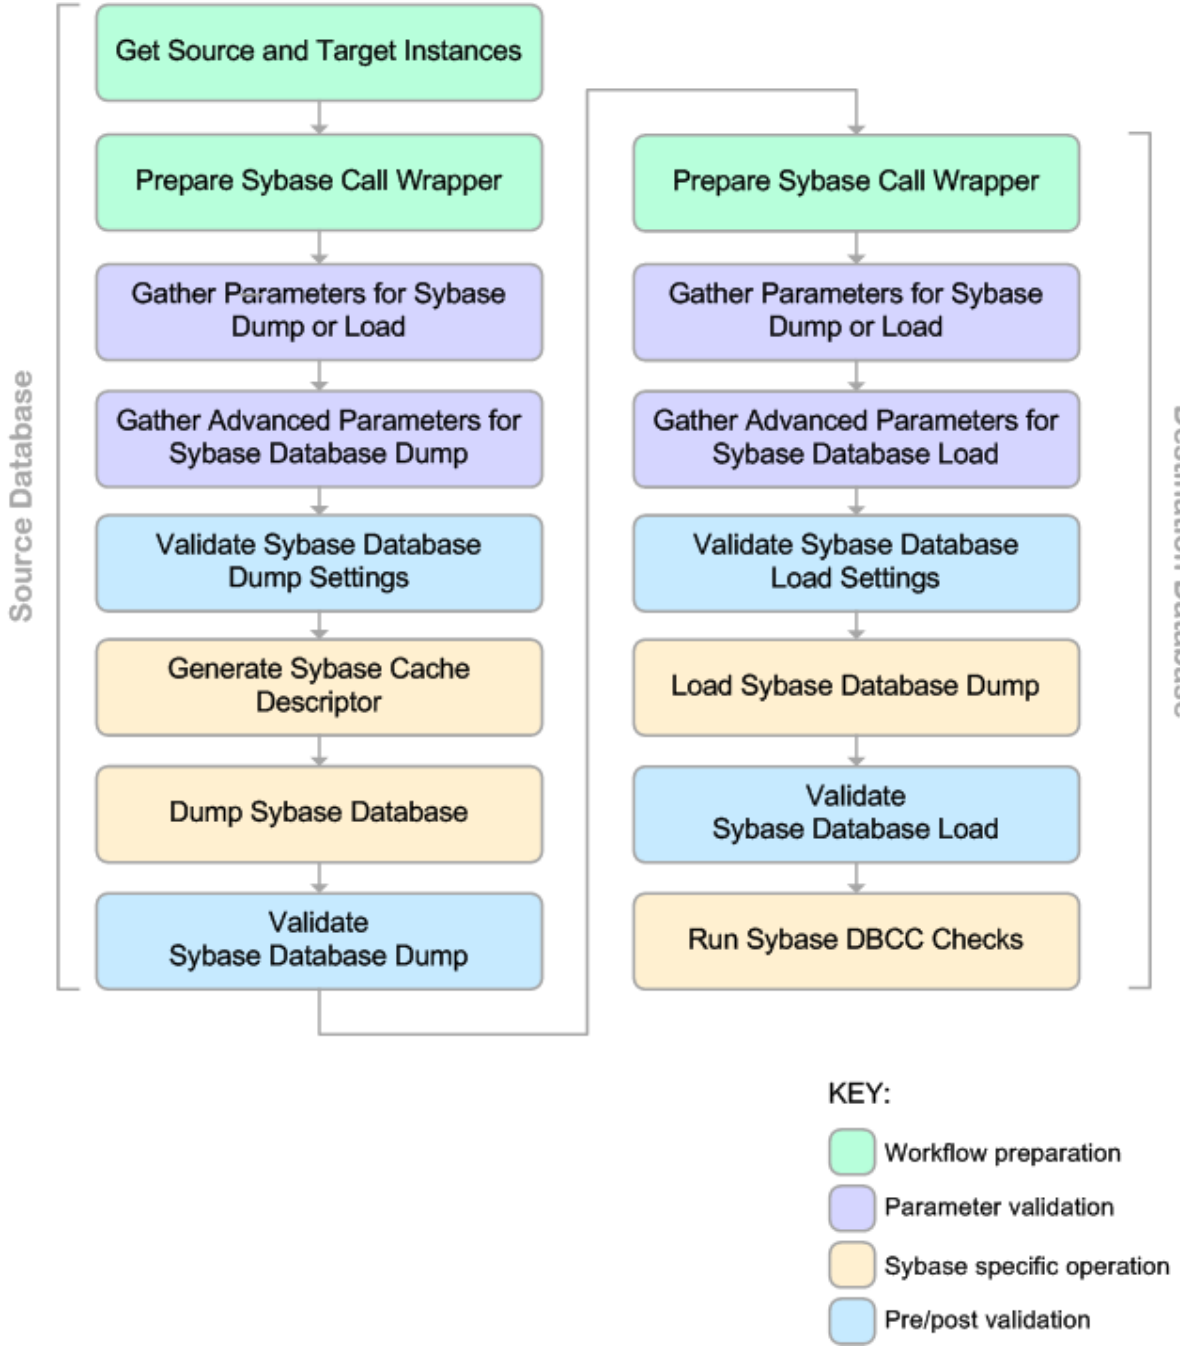

For parameter descriptions and defaults, see Parameters for Dump and Load Sybase Database on page 69.

All Sybase ASE Database Refresh Steps

The following steps are used by the workflows in this solution pack:

Database Dump Steps

- Prepare Sybase Call Wrapper on page 79
- Gather Parameters for Sybase Dump or Load on page 81
- Gather Advanced Parameters for Sybase Database Dump on page 88
- Validate Sybase Database Dump Settings on page 94
- Generate Sybase Cache Descriptor on page 96
- Dump Sybase Database on page 101
- Validate Sybase Database Dump on page 105

Database Load Steps

- Prepare Sybase Call Wrapper on page 79
- Gather Parameters for Sybase Dump or Load on page 81
- Gather Advanced Parameters for Sybase Database Load on page 83
- Validate Sybase Database Load Settings on page 92
- Load Sybase Database Dump on page 98
- Validate Sybase Database Load on page 103
- Run Sybase DBCC Checks on page 107

Database Dump and Load Steps

- Get Sybase Source and Target Instances on page 109
- Prepare Sybase Call Wrapper on page 79
- Gather Parameters for Sybase Dump or Load on page 81
- Gather Advanced Parameters for Sybase Database Dump on page 88
- Validate Sybase Database Dump Settings on page 94
- Generate Sybase Cache Descriptor on page 96
- Dump Sybase Database on page 101
- Validate Sybase Database Dump on page 105
- Prepare Sybase Call Wrapper on page 79
- Gather Parameters for Sybase Dump or Load on page 81
- Gather Advanced Parameters for Sybase Database Load on page 83
- Validate Sybase Database Load Settings on page 92
- Load Sybase Database Dump on page 98

- Validate Sybase Database Load on page 103
- Run Sybase DBCC Checks on page 107

Prepare Sybase Call Wrapper

Purpose

This step constructs the command that will be used to execute subsequent workflow steps as either the operating system administrative user (typically, root) or as the owner of the Sybase ASEinstance (typically, sybase).

Input Parameters

Parameter Name	Default Value	Required	Description
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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
Sybase OS User Name	sybase	required	OS user (typically, sybase) who owns the Sybase ASE installation directory.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

The following parameters are made available to subsequent steps in the workflow:

Parameter Name	Description
Instance Wrapper	Command that will be used to execute subsequent steps in the workflow as the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
Server Wrapper	Command that will be used to execute subsequent steps in the workflow as the OS administrative user (for example: `sudo su - root /opt/opsware/dma/jython.sh`).
Sybase OS User Name	OS user (typically, sybase) who owns the Sybase ASE installation directory.

