

HP Database and Middleware Automation

For Red Hat Enterprise Linux and SUSE Enterprise Linux

Software Version: 10.22

Concepts Guide

Document Release Date: December 2014

Software Release Date: December 2014



Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notice

© Copyright 2014 Hewlett-Packard Development Company, L.P.

Trademark Notices

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

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

Documentation Updates

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

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

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

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

Document Changes

Chapter	Version	Changes
Title Page Legal Notices	10.22	Updated version number, software release date, document release date, and copyright date range.

Support

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

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

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

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

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

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

To find more information about access levels, go to:

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

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

Contents


Contents	3
Audience	5
Document Map	6
Chapter 1: Overview	7
Chapter 2: HP DMA Objects	9
Automation Objects	9
Workflows	9
Steps	12
Functions	14
Solution Packs	15
Environment Objects	16
Organizations	16
Servers	17
Instances	17
Databases	19
Custom Fields	21
Smart Groups	22
Bridge Objects	24
Deployments	24
Policies	27
Deployment Runs	28
Connector Object	29
Connector	29
Chapter 3: Platform Architecture	30
HP DMA Server	32
HP Server Automation Server	33
Oracle Database	34
Managed Servers	34
Chapter 4: How the HP DMA Platform Works	35

Typical flow of HP DMA workflow execution	36
Workflow Execution Script (WEST)	38
Workflow Execution Engine	39
Custom Fields used with WEST	40
REST APIs	41
DMA Baseline Data	42
Automation Platform Extension (APX)	43
Chapter 5: Permissions Model	44
Roles	45
Capabilities	46
HP DMA User	47
SA User Groups for HP DMA	47
Special User: HP DMA Initial Administrator	48
Special User: HP DMA Connector User	48
Permissions	49
Permissions for Automation Objects	49
Permissions for Environment Objects	50
Permissions for the Bridge Objects	50
Additional Information	51
Chapter 6: Special Types of Workflows	52
Bridged Execution Workflows	52
Master Workflows	56
Appendix A: HP Software Documentation	61
Glossary	63

Audience

This guide is intended for anyone who wants to become more familiar with HP Database and Middleware Automation (HP DMA): installers, administrators, workflow creators, workflow runners, and others.

Since many HP DMA concepts build upon other concepts, this guide contains many references to other sections within this guide and to other documents.

The  icon indicates tips of special interest to those who want to dig deeper into HP DMA—so that you may better understand what is going on behind the scenes or how to develop your own HP DMA content.

Document Map

The following table shows you how to navigate this guide:

Topic	Description
Overview	Describes HP DMA at a high level, including the foundational HP DMA Platform and Content.
HP DMA Objects	Describes the different HP DMA objects: Automation objects, Environment objects, Bridge objects, and the Connector object.
Platform Architecture	Provides a diagram of the HP DMA architecture and describes the different elements: the HP DMA server, the HP Server Automation server, the back-end database, and managed servers.
How the HP DMA Platform Works	Describes the behind-the-scenes HP DMA processes at a high level.
Permissions Model	Describes the HP DMA permissions model.
Special Types of Workflows	Describes two special types of HP DMA workflows (Bridged Execution Workflows and Master Workflows), explaining how they work and providing examples.
HP Software Documentation	Provides information to obtain HP DMA and SA documentation.

Chapter 1: Overview

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 administration tasks that take up 60-70% of a database or application server administrator's day. Automating these tasks enables greater efficiency and faster change delivery with higher quality and better predictability. HP DMA provides role-based access to automation content. This enables you to better utilize resources at every level:

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

An HP DMA workflow performs a specific automated task—such as provisioning databases or application servers, patching databases 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 HP DMA parameters. Related HP DMA workflows are grouped together in solution packs.

HP DMA Foundational Concepts

The HP DMA Platform and HP DMA Content are the most basic, fundamental concepts within HP DMA.

HP DMA Platform

The platform is the foundation on which HP DMA runs.

Think of the platform as the operating system of your smart phone. You do not need to know all of the internals of the platform to use the smart phone but the operating system is essential for your smart phone to work.

The automation platform consists of a workflow engine, the server/instance/database environment, logging of the executed automation, and Role-based access.

HP DMA Content

The content runs on top of the platform and is responsible for the automation.

Think of the content as applications on your smart phone. On top of the phone's operating system you add the applications. You can add or remove a variety of applications. You are not required to install any particular application. Yet the applications are what make the smart phone useful and fun.

Chapter 2: HP DMA Objects

This chapter describes the HP DMA objects. Objects are basic to understanding HP DMA.

The objects fall into these categories:

- [Automation Objects](#)
- [Environment Objects](#)
- [Bridge Objects](#)
- [Connector Object](#)

Automation Objects

The umbrella term [automation objects](#) refers to those items to which role-based permissions (see [Permissions Model](#)) can be assigned. The HP DMA automation objects include:

- [Workflows](#)
- [Steps](#)
- [Functions](#)
- [Parameters](#)
- [Solution Packs](#)

Workflows


In HP DMA, a [workflow](#) executes a process —such as installing a software product or checking a database instance for compliance with a specific security benchmark.

The workflow is the primary automation object of HP DMA, and automates the process followed for an operational procedure. Workflows contain steps (see [Steps](#)), which are linked together to form business logic for a common task. Workflows connect existing tasks to perform a new business process built on existing best practices and processes.

Workflow Steps

A workflow consists of a set of steps and the paths that should be taken between the steps. Each step returns an exit code that determines the next step (or steps) to run. If there are multiple follow-on steps, all but one of the steps should specify the exit code required to run that step. If the exit code does not match any of the follow-on steps, the workflow will run the follow-on step without an exit code specification.

For example, a workflow contains steps A, B, and C (plus other steps not shown). The workflow specifies that, after Step A, the next steps are either B or C. If step A's exit code is 0, then Step B will execute next. If Step A's exit code is anything other than 0, Step C will execute next.



If there are two or more follow-on steps that require the same exit code from a previous step, then all of those steps will be executed..

Parameter mappings

The workflow also contains parameter **mappings** (see [Parameters](#)) between the steps.

A step's input parameters can be mapped to any of the following:

- User selected—the user specifies a value for the parameter in the deployment

Step A 2, 3

Parameter 1:

Parameter 2:

- The output parameter of a previous step

Step B

Parameter 1:

Parameter 2:

- An HP DMA built-in

Step C

User:

- A previously-defined policy parameter (see [Policies](#)) when the workflow is deployed

Step C

Password:

User:

- A previously-defined Custom Field (see [Custom Fields](#)) when the workflow is deployed


Step C

Operating System:

There are no output parameters for workflows, only for steps.

Success/Failure

Workflows are marked as Success if the last step executed is the Success step. Similarly, they are marked as Failure if the last step executed is the Failure step.



Source of workflows

Workflows are supplied in solution packs (see [Solution Packs](#)), but the user must copy the workflows before using them. HP DMA requires the user to copy the workflows to reduce the impact of future solution pack updates.

On the History tab you can view the installation history for HP DMA-supplied workflows and the change history for custom workflows.

Versioning

In the course of HP DMA development, improvements to workflows could cause customers' existing deployments and customizations to break (for example, by changing steps, parameters, or mappings). In this case, HP DMA releases the upgraded workflow with a version appended to the name. For example: Workflow Name v2.



Development tips

In HP DMA, users with the Workflow Creator or Administrator capability (see [Capabilities](#)) can copy or create workflows.

When you write your own workflows, ensure that the workflows will execute a single Success or Failure step.

Steps

A workflow consists of a sequence of **steps**. Each step performs a very specific task. Steps can be shared among workflows. For example, the step Download Software is used by many database and middleware workflows; and the step Discover Oracle Databases is used by many Oracle workflows.

All HP DMA users can view steps.

Steps are reusable scripts that contain the actual code used to perform a unit of work detailed in a workflow. The scripts may be in Jython or any other scripting language available on the target server.

Parameters

Steps can have input and output **parameters**, whose values will be unique to your environment.

Parameters are pieces of information—such as a file name or a user name—that a step requires to carry out its action.

Parameter values can be set in multiple ways:

- The step's parameter definition provides a default value.
- The step code assigns a value, all the time or only if the parameter does not already have a value.
- The user specifies values for **User Selected** workflow parameters in the deployment (see [Deployments](#)).
- The user specifies values for **Enter at Runtime** deployment parameters on the target system when the workflow is initiated.

For additional information see [Parameter mappings](#).

HP DMA workflows validate many parameter values—usually in a validate step—to verify that the values are acceptable.

If you provide valid values for the input parameters that the scenario requires, then the workflow will be able to accomplish its objective. Output parameters from one step often serve as input parameters to another step—this parameter mapping occurs at the workflow level (see [Parameter mappings](#)).

Source of steps

Steps are supplied as part of a solution pack (see [Solution Packs](#)). The steps that are delivered are locked and cannot be modified. You can reuse these steps in workflows that you customize.

On the History tab you can view the installation history for HP DMA-supplied steps and the change history for custom steps.



Development tips

To make changes to a step that is delivered as part of a solution pack, you must make a copy of the step, and then make the desired changes. Modifying a step that is not part of a solution pack is controlled by Role-Based Access Control (RBAC) (see [Roles](#)).

Anyone with the DMA User role can create steps in HP DMA.

Versioning

In the course of HP DMA development, improvements to steps could cause cutomers' existing deployments and customizations to break (for example, by adding or removing parameters). In this case, HP DMA releases the upgraded step with a version appended to the name. For example: Step Name v2.

Functions

Functions are reusable pieces of code and are grouped into custom content libraries. A function can be imported into both steps (see [Steps](#)) and other functions. HP DMA-supplied functions are frequently used by many steps. Any common routine or operation performed in multiple steps is a good candidate for a function. Functions may be tagged with keywords indicating the language in which they are written and the operating system with which they work. HP DMA imports functions into the step code just prior to step execution—either by injecting the function directly where the replacement should occur or by using Python-style imports. Function scripts are usually written in Jython.

Functions are built and grouped based upon their main functionality, or the targeted database/middleware application. Here are some HP DMA-supplied function libraries:

Function library	Description
ostools	Contains many tools/methods for interactions with different operating systems, including general file handling, user permissions, and command-running.
commonvalidation	Contains validation methods common between different databases and operating systems, including file location validation and email address validation.
oraclevalidation	Contains Oracle specific validation tasks, including validate Oracle Home and validate Oracle user.
steplog	Contains step output and error output handling designed to make output easier and more consistent. This module features a debugging level to control the amount of output.

