Database and Middleware Automation

Software Version: 10.50

Linux, Solaris, AIX, and HP-UX

Planning Guide

Document Release Date: June 2016

Software Release Date: June 2016

Hewlett Packard Enterprise

Legal Notices

Warranty

Restricted Rights Legend

Copyright Notice

Trademark Notices

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/.

This site requires that you register for an HP Passport and to sign in. To register for an HP Passport ID, click **Register** on the HPE Software Support site or click **Create an Account** 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 HPE sales representative for details.

Support

Visit the HPE Software Support site at: https://softwaresupport.hpe.com/.

Most of the support areas require that you register as an HP Passport user and to sign in. Many also require a support contract. To register for an HP Passport ID, click **Register** on the HPE Support site or click **Create an Account** on the HP Passport login page.

To find more information about access levels, go to: https://softwaresupport.hpe.com/web/softwaresupport/access-levels.

HPE Software Solutions Now accesses the HPSW Solution and Integration Portal website. This site enables you to explore HPE Product Solutions to meet your business needs, includes a full list of Integrations between HPE Products, as well as a listing of ITIL Processes. The URL for this website is https://softwaresupport.hpe.com/km/KM01702731.

Contents

Contents

Planning

This section provides resources to help you plan using HPE DMA.

Topic	Description
"Key concepts" on page 8	Provides conceptual information on HPE DMA
"Requirements" on page 66	Lists minimum system requirements to install and work with DMA.
"Support matrix" on page 71	Lists workflow support information.
"Performance and sizing" on page 72	Lists minimum hardware requirements to install and work with HPE DMA.
"Roles, Capabilities, and Permissions" on page 75	Outlines various user privileges.

Key concepts

Database and Middleware Automation (DMA) software automates tasks like provisioning and configuring, 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.

Automating these tasks enables greater efficiency and faster change delivery with higher quality and better predictability. 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, configuring, and compliance checking.
- Subject matter experts can define, enforce, and audit full stack automation across network, storage, server, database, and middleware.

An 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 DMA parameters. Related DMA workflows are grouped together in solution packs.

DMA Foundational Concepts

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

DMA Platform

The platform is the foundation on which 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.

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.

DMA objects

This section describes the DMA objects. Objects are basic to understanding DMA.

The objects fall into the following categories:

- "Automation objects"
- "Environment objects" on page 16
- "Bridge objects" on page 26
- "Connector object"

Automation objects

The umbrella term automation objects refers to those items to which role-based permissions can be assigned. The DMA automation objects include:

- "Workflows"
- "Steps"
- "Functions"
- "Parameters"
- "Solution Packs"

Workflows

In 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 DMA, and 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 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.

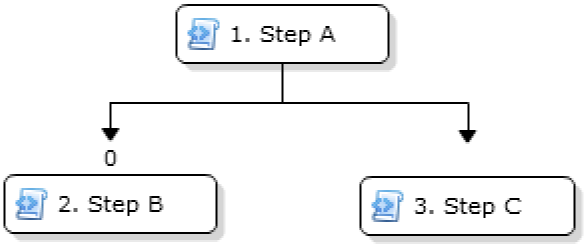

If there are two or more follow-on steps that require the same exit code from a previous step, then all of those steps will be executed..

Parameter mappings

The workflow also contains parameter mappings between the steps. For more information about parameters, see "Parameters"

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

- User-selected values for deployment parameters

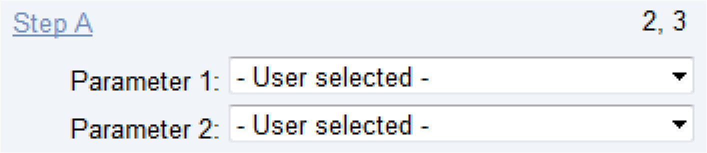

Step A 2, 3

Parameter 1: - User selected -

Parameter 2: - User selected -

- Output parameters from the previous step

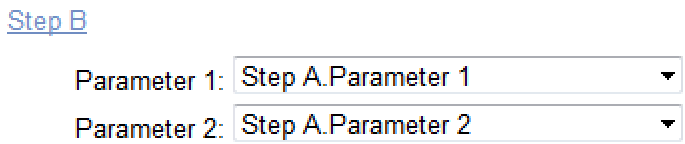

Step B

Parameter 1: Step A.Parameter 1

Parameter 2: Step A.Parameter 2

- DMA built-in user

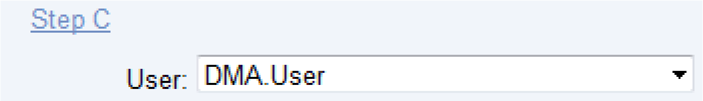

Step C

User: DMA.User

- Previously-defined policy parameters when the workflow is deployed. For more information about policies, see "Policies".

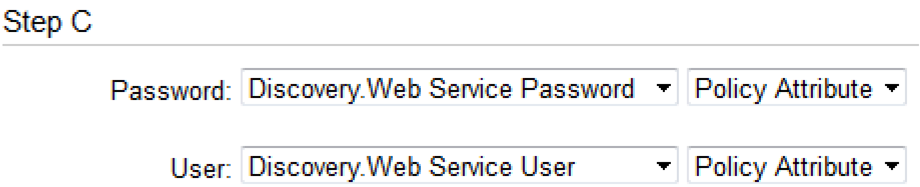

Step C

Password: Discovery.Web Service Password Policy Attribute

User: Discovery.Web Service User Policy Attribute

- Previously-defined Custom Field (see "Custom Fields") when the workflow is deployed

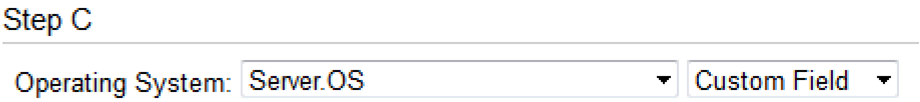

Step C

Operating System: Server.OS Custom Field

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.