Return Codes

0 = Call wrappers were successfully created.

1 = Either or both call wrappers could not be created.

Used By Workflows

Gather Parameters for Sybase Dump or Load

Purpose

This step gathers the input parameters required for a basic database refresh. For more customized refresh scenarios, see Gather Advanced Parameters for Sybase Database Load on page 83.

Input Parameters

Parameter Name	Default Value	Required	Description
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.
ASE SysAdmin Password	password	required	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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).

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

The following parameters are made available to subsequent steps in the workflow:

Parameter Name	Description
ASE SysAdmin Password	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
ASE SysAdmin Username	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database load operation.

Parameter Name	Description
Sybase ASE Home Directory	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).

Return Codes

0 = No errors occurred during the parsing of the input parameters.

1 = One or more parameter values could not be read.

Used By Workflows

- Dump Sybase Database on page 26
- Load Sybase Database Dump on page 36
- Dump And Load Sybase Database on page 48

Gather Advanced Parameters for Sybase Database Load

Purpose

This step gathers the parameters required for a customized database load scenario. Any parameters that are hidden (not explicitly exposed) in the deployment are assigned default values.

The basic parameters are gathered in the Gather Parameters for Sybase Dump or Load on page 81 step.

Input Parameters

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	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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.
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).

Parameter Name	Default Value	Required	Description
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	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).
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.

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).
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.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

Parameter Name	Description
ASE SysAdmin Password	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
ASE SysAdmin Username	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database load operation.
Cache Dump File	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.
DBCC Checks	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	The directory (with absolute path) where you want to store the DBCC results (output files) for post-load checks. This directory must exist.
Database Instance Name	The name of the Sybase ASE instance where the database will be loaded from the dump file (or files).

Parameter Name	Description
Dump Device Name	Not used in this release.
Dump File List	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 Operating System	The value of this parameter will be derived from the database dump file header to determine whether this is a cross-platform (byte-swapped architecture) refresh. **Caution:** Do not unmap this parameter—its value is derived by the workflow.
Dump File Password	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).
Dump File Path	Not used in this release.
Local Backup Instance Name	Not used in this release.
Local Backup Instance Port	Not used in this release.
Remote Backup Instance Name	Not used in this release.
Remote Backup Instance Port	Not used in this release.
Sybase ASE Home Directory	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).

Parameter Name	Description
Target Database Instance Name	Name of the Sybase ASE instance where the dump file (or files) will be loaded.
Target Database Name	Name of the database where the dump file (or files) will be loaded.
Target Database Page Size	Page size of the target database server (in kilobytes).

All of the input parameters with the exception of Call Wrapper are also output parameters.

Return Codes

0 = No errors occurred during the parsing of the input parameters and assignment of default values.

1 = One or more parameter values could not be read.

Used By Workflows

- Load Sybase Database Dump on page 36
- Dump And Load Sybase Database on page 48

Gather Advanced Parameters for Sybase Database Dump

Purpose

This step gathers the parameters required for a customized database dump scenario. Any parameters that are hidden (not explicitly exposed) in the deployment are assigned default values.

The basic parameters are gathered in the Gather Parameters for Sybase Dump or Load on page 81 step.

Input Parameters

Parameter Name	Default Value	Required	Description
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: `sudo su - root /opt/opsware/dma/jython.sh`) or the owner of the Sybase ASE installation (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
ASE SysAdmin Username	sa	required	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database dump operation.
ASE SysAdmin Password	password	required	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
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).

Parameter Name	Default Value	Required	Description
Dump File Password	no default	optional	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).
Dump File List	`/var/tmp/<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.	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.
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.
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.
Dump File Compression Level	7	optional	Compression level (1-9) to apply to the dump file (or files) that will be created.