Any HP DMA user can view functions.

Source of functions

Functions are supplied as part of a solution pack (see [Solution Packs](#)). Functions that are part of a solution pack cannot be modified, but can be used in user-created steps and functions.

On the History tab you can view the installation history for HP DMA-supplied functions and the change history for custom functions.



Development tips

Only users with the DMA Admin role (see [Roles](#)) can create or modify functions.

Usually you create a new function when the code is used by multiple steps or has potential to be used in multiple steps.

Solution Packs

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

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

Solution packs are organized by function, for example: database patching or application server provisioning.

When you purchase or upgrade HP DMA, you are granted access to download solution packs.

When a solution pack is imported, the entire solution pack is inserted into the HP DMA back-end database (see [Oracle Database](#)). Solution packs can be imported, deleted, and rolled back.

Environment Objects

Environment objects refer to items describing either aspects of managed objects or the managed objects themselves. Environment objects are stored in the HP DMA repository.


The HP DMA environment objects include:

- [Organizations](#)
- [Servers](#)
- [Instances](#)
- [Databases](#)
- [Custom Fields](#)
- [Smart Groups](#)

Tip: Running the Discovery workflow will populate the metadata with the "discovered" servers, instances, and databases.

Organizations

An [organization](#) is a logical grouping of servers.



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

An organization is effectively just a named group of servers that can be used for permissions (see [Permissions Model](#)) or Custom Fields (see [Custom Fields](#)). There is no hierarchy or grouping of organizations.

Only users with the DMA Admins role (see [Roles](#)) can create or set permissions on organizations.

Servers

A **server** in HP DMA is an HP Server Automation (SA) Managed Server (see [Managed Servers](#)).

A server can only belong to a single organization.

HP DMA uses the SA-assigned host name—at the time the server is added to HP DMA. If the SA name changes, the server is not automatically updated in HP DMA; you must remove the server from HP DMA and re-add it to reflect the changed name.

To be added to HP DMA, the Managed Server must have the "DMA Client Files" Software Policy attached and remediated in SA. The user must also be able to "Read" the Managed Server in SA.

Tip: For more information see "Integrate HP DMA with HP Server Automation" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Permissions for a server are inherited from the organization that the server is associated with (see [Permissions Model](#)).


Tip: For more information about valid server operating systems and architectures, see "Supported Target Platforms" in the *HP DMA Support Matrix*, available at: support.openview.hp.com/selfsolve/manuals

Instances

An **instance** in HP DMA has different definitions depending on the supported product:

Product	Instance definition
Oracle	A set of processes running in memory that stores cached data and the engine that allows access to the underlying database.
SQL Server	A copy of the <code>sqlservr.exe</code> that runs as an operating system service. An instance can manage several system and user databases. A server can run multiple instances. Applications connect to the instance to process its databases.
Sybase	An SAP Adaptive Server. The Sybase client/server architecture that manages multiple databases and users, tracks the location of data on disks, maps logical to physical data storage, and maintains caches.
DB2	A database server instance. A logical database server environment with its configuration files, directories, and authorized users.
Middleware (WebSphere, WebLogic, JBoss)	A cell or domain is represented as an instance.

The following diagram shows the relationships between organizations, servers, and instances within HP DMA. Every server belongs to one organization. A server can have—but does not need to have—host instances.



An Instance can be in multiple servers, but the servers must be in the same organization. For example, in the diagram Instance1 is associated with both Server1 and Server2.

Instances have the following properties:

- General properties: name and type of instance
- Servers
- Connection
- Databases


Permissions for an instance are inherited from the organization and the server that the instance is associated with (see [Permissions Model](#)).

Databases

A **database** in HP DMA has different definitions depending on the supported product:

Product	Database definition
Oracle	A collection of physical files storing database objects: tables, view definitions, triggers, and more.
SQL Server	A collection of tables that store a specific set of structured data.
Sybase	A system for storing and retrieving data from two-dimensional tables that use SQL.
DB2	A collection of interrelated or independent data items that are stored together to serve multiple applications.
Middleware (WebSphere, WebLogic, JBoss)	Anything that is not an instance is represented as a database within HP DMA: application servers, node agents, node managers, clusters, and web servers.

The following diagram shows the relationships between organizations, servers, instances, and databases within HP DMA. Every server belongs to one organization. A server can have—but does not need to have—host instances. An instance can have—but does not need to have—host databases:



A database can be in multiple Instances, but the servers must be in the same organization.

Permissions for a database are inherited from the organization, server, and instance that the database is associated with (see [Permissions Model](#)).

Custom Fields

Custom Fields are used to customize workflows (see [Workflows](#)) or to show information about the environment. Custom Fields can be used in workflow steps (see [Steps](#)) to automatically supply information that is specific to an organization, server, instance, or database.

Custom Fields can be any of the following types:

Custom Field Type	Description
Text	The Text type of Custom Field stores text. The size of the column in the database is 1000 bytes.
Multi-lineText	The Multi-line Text type of Custom Field also stores text. In the user interface the user has a larger input box and can enter multiple lines.
Password	The Password type of Custom Field stores an encrypted password. If the value of a password type is null or empty, then nothing is displayed. If a value is specified, it is obfuscated. The Custom Field type cannot be changed to or from type Password.
List	The List type of Custom Field stores a set of possible values. For example, the Custom Field my_flag can only have values "TRUE" or "FALSE". By using a List type custom field and setting those as possible options, the user cannot mistype a value. No value is always a valid option for type List.

Anyone with the DMA User role can create Custom Fields, but a user must have Write permission on the Organization to set the value of the Custom Field (see [Permissions Model](#)).

Precedence

The same Custom Field can be defined at the organization level and lower levels—servers, instances, and databases. This enables you to use the Custom Field with some targets but not others. If the Custom Field is defined at both the organization level and a lower level, then the lower level value takes precedence over the organization level value.

The following shows the precedence of a Custom Field for an organization and a server in that organization:

Custom Field Precedence	Server value is specified	Server value is not specified
Organization value is specified	Use the Custom Field value for the server	Use the Custom Field value for the organization
Organization value is not specified	Use the Custom Field value for the server	Do not use a Custom Field value for this server

Smart Groups

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

Smart Group selection criteria is based on type, attributes, and potentially parent object attributes.

Valid comparison operators are:

- equals
- does not equal
- contains
- does not contain
- ends with
- starts with
- is empty
- is not empty

If more than one criteria is used in a Smart Group definition, the included members must meet ALL specified criteria.

Smart Groups are associated with a role (see [Roles](#)). The role's permissions are used when evaluating the Smart Group.

Smart Groups are typically used in deployments (see [Deployments](#)). You can also use Smart Groups to gather information, for example, to identify and quantify the Smart Group members.

Smart Groups are evaluated at run time and whenever the environment changes.

As the Smart Group criteria becomes more complex, the HP DMA run time slows down due to the time required to evaluate the Smart Group members.

Example 1

'My Servers' Smart Group for role 'DMA Admins':

- Server.OS contains 'linux'
- Organization.name equals 'New York'

Example 2

'My Instances' Smart Group for role 'DMA Users':

- Organization.Name does not equal 'New York'
- Instance.Url contains 'mycompany.com'
- Instance.DatabaseNames starts with 'ORA'
- Instance.oracle version equals '12.1.0'

Tip: You can use multiple Smart Groups in a deployment to merge the Smart Groups—allowing members of ANY of the selected Smart Groups to be available targets.

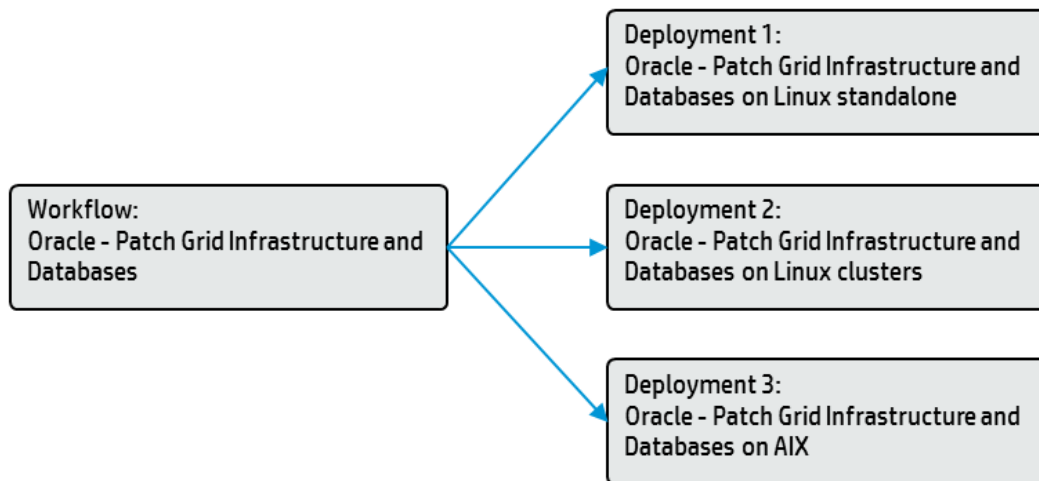
Bridge Objects

The objects that "bridge" between automation objects (see [Automation Objects](#)) and environment objects (see [Environment Objects](#)) include:

- [Deployments](#)
- [Policies](#)
- [Deployment Runs](#)

Deployments

A **deployment** associates a workflow (see [Workflows](#)) with the target environment (servers, instances, or databases) (see [Environment Objects](#)) where the workflow will run. To run a workflow, you execute a specific deployment. A deployment is associated with one workflow, but a workflow can have many deployments, each with its own targets (see [Deployment Targets](#)) and parameter settings:

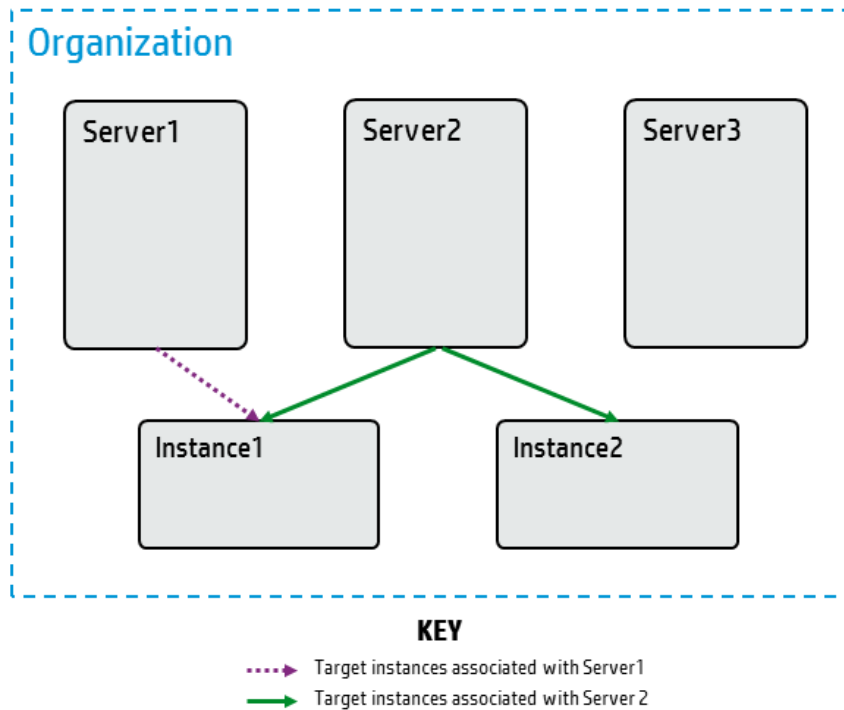


You must save a deployment before you can run the workflow. You can re-use a saved deployment many times.

Deployment Targets

A deployment represents a workflow (see [Workflows](#)) that can run on a designated set of targets. Depending on the level of the workflow (Server, Instance, or Database), the targets can either be specific objects of that level or a Smart Group (see [Smart Groups](#)) for that level that is available for one of the user's roles (see [Roles](#)).

The following is an example of instances used as deployment targets:



Users can only add specific targets if they have Write permission on the deployment and Deploy permission on a target's Environment (see [Permissions Model](#)).

Users can only execute the deployment if they have Execute permission on the deployment and Read permission on the target's Environment. HP DMA is designed so that you can have weaker users who can set values on a Deployment, but cannot add targets.

When a user edits a deployment, only the targets he has permission to see are displayed and saved. This can cause the list of available deployment targets to be "truncated" if a user with greater permissions creates a deployment, but later a user with fewer permissions saves it.

Deployment Parameters

The deployment can assign the following values to the workflow's parameters (see [Parameters](#)):

- Fixed Value - The specified value is used when the workflow runs.
- Custom Field - The value of the given Custom Field is used.

Note: This value varies depending upon the object on which the workflow runs, for example Server.name will be the name of the target's server.

- Policy Attribute - Current value of the specified policy attribute is used.
- Runtime Value - The user enters this value when the deployment is executed.

You can customize a deployment by specifying values for any workflow parameters that are designated User selected in the workflow.

Exposing Advanced Deployment Parameters

HP DMA workflows automatically "expose" the commonly used, required parameters. These parameters are usually inputs to a step called Gather Parameters for XYZ. When you create a deployment, HP DMA displays these parameters by default, and then you can specify values for them.

Many HP DMA workflows also have additional, advanced parameters. Advanced parameters cover less common workflow use cases and allow you to tailor advanced database and middleware variables to the specific needs of your IT organization. Advanced parameters are optional. To keep the user interface simple, they are "hidden." You can make advanced parameters available, as needed, for your use case and specify values appropriate for advanced functionality. These parameters are usually inputs to a step called Gather Advanced Parameters for XYZ.

To use these "hidden" parameters, you need to "expose" the parameters by changing the parameter mapping (see [Parameter mappings](#)) in the deployable copy of the workflow from No Value (or similar) to User selected:

The screenshot shows a configuration panel for a workflow step. At the top left, it says "Step A" and at the top right, "2, 3". Below this, there are two rows of parameter configurations. Each row consists of a label (e.g., "Parameter 1:") followed by a dropdown menu. Both dropdown menus currently display the text "- User selected -".

Once the parameter is available, you can specify a value—just as you would for any other User selected parameter—when you create a deployment of the workflow (see [Deployment Parameters](#)).

Scheduling Deployments

You can schedule deployments to run at a future time as long as they do not have Runtime Value parameters.

HP DMA handles scheduled deployments as follows:

- The scheduled deployment runs in the time zone of the HP DMA server.
- A scheduled deployment will not run if the deployment is already running on the target at the time when the deployment is scheduled.

Policies

Policies are reusable sets of attributes that can be used as parameter values in deployments (see [Deployments](#)). Policies allow the user to change a value in one place and have the new value used in multiple places, even across an entire enterprise.

Policies can contain fixed values or reference Custom Fields (see [Custom Fields](#)).

Advantages of policies:

- Deployments can reference policy attributes to change the automation behavior in a standard way.
- Policies enable HP DMA to manage groups of hundreds or thousands of servers at a time without needing to configure each individual server.
- Policies can be used to keep password parameter values secure. For example, this allows an HP DMA administrator to specify the password value in a policy and a user with the Workflow Runner role (see [Roles](#)) to execute a deployment that requires the password value—without the workflow runner knowing or entering the password value.
- Policies allow the user to specify values that change frequently—for example, passwords that must be changed regularly—in one place.

Policies are supplied as part of a solution pack (see [Solution Packs](#)). You cannot change the parameter definitions of policies that are part of an HP DMA solution pack, but you can set their values as appropriate for your environment. An HP DMA user can also create (or extract) a new policy to use common values across multiple workflows and deployments.

Policies have Role-Based Access Control (RBAC) for read and write (see [Roles](#)).

Workflows (see [Workflows](#)) cannot reference policies; only deployments can (see [Deployments](#)).

Policy values are evaluated when a deployment is started.

Deployment Runs

HP DMA users frequently talk about "Running" a workflow (see [Workflows](#)). Technically, it is the deployment (see [Deployments](#)) that is executed since the deployment associates a workflow (see [Workflows](#)) with targets (see [Deployment Targets](#)) and parameter values (see [Deployment Parameters](#)).

When multiple targets are selected for a deployment and the Run button is clicked, a **Deployment Run** is created for each target.

The HP DMA Console and History pages display deployment runs. The Console page shows running deployments and very recently completed deployments. The status updates throughout deployment execution. The History page displays both running deployments and completed deployments, and does not update unless the user selects it. The History page has powerful filtering capabilities to allow users to view the deployments of interest.

Connector Object

The Connector is the only object that connects HP DMA to HP Server Automation (the server management tool).

Connector

HP DMA includes a **Connector** component that enables HP DMA to communicate with HP Server Automation. You must configure the Connector before you can perform the following tasks:

- Configure roles (see [Roles](#)) and capabilities (see [Capabilities](#))
- View managed servers (see [Managed Servers](#))
- Add targets (see [Deployment Targets](#)) to a deployment (see [Deployments](#))
- Run a deployment (see [Deployments](#)) against a target

The initial HP DMA administrator (see [Special User: HP DMA Initial Administrator](#)) initially configures the Connector. Subsequently, any user with the DMA Admin role (see [Roles](#)) can modify the connector.

The configuration consists of the following:

- SA system name (or IP address)
- The SA username
- The SA password to connect to SA

Every time a user logs in to HP DMA, HP DMA retrieves the user and user group information from SA via the Connector.

Chapter 3: Platform Architecture

 **Dig deeper**

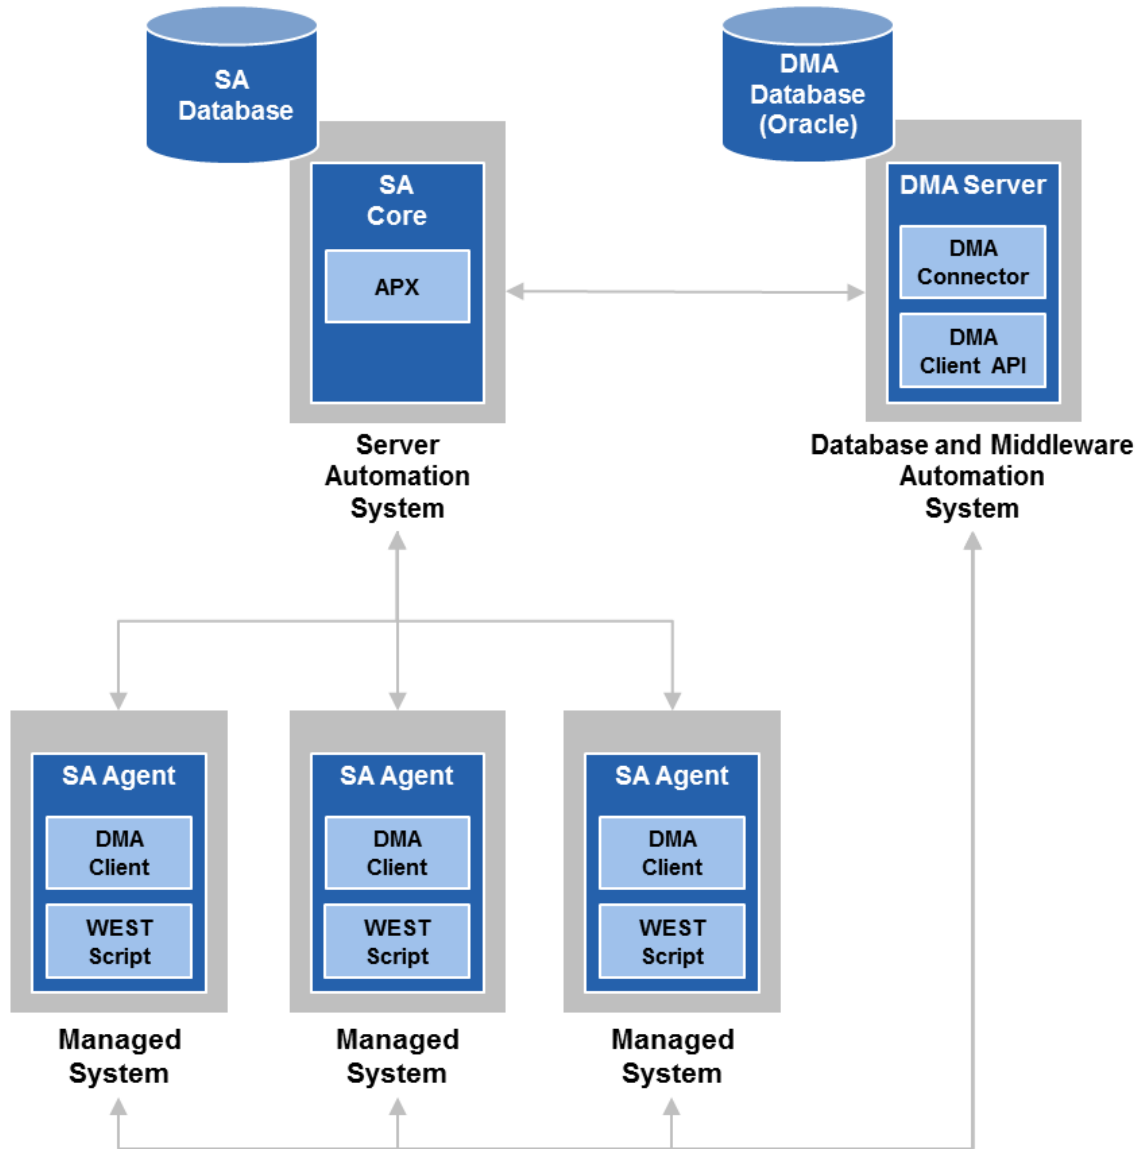
For the user who desires to dig deeper into the HP DMA platform architecture.