Output Parameters

Parameter Name	Description
ASE SysAdmin Username	The Sybase ASE user who can perform all administrative operations (typically sa). This user will perform the database dump operation.
ASE SysAdmin Password	Password for the Sybase ASE user specified in the ASE SysAdmin Username parameter.
Sybase ASE Home Directory	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).
Dump File Password	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).
Dump File List	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	Compression level (1-9) to apply to the dump file (or files) that will be created.
Source Database Instance Name	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	Name of database from which the dump file (or files) will be created.
Cache File	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.

Return Codes

0 = No errors occurred during the execution of this step.

1 = One or more errors occurred.

Used By Workflows

- Dump Sybase Database on page 26
- Dump And Load Sybase Database on page 48

Validate Sybase Database Load Settings

Purpose

This step performs the following validation checks prior to the database load operation:

- Parameters are present and set correctly.
- Target database container exists.
- Target database is online.
- Dump file and target database page sizes match.

Input Parameters

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.
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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.
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).
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.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

The following parameters are made available to subsequent steps in the workflow:

Parameter Name	Description
Dump File Path	Not used in this release.

Return Codes

0 = All parameter values specified are valid.

1 = One or more parameter values are invalid.

Used By Workflows

- Load Sybase Database Dump on page 36
- Dump And Load Sybase Database on page 48

Validate Sybase Database Dump Settings

Purpose

This step performs the following validation checks prior to the database dump operation:

- Parameters are present and set correctly.
- Source database is online.
- Source database backup server is online.
- Sufficient space is available to create the database dump.

Input Parameters

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 operation.
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
Dump File List	`/var/tmp/<`*`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.	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.

Parameter Name	Default Value	Required	Description
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).
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.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

There are no output parameters for this step.

Return Codes

0 = All parameter values specified are valid.

1 = One or more parameter values are invalid.

Used By Workflows

- Dump Sybase Database on page 26
- Dump And Load Sybase Database on page 48

Generate Sybase Cache Descriptor

Purpose

This step generates the cache descriptor file for the source database. This is used to replicate the cache objects on the destination database server.

Input Parameters

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 operation.
Cache File	no default	optional	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).
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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).
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.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

There are no output parameters for this step.

Return Codes

0 = All parameter values specified are valid.

1 = One or more parameter values are invalid.

Used By Workflows

- Dump Sybase Database on page 26
- Dump And Load Sybase Database on page 48

Load Sybase Database Dump

Purpose

This step performs the Sybase ASE database load operation.

Where applicable, this step will also refresh the database cache, backup and restore target database users, and rebuild database indices.

Note: If a Cache Dump File is used, that file must be created such that each row is on one line. For example:

Cache Name	Config Size	Run Size	Overhead
----------	----------	----------	----------
default data cache	0.00 Mb	8.00 Mb	0.94 Mb
test_cache	3.00 Mb	3.00 Mb	0.41 Mb

This can be accomplished by using the "`-w 250`" option to the `isql` command.

Input Parameters

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	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).

Parameter Name	Default Value	Required	Description
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.
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 Operating System	no default	optional	The value of this parameter will be derived from the database dump file header to determine whether this is a cross-platform (byte-swapped architecture) refresh. **Caution:** Do not unmap this parameter—its value is derived by the workflow.
Dump File Password	no default	optional	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).
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).
Target Database Instance Name	NY_DS	required	Name of the Sybase ASE instance where the dump file (or files) will be loaded.

Parameter Name	Default Value	Required	Description
Target Database Name	mytestdb	required	Name of the database where the dump file (or files) will be loaded.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

This step has no output parameters.

Return Codes

0 = The specified database was successfully refreshed, and all necessary post-refresh tasks were successfully completed.

1 = Either the refresh operation or one of the post-refresh steps failed.

Used By Workflows

- Load Sybase Database Dump on page 36
- Dump And Load Sybase Database on page 48

Dump Sybase Database

Purpose

This step performs the Sybase ASE database dump operation.

Input Parameters show

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.
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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	`/var/tmp/<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.	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.

Parameter Name	Default Value	Required	Description
Dump File Password	no default	optional	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).
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).
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.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

This step has no output parameters.