This chapter describes the architecture of the HP DMA platform and walks you through each of these components:

- [HP DMA Server](#)
- [HP Server Automation Server](#)
- [Oracle Database](#)
- [Managed Servers](#)

Architecture

Here is a simplified diagram of the various components of the HP DMA architecture and how they are connected:



Tip: For more information about how the HP DMA components interact, see [How the HP DMA Platform Works](#).

HP DMA Server

The HP DMA Server must be installed on a Linux-based operating system. It can be either a physical machine or a virtual machine.

For more information see "Supported Products and Platforms" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

To access the HP DMA user interface, go to: `https://<DMA_SERVER>:8443/dma/login`

Here, `<DMA_Server>` is the fully qualified host name of your HP DMA server.

REST application programming interfaces (APIs) documentation (see [REST APIs](#)) is available on all HP DMA Servers at `https://<DMA_SERVER>:8443/dma/api`.

HP DMA Servers can be clustered. You can connect multiple HP DMA Servers to a single back-end database (see [Oracle Database](#)).

Tomcat

HP DMA is shipped with Apache Tomcat, an open source software implementation that powers numerous large-scale, mission-critical web applications.

The Tomcat context file, `dma.xml`, contains the data used to configure HP DMA and the connection information for the database. It is created when you run the `dmaBaselineData` command (see [DMA Baseline Data](#)).

Tip: For more information see the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

HP Server Automation Server

HP DMA uses HP Server Automation (SA) as an agent infrastructure.

HP DMA integrates with SA to perform the following tasks:

- Authenticate users
- Associate users with groups
- Determine user privileges
- Acquire knowledge of servers

Note: Any server that will be used as an HP DMA target (see [Deployment Targets](#)) must be managed by SA. It must also have the DMA Client Files software policy attached.

- Send requests to execute workflows on servers
- Communicate securely
- Stores common files in the software repository

Note: The [software repository](#) contains any files that a workflow (see [Workflows](#)) might need to carry out its purpose (for example, software binaries or patch archives). If the files that a workflow requires are not stored locally on each target server, the workflow looks for them in the software repository.

Integration

Before HP DMA is ready to use, the SA Administrator and HP DMA Administrator must perform a series of integration steps on your SA system as well as on your HP DMA server.

Tip: For more information see "Supported Products and Platforms" and "Integrate HP DMA with HP Server Automation" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

HP DMA only talks to a single SA Server (SA Core).

SA functionality used by HP DMA

HP DMA uses the following SA functionality:

- An SA Application Programming Extension (see [Automation Platform Extension \(APX\)](#))
- The DMA Client Files software policy that includes the DMA Client (a script called by the SA Agent) and the run-time configuration necessary for HP DMA workflows

Oracle Database

HP DMA uses an Oracle Database as a back-end database.

The database contains:

- HP DMA automation—workflows, steps, functions
- HP DMA environment information—the organizations, servers, instances, databases
- HP DMA bridge information—the deployments and policies
- The history of HP DMA workflows that were run

Tip: For more information see "Supported Products and Platforms" and "Create and Configure the Oracle Database" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Managed Servers

HP DMA requires **managed servers** to run. The HP DMA environment (see [Environment Objects](#)) contains information about the managed servers. The servers are targets for HP DMA deployments (see [Deployments](#)).

HP DMA receives its information about servers from HP Server Automation (SA). The managed servers that you can "see" as available in HP DMA are the subset of servers that both the Connector to SA can "see" (see [Connector](#)) and your roles permit you to "see" (see [Roles](#)).

Managed servers have the following SA attributes:

- Object ID
- Name
- Hostname

Tip: For more information about how to configure SA managed servers for HP DMA see "Integrate HP DMA with HP Server Automation" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Chapter 4: How the HP DMA Platform Works

Dig deeper

For the user who desires to dig deeper into how the HP DMA platform works.

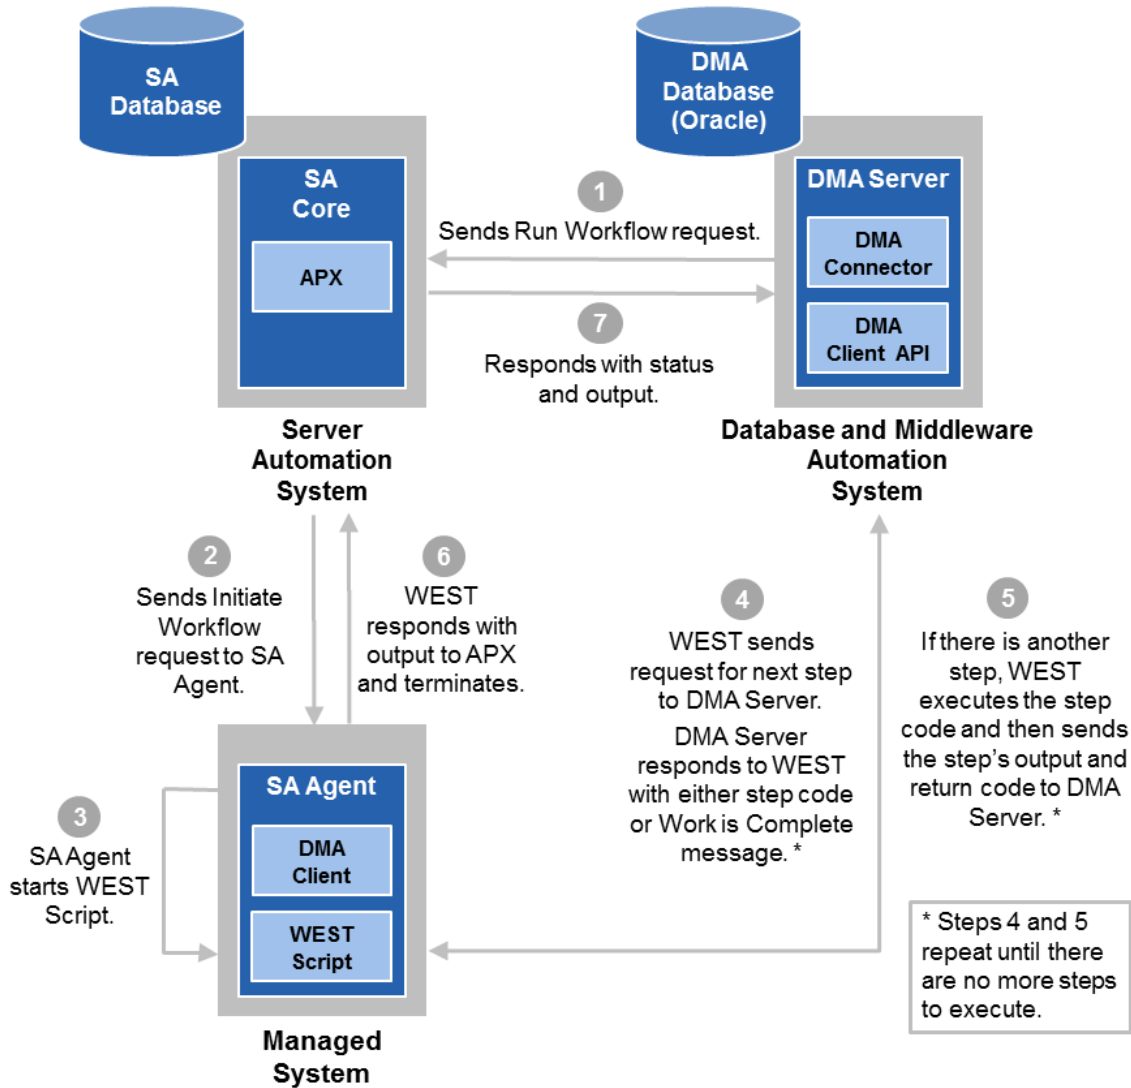
This chapter describes how the HP DMA platform works behind-the-scenes. It covers the following topics:

- [Typical flow of HP DMA workflow execution](#)
- [Workflow Execution Script \(WEST\)](#)
- [REST APIs](#)
- [DMA Baseline Data](#)
- [Automation Platform Extension \(APX\)](#)

Typical flow of HP DMA workflow execution

This section shows how—and in what order—the HP DMA components work together to execute a workflow within a simple HP DMA architecture (see [Platform Architecture](#)).

Tip: New terms are explained later in this chapter.



Note: This diagram is based on the HP DMA architecture diagram—[Architecture](#)—but with a single Managed System.

Process Description

- 1 The HP DMA Server (see [HP DMA Server](#)) communicates with the SA Server (see [HP Server Automation Server](#)) to start the HP DMA APX (see [Automation Platform Extension \(APX\)](#)).
- 2 The HP DMA APX in turn starts the DMA Client on the managed server (see [Managed Servers](#)) and provides the options for WEST (see [Workflow Execution Script \(WEST\)](#)) execution. If more than one managed server is used, the HP DMA APX starts the DMA Client on each one.
- 3 The HP DMA APX communicates with the SA Agent on the managed server and starts WEST Script.

Note: Steps 4 and 5 repeat until all workflow steps complete.

- 4 WEST sends an HTTP message—communicating via REST APIs (see [REST APIs](#))—directly to the HP DMA Server requesting the next workflow step (see [Steps](#)).

The HP DMA Server determines which step should run next—the first step of the workflow, the next step as determined by the previous step's return code, or no additional steps. Then the HP DMA Server responds with one of the following:

- The HP DMA Server responds with the information for the step to be executed. This includes the step name, call wrapper, step code, and function code. WEST creates a working directory, starts the queue monitor thread, and uses pipes for input and output.
- If there are no additional steps to execute—the workflow is complete—the HP DMA Server responds telling WEST to stop.
- If there currently is no more input for WEST, the HP DMA Server responds directing WEST to wait. This can occur when running a Bridged Execution Workflow where different steps run on different managed servers (see [Bridged Execution Workflows](#)).

- 5 WEST executes the step code. While the step is executing, WEST sends messages—regular "heartbeats", output information, and output parameter values—to the HP DMA Server. When the step has completed executing, WEST sends the step return code to the HP DMA Server. The information will be available in the Step Output, Step Errors, and Step Header tabs on the History page in the HP DMA Server UI.
- 6 The HP DMA APX gathers the output from WEST and then WEST terminates.
- 7 The HP DMA APX sends the workflow status and the output information to the HP DMA Server. This information will be available in the Connector Output and Connector Errors tabs on the History page in the HP DMA Server UI.

Workflow Execution Script (WEST)

Each HP DMA managed server (see [Managed Servers](#)) uses a program called **Workflow Execution Script (WEST)** to communicate with the HP DMA server (see [HP DMA Server](#)). WEST does the following:

- Executes workflow steps on the Managed Servers
- Communicates with the HP DMA Server via HTTPs
- Provides the output (stdout, stderr, return code, and end time) for step execution

WEST is installed on each managed server when you attach and remediate the DMA Client Files software policy on that target.

Tip: For more information see "Integrate HP DMA with HP Server Automation" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

WEST is the main script (and libraries) that are contained by the DMA Client. The script starts, and then handles the complete cycle of a workflow execution on the managed server.

Under certain circumstances, you may need to manually terminate WEST on a managed server. This would be necessary, for example, if the HP DMA server name was specified incorrectly when the `dmaBaselineData` command was executed, and a workflow execution was subsequently attempted.

Workflow Execution Engine

The **Workflow Execution Engine** is effectively a state machine driven by WEST. It handles all workflow processes except for aborting or canceling a workflow. WEST drives each state transition by notifying the HP DMA Server of the exit code for each step's execution (see [Deployment Runs](#)). The Workflow Execution Engine manages the states of steps (see [Steps](#)) and workflow deployments (see [Deployments](#)).

Step Execution States

State	Description
Aborted	Only applies to steps of type Script. Occurs if the deployment aborts after the server initiates a step. Most likely to occur for the first step of a workflow.
Initiated	Only applies to steps of type Script. The step has been prepared for execution but execution has not yet started.
Finished	Applies to all step types. The step has completed.
Running	Applies to all steps types. The step is currently executing.

Deployment Execution States

State	Description
Aborted	HP DMA aborted the deployment due to an unexpected or unrecoverable error condition.
Cancelled	The HP DMA user canceled the deployment before it had completed.
Failure	The deployment completed and the last completion step was Failure.
Finished	The deployment completed but neither the Failure or Success step was executed (see Success/Failure).
Running	The deployment is currently executing.
Skipped	The deployment was scheduled to run at the given time, but was not started because the deployment was already running on the target server.
Success	The deployment completed and the last completion step was Success.

Custom Fields used with WEST

The HP DMA Client uses the following Custom Fields (see [Custom Fields](#)) with WEST:

Custom Field	Description
agent_username_win	Username for a renamed local administrator (Windows).
domain_username_win	Username for the domain user (Windows).
domain_password_win	Password for the domain user (Windows).
west_message_size	Maximum size in characters of the HTTP messages. Default is 262144 characters.
west_timeout	Time out in seconds for HTTP messages. Default is 60 seconds.
west_retries	Number of retries. Default is unlimited.
west_heartbeat	Interval for heartbeat messages. Default is 180 seconds.
west_proxy_in_use	Flag to indicate whether or not a proxy is being used. Valid values are TRUE and FALSE. If specified, then <code>west_proxy_address</code> must also be specified.
west_proxy_address	Proxy address. If specified, <code>west_proxy_in_use</code> must also be specified. If using the SA Gateway Network, then set the value to <code>sa_auto_select</code> .
west_verbose	Flag to indicate whether or not to turn on debug logging in the HP DMA Client log. Valid values are TRUE and FALSE.
west_keep	Flag to indicate whether to save the workflow execution folder and all files contained in it. Valid values are TRUE and FALSE. If set to FALSE the folder and its files will be deleted.

Tip: For examples of these Custom Fields, see the *HP DMA Administrator Guide* and the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

REST APIs

The HP DMA REST application programming interfaces (APIs) are a collection of tools that you can use to operate HP DMA programmatically.

Representational State Transfer (REST) is a software design philosophy that focuses on resources and adheres to a set of specific architectural constraints. A resource is a piece of information that can be uniquely identified and described by a specific type of representation. A resource may map to a thing (for example: a physical object, a concept, a phrase), but it is not the thing itself. Multiple resources can map to a single thing.

For example, consider the book *War and Peace* by Leo Tolstoy. The following resources identify this book: number seven on my list of favorite books, the last book that I read, the thickest book on my bookshelf.

The HP DMA APIs enable you to read and—in some cases—create, modify, and delete HP DMA objects: environment objects, automation objects, the set up, and solution packs.

Tip: *HP DMA API Reference WebHelp* is available on all HP DMA Servers at:

`https://<DMA_SERVER>:8443/dma/api`

Here, <DMA_SERVER> is the fully qualified host name of your HP DMA server.

The documentation gives examples of the XML and lists valid methods and query parameters.

DMA Baseline Data

For HP DMA to be usable, the administrator must run the **DMA Baseline Data** command, `dmaBaselineData.sh`. Depending on the options, the command performs the following tasks:

- Creates the HP DMA database tables
- Loads the "baseline" (initial) data into the HP DMA tables
- Creates the `dma.xml` file that specifies the database connection information (see [Oracle Database](#))
- Specifies the Java Database Connectivity (JDBC) Connection String used to connect to the database
- Updates the HP DMA database tables when HP DMA is upgraded to a new release
- Generates and overwrites HP DMA keys

HP DMA Installation

To initially install HP DMA you need to run `dmaBaselineData.sh`. The command creates the context file, `dma.xml`, with the database connection information, creates the tables, loads the baseline data, creates the public and private keys, and creates the default organization.

Tip: For more information see "Install the HP DMA Server" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

HP DMA Upgrade

To upgrade HP DMA to a newer HP DMA release, you must also run `dmaBaselineData.sh`. To update, you typically run `dmaBaselineData.sh` with no options so that it reads the database connection information from the `dma.xml` file, and then makes the required updates.

Tip: For more information see "How to Upgrade HP DMA" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Troubleshooting

Tip: For information about errors that you may encounter related to baselining, see "Common Baseline Errors" in either the *HP DMA Installation Guide* or the *HP DMA Troubleshooting Guide*, available at: support.openview.hp.com/selfsolve/manuals

HP DMA Baseline Options

Tip: For a complete list of baseline options, see "HP DMA Baseline Options" in either the *HP DMA Installation Guide* or the *HP DMA Troubleshooting Guide*, available at: support.openview.hp.com/selfsolve/manuals

Automation Platform Extension (APX)

HP DMA needs HP Server Automation **APXs (Automation Platform Extensions)** to start WEST (see [Workflow Execution Script \(WEST\)](#)) and to run a workflow (see [Deployment Runs](#)). The APX layer allows HP DMA to add users and roles to SA so that SA can communicate with HP DMA.

The `westapx` verifies that the DMA Software Policy is attached and remediated to the managed servers (see [Managed Servers](#)) and to start WEST (see [Workflow Execution Script \(WEST\)](#)). It also validates that the managed server are in an SA lifecycle of "Managed".

The `updateWinAdmin` APX sets the Windows Administrator that WEST uses. This APX is only needed when the default Windows Administrator has been renamed or replaced.

The APX files are zipped together and delivered as part of the HP DMA Client for SA.

Note: The APX completes after WEST finishes running a workflow. Thus, there is a delay in receiving the `stdout` and `stderr` output messages.

Additional Information

For information about importing the APX, see "Import the HP DMA APX" in the *HP DMA Installation Guide*. The guide describes how to load the APXs depending on the version of HP Server Automation that you use. On Enterprise SA it can be loaded manually using the `apxtool` tool or via Live Network Connector (LNC). On SAVA the APXs can only be loaded using the Live Network Connector (LNC).

For troubleshooting information see "APX Tool Configuration Error" in the *HP DMA Installation Guide* or the *HP DMA Troubleshooting Guide*.

For HP Server Automation APX documentation see the following:

- "Creating Automation Platform Extensions (APX)" in the *SA Platform Developer Guide*
- "Running SA Extensions" in the *User Guide: Server Automation*

These documents are available at: support.openview.hp.com/selfsolve/manuals

Chapter 5: Permissions Model

This chapter describes the permissions model employed by HP DMA and addresses the following topics:

- [Roles](#)
- [Capabilities](#)
- [HP DMA User](#)
- [Permissions](#)

HP DMA uses Role-Based Access Control (RBAC) for its permissions model to do the following:

- Provide granular controls of features and functions available on a per group basis
- Enable granular control of which servers/assets a user may view, access, or execute against
- Track the history for all automation objects (see [Automation Objects](#))

Permissions are created for specific objects and specific roles. A user's permissions for a given object are the union of both the user's permissions for the object (see [Permissions](#)) and the user's roles (see [Roles](#)).

Roles

The HP DMA permission model is based on **roles**—called Role-Based Access Control (RBAC).

Obtained from SA

HP DMA obtains the complete set of available roles from HP Server Automation (SA) through the Connector (see [Connector](#)). The Connector retrieves the SA public User Groups—including the groups that your SA administrator configures for HP DMA.

Commonly defined roles are:

- DMA Admins
- DMA Users
- DMA Workflow Developers

You can create additional roles based on your enterprise's needs.

Each HP DMA user has one or more roles. The roles that each user has are defined in SA, and cannot be changed within HP DMA.

HP DMA warns you if a particular role no longer exists in SA.

Made Available in HP DMA

While you are logged in as an HP DMA administrator, you need to register the roles that you want HP DMA to use by changing the status of SA public user groups from Available to Registered.

The available roles are saved in the HP DMA database (see [Oracle Database](#)).