Return Codes

0 = The dump file was successfully created.

1 = The dump operation failed.

Used By Workflows

- Dump Sybase Database on page 26
- Dump And Load Sybase Database on page 48

Validate Sybase Database Load

Purpose

This step performs the following post-load checks:

- The destination database is online.
- No database objects were corrupted.
- User-defined database tables exist in the restored database.

Input Parameters

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.
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
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).
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.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

This step has no output parameters.

Return Codes

0 = All post-refresh validation checks passed.

1 = One or more checks failed.

Used By Workflows

- Load Sybase Database Dump on page 36
- Dump And Load Sybase Database on page 48

Validate Sybase Database Dump

Purpose

This step determines whether all parameters required for the dump operation have values.

Input Parameters

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 operation.
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
Dump File List	`/var/tmp/<`*`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.	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.
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).
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.

Parameter Name	Default Value	Required	Description
Source Database Name	mytestdb	required	Name of database from which the dump file (or files) will be created.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

This step has no output parameters.

Return Codes

0 = All required parameters have values.

1 = One or more required parameters has not been specified.

Used By Workflows

- Dump Sybase Database on page 26
- Dump And Load Sybase Database on page 48

Run Sybase DBCC Checks

Purpose

This step runs the specified database consistency checker (DBCC) checks against the restored database.

Input Parameters

Parameter Name	Default Value	Required	Description
Call Wrapper	no default	required	Command that will execute the specified step as the owner of the Sybase ASE software (for example: `sudo su - sybase /opt/opsware/dma/jython.sh`).
Database Server Name	NY_DS	required	Name of the Sybase ASEinstance where the database that was loaded from the dump file resides.
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. If no DBBC checks are specified, this step doe not do anything.
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.
Role Password SQL Statement	no default	required	Not used in this release.
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 User Name	sybase	required	OS user (typically, sybase) who owns the Sybase ASE installation directory.
Sybase User Password	sybase123	required	Password for the OS user (typically, sybase) who owns the Sybase ASE installation directory.

For detailed descriptions of all input parameters used by this workflow, see the Reference Information on page 61.

Output Parameters

This step has no output parameters.

Return Codes

0 = The specified DBCC checks were successfully performed. Note: This does not indicate whether the checks passed or failed. You must look at the files in the DBCC Error Directory to see the results of each check.

1 = One or more of the specified DBCC checks could not be performed.

Used By Workflows

- Load Sybase Database Dump on page 36
- Dump And Load Sybase Database on page 48

Get Sybase Source and Target Instances

Purpose

This step collects the source and target instance information specified at run time and stores this information in parameters that are consumed by subsequent steps in the workflow.

This step is used in bridged execution workflows, which are supported on HP Server Automation version 9.11 (and later).

Input Parameters

Parameter Name	Default Value	Required	Description
Source Instance	no default	required	The Adaptive Server instance where the dump file will be created. You specify this when you run the workflow. **Note:** The Source Instance that you specify at run time must match the Source Database Instance Name that you specify in the deployment.
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.

Output Parametersshow

Parameter Name	Description
Source Instance	See the Input Parameters.
Target Instance	See the Input Parameters.

Return Codes

0 = No errors occurred during the execution of this step.

1 = One or more errors occurred.

Used By Workflows

- Dump And Load Sybase Database on page 48

Chapter 5

Other Reference Information

The following topics provide additional information pertinent to the workflows in this solution pack:

- Sybase ASE Product Documentation below
- Using this Solution Pack With HP Server Automation on next page

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

Using this Solution Pack With HP Server Automation

HP Database and Middleware Automation (HP DMA) version 1.00 is compatible with HP Server Automation version 9.02 (and later 9.0x versions).

For information about running HP DMA workflows from HP Server Automation versions prior to 9.10, refer to the following documents:

- *HP Server Automation Application Deployment User Guide* (version 9.02 and later 9.0x versions)
- *HP Database and Middleware Automation User Guide* (version 1.00)

HP Database and Middleware Automation version 9.10 is compatible with HP Server Automation version 9.10 (and later).

For information about running HP Database and Middleware Automation workflows from HP Server Automation version 9.10 (and later), refer to the following documents:

- *User Guide: Application Deployment Manager*
- *User Guide: Database and Middleware Automation User Guide*

These guides are included in the HP Server Automation documentation library (version 9.10 and later).

Chapter 6

Tips and Best Practices

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

Using a Policy to Specify Parameter Values below

Using 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:

- Passwords are obfuscated (not displayed in clear text).
- The policy can be used in any deployment.
- It is faster and less error-prone than specifying parameter values manually.

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.

If you are using HP Server Automation, see the *User Guide: Database and Middleware Automation*. This guide is included in the HP Server Automation documentation library (SA version 9.10 and later).

If you are using HP DMA 1.00, see "Policies" in the *HP Database and Middleware Automation User Guide* for more information.

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 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-delimited list of values (or a large amount of text not suitable for a Text attribute).
- A Password attribute contains simple text, but it is obfuscated 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.

> **Note:** Extracted policies only use Text type attributes. Therefore, passwords are not obfuscated when you specify them in an extracted policy. You can, however, delete an automatically extracted attribute and then add a new one of type Password.

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 *User Guide: Database and Middleware Automation* for details. This guide is included in the HP Server Automation documentation library (SA version 9.10 and later).

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 text box to the right of the parameter name, type the first few characters of the policy name.

      A drop-down list of policy attributes appears.

   b. From the drop-down list, select the attribute that you want to reference.

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

Chapter 7

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 Server Automation on next page

For additional information, refer to the “Troubleshooting” chapter in the *HP Server Automation User Guide: Database and Middleware Automation*.

If you are using HP Database and Middleware Automation version 1.00, see the *HP Database and Middleware Automation Installation Guide*.

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 (Read or Write) permissions for organizations, workflows, steps, policies, and deployments. Deployments have an extra permission: Execute. Users are assigned to roles, and they gain access to these items according to the permissions defined for their roles.

Note: The following information pertains only to HP DMA 1.00:

Roles can be defined in one of two ways: native or LDAP groups.

- Native roles define groups of HP DMA users in the repository.
- LDAP groups are retrieved from the LDAP server configured in the Setup > Expert Engine area. No user information is stored in the repository for LDAP groups. This allows you to use your corporate directory for defining users and their permissions making security audits easier.

Roles are assigned on the Roles tab of the Setup page. See "Roles" in the *HP Database and Middleware Automation User Guide* (version 1.00) for more information.

Make sure that the HP DMA users in your environment are assigned roles that grant them the permissions they need to accomplish their tasks. For example:

- To view a workflow, your role must have Read permission for that workflow.
- To view a deployment, your role must have Read permission for that deployment.

- To edit a workflow, your role must have Write permission for that workflow.
- To run a deployment, your role must have Execute permission for that deployment.

Permissions determine what features and functions are available and active in the HP DMA UI. For a detailed breakdown, see the HP Database and Middleware Automation *User Guide*.

Note: In HP Server Automation, roles and permissions work differently. Both roles and permissions are assigned by the SA administrator. See the HP Server Automation *Administration Guide* and the *User Guide: Database and Middleware Automation* for more information. Both guides are included in the HP Server Automation documentation library (SA version 9.10 and later).

Discovery in HP Server Automation

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

In HP DMAversion 1.00, discovery is automatically activated when an agent is started on a target machine.

In HP Server Automation, you must explicitly initiate the process of discovery—it is not automatic. Refer to the *User Guide: Database and Middleware Automation* for instructions. This guide is included in the SA documentation library (version 9.10 and later).

Glossary

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 Server Automation version 9.11 (and later).

C

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.

D

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 databse.

DESTINATION

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

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 it's 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

Oracle 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.

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.

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.

S

source

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

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.