A role can have multiple capabilities (see [Capabilities](#)).

Tip: For more information see "Set Up the SA Groups and Users" and "Set Up HP DMA" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Capabilities

Capabilities are collections of related privileges. They control what actions the specific roles (see [Roles](#)) are allowed to perform—mapping between a role and objects (see [HP DMA Objects](#)).

You must assign capabilities to each role that you register. Initially, the HP DMA initial administrator (see [Special User: HP DMA Initial Administrator](#)) assigns capabilities to roles.

After that, any user with the DMA Admins role can set capabilities.

The capabilities are:

- Login - User can log in to HP DMA
- Workflow Creator - User can create workflows (see [Workflows](#))
- Administrator - The Administrator is the Super User in HP DMA and is implicitly a workflow creator. Anyone with this capability can view, edit, or perform any task in HP DMA. The Administrator capability is (in most companies) synonymous with the DMA Admins role, yet multiple roles can be Administrators.

Capabilities are used to allow users to perform additional tasks in HP DMA that are not representable by Permissions (see [Permissions](#)).

Tip: For more information see "Set Up the SA Groups and Users", "Start HP DMA" and "Set Up HP DMA" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

HP DMA User

The concept of a "user" comes from HP Server Automation (SA). Users are configured in SA—added, deleted, and modified. SA authenticates the user and password.

HP DMA does not store "users". On login, the user and group information is retrieved from SA via the Connector (see [Connector](#)). Based on the Roles (see [Roles](#)) and the Capabilities (see [Capabilities](#)) that are configured, HP DMA determines whether a user is allowed to log in and what actions a user can perform.

HP DMA calls the Connector (see [Connector](#)) regularly to reevaluate the user Roles and Capabilities. If the user gains or loses Administrator capability, the user is forced to log in again—ensuring that the proper menus are displayed. If the user loses the Login Access capability, the user is logged out.

SA User Groups for HP DMA

The following chart gives examples of the SA user groups that must be set up by an SA administrator before HP DMA can be used. The user groups—along with the associated roles (see [Roles](#)) and capabilities (see [Capabilities](#))—are required to use and manage HP DMA in your environment.

Group Type	SA User Group - HP DMA Role	Capability Required	Description
HP DMA administrators	DMA Admins	Administrator	Users in this group will perform HP DMA administrative duties.
Users who will create HP DMA workflows	DMA Workflow Creators	Workflow Creator	Users in this group will have the ability to create HP DMA workflows. Note: Once a workflow is created, it can be modified using Role-Based Access Control (RBAC) as needed.
Users who will run HP DMA workflows	DMA Workflow Runners	Login Access	Users in this group will have the ability to run HP DMA workflows.

Tip: For more information see "Set Up the SA Groups and Users" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Special User: HP DMA Initial Administrator

The HP DMA initial administrator accomplishes the initial configuration of the Connector (see [Connector](#)). Since the Connector has to be configured by a user with the DMA Admin role (see [Roles](#)) and HP DMA has no idea of what users are DMA Admins without the Connector, HP DMA requires a specific user to log in and perform the initial configuration when HP DMA is started after the initial install. Effectively, this is the only HP DMA user. This user is always `dma_initial_admin`.

The HP DMA initial administrator does the following:

- Sets the initial password
- Configures the Connector (see [Connector](#))
- Registers the HP DMA Roles (see [Roles](#))
- Assigns the HP DMA Capabilities (see [Capabilities](#))

The HP DMA initial administrator remains after the initial configuration and can be used again. The HP DMA initial administrator can be used in these situations:

- If the connector user password changes
- If HP DMA cannot connect to SA
- If the SA connector needs to be reconfigured

Tip: For more information see "Start HP DMA" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Special User: HP DMA Connector User

A single SA user—the HP DMA connector user—is required to configure the HP DMA Connector (see [Connector](#)) to SA (see [HP Server Automation Server](#)). This user is used by HP DMA to connect to SA whenever a specific, personalized SA account cannot be used—for example, to verify whether a login is allowed.

This user requires special permissions that are described in the *HP DMA Installation Guide*.

Tip: For more information see "Set Up HP DMA" in the *HP DMA Installation Guide*, available at: support.openview.hp.com/selfsolve/manuals

Permissions

A **permission** is a class used to associate permissions (Read, Write, Execute, Deploy) between a given role (see [Roles](#)) and a given object—not between a user and an object.

If you delete a role, all of the associated permissions are also deleted.

The permissions are discussed by type of object:

- [Permissions for Automation Objects](#)
- [Permissions for Environment Objects](#)
- [Permissions for the Bridge Objects](#)

Permissions for Automation Objects

The permissions for automation objects (see [Automation Objects](#)) are different depending upon whether the automation object comes from a solution pack (see [Solution Packs](#)) or is one that you created (or copied).

Automation Object	Permissions for solution pack objects	Permissions for created objects
Workflow	Locked. Cannot change to READ or WRITE.*	READ, WRITE
Step	Locked. Cannot change to WRITE.*	WRITE

* The HP DMA-supplied workflows and steps come locked and cannot be changed to READ/WRITE. If you want to change a workflow or step, you need to copy it and make the desired changes in the copy.

Permissions for Environment Objects

Environment objects (see [Environment Objects](#))—servers, instances, and databases—inheriting permissions from the organization.

Environment Object	Permissions
Organization	READ, WRITE, DEPLOY
Server	Inherited through organization
Instance	Inherited through server and organization
Database	Inherited through instance, server, and organization

Organizations have DEPLOY permission instead of EXECUTE permission. Having DEPLOY permission on an organization means that you can create deployments on managed servers in the organization. HP DMA administrators can use the DEPLOY permission to limit what target servers a user can run a workflow on.

Having WRITE permission on an organization means that you can change the organization—or any of the objects in it.

To create an Organization or update its permissions through the HP DMA user interface, you must have the DMA Admin role (see [Roles](#)).

A user must have DEPLOY permission on an Organization to add servers to a Deployment (see [Deployments](#)), either manually or through a Smart Group (see [Smart Groups](#)). A user is allowed to execute a deployment on a given server if the user has READ permission on the Organization.

Permissions for the Bridge Objects

The following permissions are available for the Bridge Objects (see [Bridge Objects](#)):

Bridge Object	Permissions for solution pack objects	Permissions for created objects
Deployment	<i>Not delivered in solution packs.</i>	READ, WRITE, EXECUTE
Policy	Can change to READ and WRITE. Cannot add or delete attributes.	READ, WRITE

Deployments are the only object with EXECUTE permission.

Additional Information

This section explains the relationships of permissions between automation objects (see [Automation Objects](#)), environment objects (see [Environment Objects](#)), and the bridge object (deployment, see [Deployments](#)).

Use Case	Automation Permissions	Environment Permissions	Deployment permissions
To execute a deployment		READ permission on the organization that contains the target	EXECUTE permission on the deployment
To edit or add targets to a deployment		READ or DEPLOY permission on the organization	WRITE permission on the deployment
To create a deployment	READ permission on the workflow	READ or DEPLOY permission on the organization	WRITE permission on the deployment

Interesting cases

If your role(s) have READ permission on an organization and WRITE permission on a deployment, then you cannot "see" servers in the organization. At least one of your roles must have DEPLOY permission on the organization.

If your role(s) have READ/WRITE/DEPLOY permissions on an organization but only READ permission on a deployment, then you can only look at the deployment. You cannot run it since at least one of your roles must have EXECUTE permission for the deployment.

Permissions on the workflow only matter if you are creating a new deployment. To create a deployment, at least one of your roles must have READ permission on the workflow. If another user's role(s) have READ permission on the deployment but do not have READ permission on the workflow, then that user can "see" the deployment but not the associated workflow,.

Users may have different views of targets based on what organizations their roles have DEPLOY permission.

If User 1's role has DEPLOY permission on Organizations A and B, and then User 1 creates a deployment XYA on A and B and adds servers from both A and B to XYZ. If User 2's role has WRITE permission on deployment XYZ and DEPLOY permission on Organization A but not on B, and then User 2 resaves XYZ. In this situation, deployment XYZ no longer contains servers from Organization B because User 2's role could not "see" the servers.

Chapter 6: Special Types of Workflows

This chapter describes two unique types of HP DMA workflows:

- [Bridged Execution Workflows](#)
- [Master Workflows](#)

Bridged Execution Workflows

Bridged Execution Workflows are multi-target workflows. In other words, the workflows run on multiple systems. A bridged execution workflow includes some steps that run on one target and other steps that run on another target. Several target can be used (see [Deployment Targets](#)).

How a bridged execution workflow works

The workflow defines which step runs on which target. If a target is not specified, the workflow runs on the default target.

Steps in bridged execution workflows have several differences. Re-targetable steps have a box checked to indicate Targetable and have an additional parameter, Step Target. If the Targetable box is not checked, the step can only run on the primary target.

Running a bridged execution workflow has several differences from running a regular HP DMA workflow. At runtime you are asked to specify the different targets. Internal to the run, HP DMA starts WEST (see [Workflow Execution Script \(WEST\)](#)) on all of the target servers. WEST on each target server repeatedly "asks" the HP DMA server (see [HP DMA Server](#)) if there is anything to do. The HP DMA server orchestrates the steps so that only one step runs at a time. For more information see [Typical flow of HP DMA workflow execution](#).

If you watch the HP DMA Console while the bridged execution workflow is executing, you will observe a single deployment name that has different steps running on different target servers. Only one step runs at a time. Each step runs on the target defined in the workflow parameter mapping.

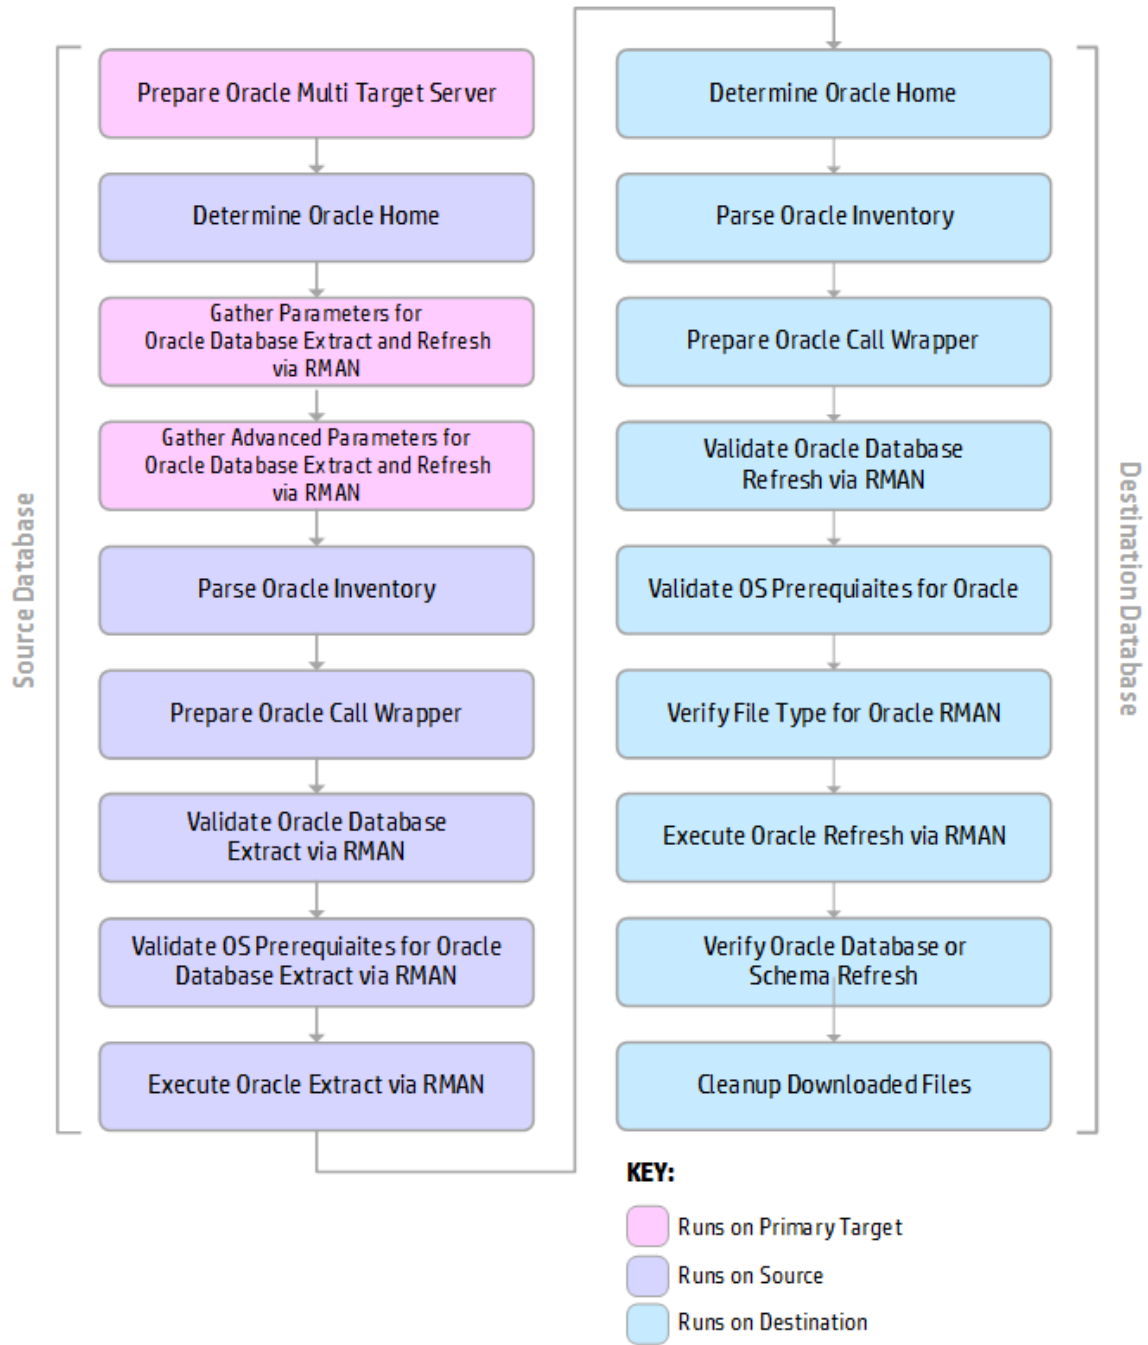
Bridged execution workflows are supported on HP Server Automation version 9.11 (and later).

Example

An example of a bridged execution workflow is Oracle - Extract and Refresh Database via RMAN, found in the HP DMA 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.

The following diagram shows all of the steps in the workflow. The color codes indicate which steps run on which targets:



The first step of the workflow, Prepare Oracle Multi Target Server, has input parameters Source and Destination that are specified at runtime. The values of Source and Destination are mapped to Step Target in subsequent steps so that the steps run on the appropriate target. In this case, all targets are at the same level—Oracle instances.

Prepare Oracle Multi Target Server

General Action **Parameters**

Parameters

INPUT PARAMETERS

- Destination
- Source

OUTPUT PARAMETERS

- Destination
- No Value
- Server Wrapper
- Source

One of the subsequent steps, Parse Oracle Inventory, is re-targetable because the box is checked:

Parse Oracle Inventory

General **Action** Parameters

Properties

Name: Parse Oracle Inventory
Tags: Instance, Prepare
Type: Oracle
Category: Script
Targetable: ←

The workflow runs the step Parse Oracle Inventory two times. In the first run of Parse Oracle Inventory, the Step Target parameter is mapped to Source:

▼ 8 [Parse Oracle Inventory](#) 0 7, 9

Step Target: Prepare Oracle Multi Target Server.Source ←

In the second run of Parse Oracle Inventory, the Step Target parameter is mapped to Destination:



When you run Oracle - Extract and Refresh Database via RMAN, on the Run page you need to select values for the target parameters:

Prepare Oracle Multi Target Server

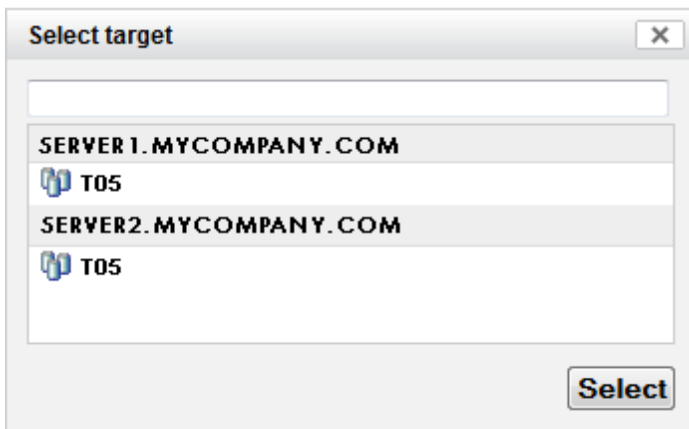
Target Parameters

Primary Target: [SELECT](#)

Destination: [SELECT](#)

Source: [SELECT](#)

You can choose targets from the list of Available Targets that you defined when you created the deployment:



HP DMA Bridged Execution Workflows

The following bridged execution workflows are delivered in HP DMA solution packs:

Solution Pack	Bridged Execution Workflows
Database Refresh	Backup and Restore MS SQL Database Dump and Load Sybase Database Oracle - Export and Refresh Database via Data Pump Oracle - Extract and Refresh Database via RMAN

Master Workflows

Master workflows are workflows that can run multiple workflows in a single execution. They are useful if you want to run several workflows in an orderly, repeatable manner.

How a master workflow works

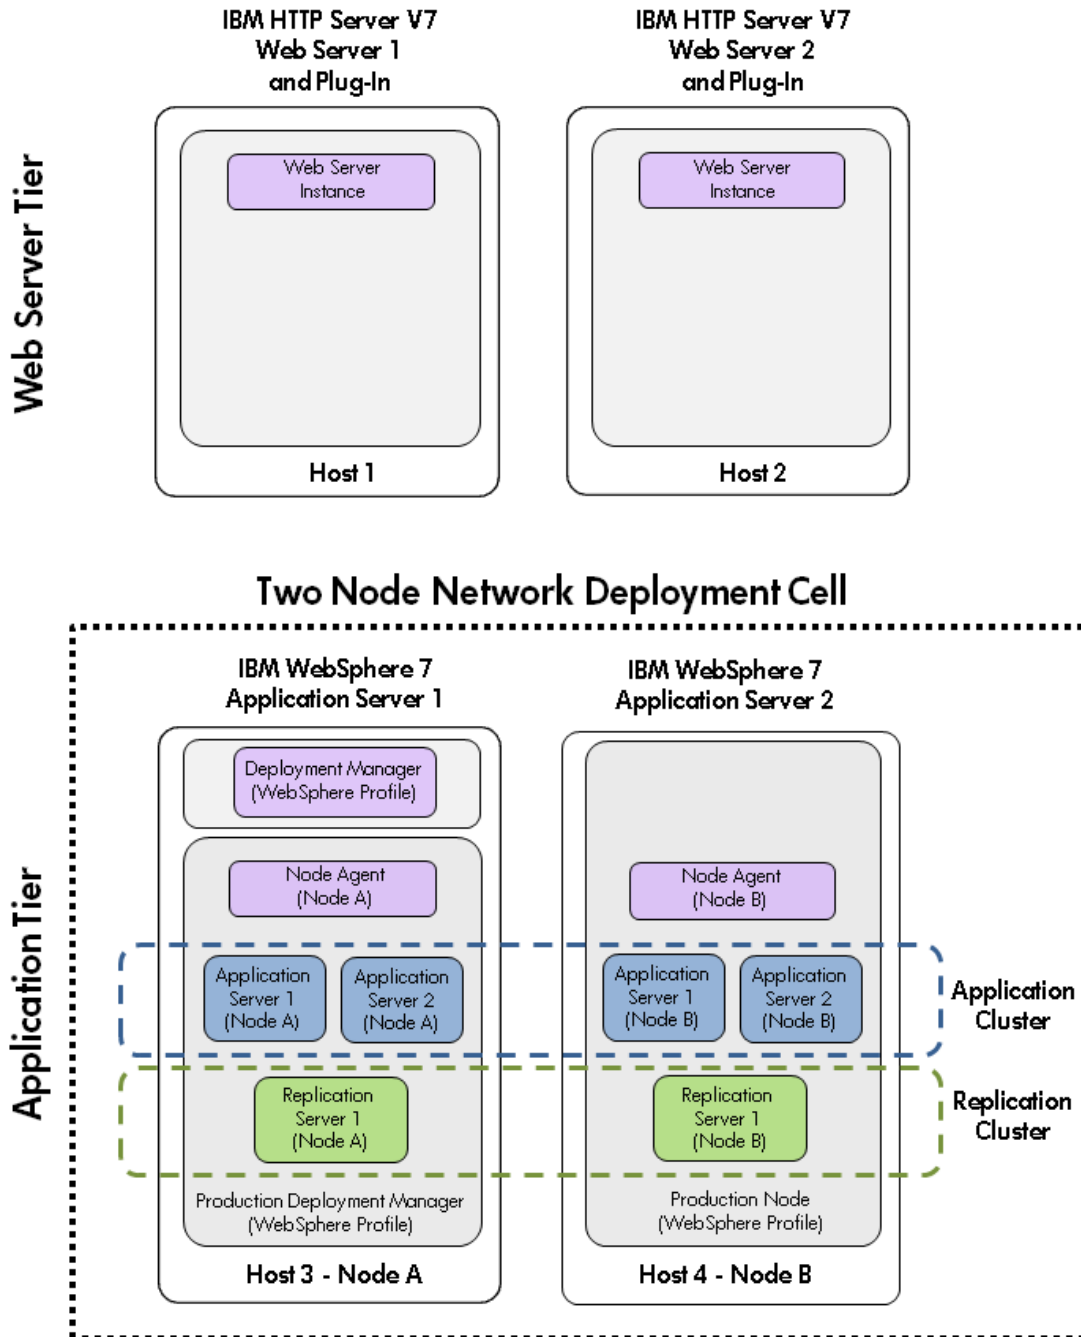
HP DMA can convert a regular workflow into a component workflow. The component workflow becomes a single step—a subflow—in a merged, composite "master workflow." This new step is normally named "Run" followed by the original workflow name.

The master workflow uses a REST API capability (see [REST APIs](#)) to run a workflow without a deployment (see [Deployments](#)), by only passing the workflow name, target server, and input parameters to the API call.

If you watch the HP DMA Console while the master workflow is executing, you will first observe the deployment for the master workflow starts. Whenever the master workflow encounters a subflow step, a deployment starts for the component workflow. When the component workflow finishes, control returns to the master workflow. All deployments run on the same target (unless multiple deployment targets are selected at runtime—then the master workflow and its component workflows run in parallel on each target).

Example

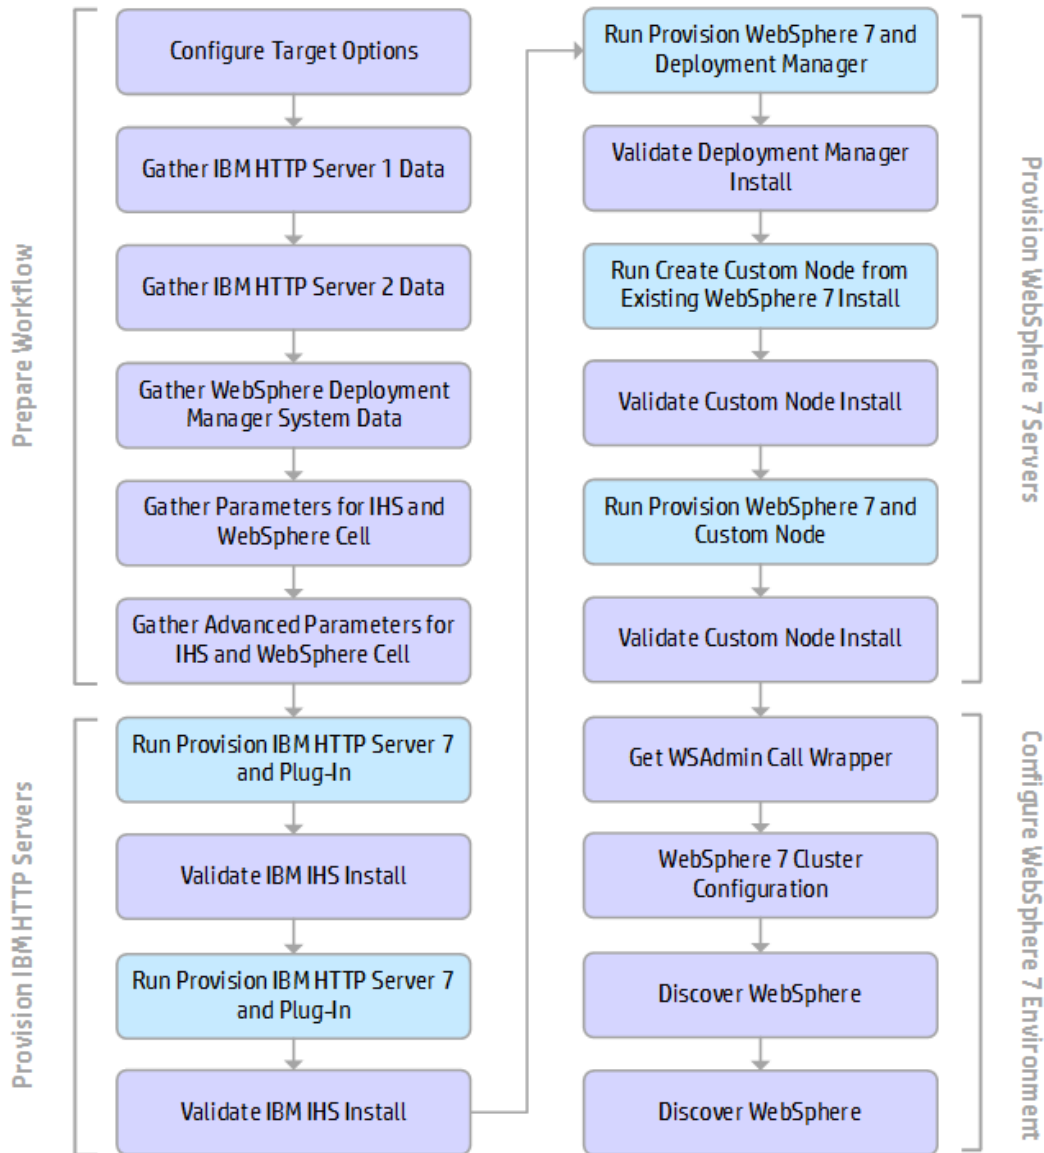
As shown below, you may want to provision a web server tier with two web servers and an application tier with a two node Network Deployment cell. Each node is an IBM WebSphere Application Server version 7 server.



You can accomplish this architecture by running five HP DMA workflows (found in the HP DMA Application Server Provisioning Solution Pack) in an orderly, repeatable manner:

HP DMA Workflow	Purpose in the Architectural Diagram
Provision IBM HTTP Server 7 and Plug-In	Provisions IBM HTTP Server 1
Provision IBM HTTP Server 7 and Plug-In	Provisions IBM HTTP Server 2
Provision WebSphere 7 and Deployment Manager	Provisions IBM WebSphere 7 Application Server 1 with a Deployment Manager
Create Custom Node from Existing WebSphere 7 Install	Provisions IBM WebSphere 7 Application Server 1
Provision WebSphere 7 and Custom Node	Provisions IBM WebSphere 7 Application Server 2

You can also accomplish this by running a single master workflow. The workflow Provision IBM HTTP Server and WebSphere 7 Two Node Cell (also found in the HP DMA Application Server Provisioning Solution Pack) creates the entire architecture. The following diagram shows all of the steps in the workflow. The subflow steps—steps that invoke the component workflows—are in blue:



- KEY:**
- Normal workflow step
 - Component workflow step (subflow)

HP DMA Master Workflows

The following master workflows are delivered in HP DMA solution packs:

Solution Pack	Examples of Master Workflows
Database Provisioning	Deploy MS SQL 2008 R2 Cluster Deploy Sybase ASE 15 SMP Server MS SQL - Install Cluster Patch
Advanced Database Provisioning	Provision Dataguard One Node RAC Provision One Node RAC
Application Server Provisioning	Provision IBM HTTP Server and WebSphere 7 Two Node Cell Add WebSphere 7 Node To Existing Cell Provision HTTP Server and WebSphere 7 StandAlone Profile

Appendix A: HP Software Documentation

HP Database and Middleware Automation Documentation

The following documents are included in the HP DMA documentation library:

- *HP DMA Installation Guide*
- *HP DMA Troubleshooting Guide*
- *HP DMA Administrator Guide*
- *HP DMA User Guide*
- *HP DMA Quick Start Tutorial*
- *HP DMA Concepts Guide*
- *HP DMA Release Notes*
- *HP DMA Support Matrix*
- *HP DMA Solution Pack User Guides*

The latest versions of these documents are available on the HP Software Support web site:

<https://softwaresupport.hp.com/>

HP DMA API Reference WebHelp is available on all HP DMA Servers at:

`https://<DMA_SERVER>:8443/dma/api`

Here, `<DMA_SERVER>` is the fully qualified host name of your HP DMA server.

HP Server Automation Documentation

The latest versions of SA documents are available on the HP Software Support web site:

<https://softwaresupport.hp.com/>

HP Live Network connector Documentation

The following documents are included in the HP Live Network connector documentation library:

- *HP Live Network connector User Guide*
- *LNC Release Notes*

The latest versions of these documents are available on the HP Live Network web site:

1. Go to the following HP Live Network connector page:

<https://hpln.hp.com/group/hp-live-network-connector>

2. Click the RESOURCES link.
3. Open Resources.
4. Open the Documentation folder.
5. Download the latest version of the documents.

Note: You must sign in to HP Live Network using your HP Passport credentials. (See [Support](#) on page 2 for more information about obtaining an HP Passport account.)

Glossary

A

automation items

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

B

bridged execution

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

C

capability

Capabilities are collections of related privileges. There are three capabilities defined in HP DMA. Login Access capability enables a user to log in to the web interface. This capability does not guarantee that this user can view any organizations or automation items—permissions are required to access those items. Workflow Creator capability

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

connector

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

cross-platform

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

custom field

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

D

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.

F

function

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

I

input parameters

A workflow has a set of required parameters for which you must specify a value. The required parameters are a subset of all the parameters associated with that workflow. The remaining parameters are considered optional. You can specify a value for an optional parameter by first exposing it using the workflow editor and then specifying the value when you create a deployment.

M

mapping

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

O

organization

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

P

parameters

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

policy

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

R

raw devices

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

operations from the user address space to the operating system kernel buffers.

role

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

S

smart group

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

software repository

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

solution pack

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

steps

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

T

target instance

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

W

workflow

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

workflow editor

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

workflow templates

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