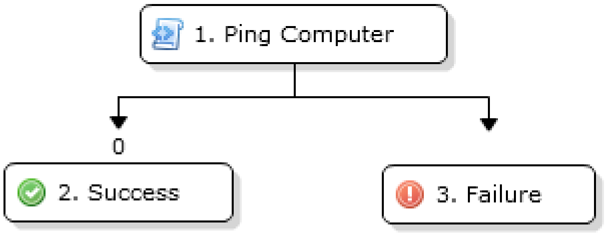

Source of workflows

Workflows are supplied in solution packs (see "Solution Packs"), but the user must copy the workflows before using them. 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 DMA-supplied workflows and the change history for custom workflows.

Versioning

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

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 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".

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 DMA-supplied steps and the change history for custom steps.

Versioning

In the course of DMA development, improvements to steps could cause existing deployments and customizations to break (for example, by adding or removing parameters). In this case, 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 and other functions. 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. 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 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 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 DMA-supplied functions and the change history for custom functions.

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 template. 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 DMA, you are granted access to download solution packs.

When a solution pack is imported, the entire solution pack is inserted into the DMA back-end database (see "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 DMA repository.

The 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.

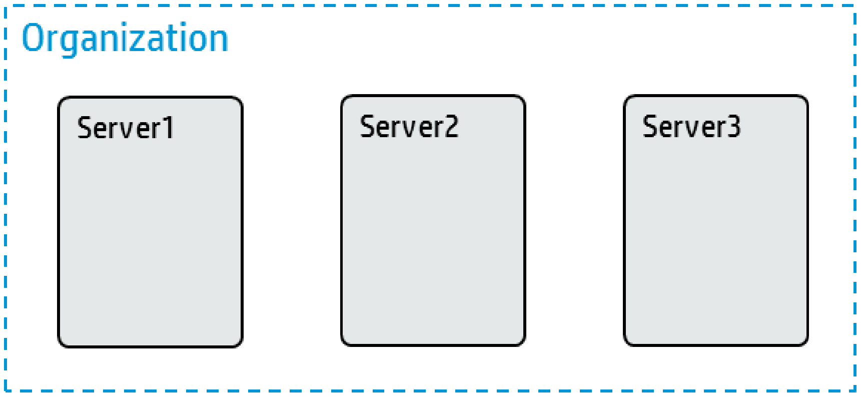

You can use organizations to separate development, staging, and production resources—or to separate logical business units. Because user security for running workflows is defined at the organization level, organizations should be composed with security in mind.

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 DMA is an Server Automation (SA) Managed Server (see "Managed servers").

A server can only belong to a single organization.

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

To be added to 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 Integrating HPE DMA with SA topic in the Installation Guide.

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 *DMA Support Matrix*, available at: https://softwaresupport.hpe.com

Instances

An Instance in 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 `sqlservr.exe` 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 DMA. Every server belongs to one organization. A server can have—but does not need to have—host instances.

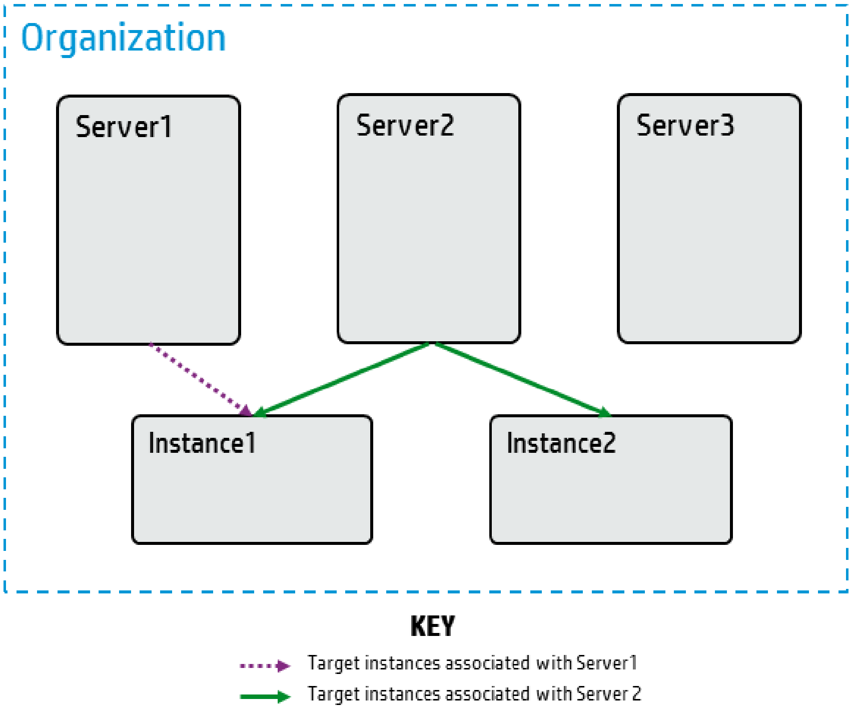
Organization
Server1
Server2
Server3
Instance1
Instance2
KEY
Target instances associated with Server1
Target instances associated with Server 2

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 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 DMA: application servers, node agents, node managers, clusters, and web servers.

The following diagram shows the relationships between organizations, servers, instances, and databases within 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:

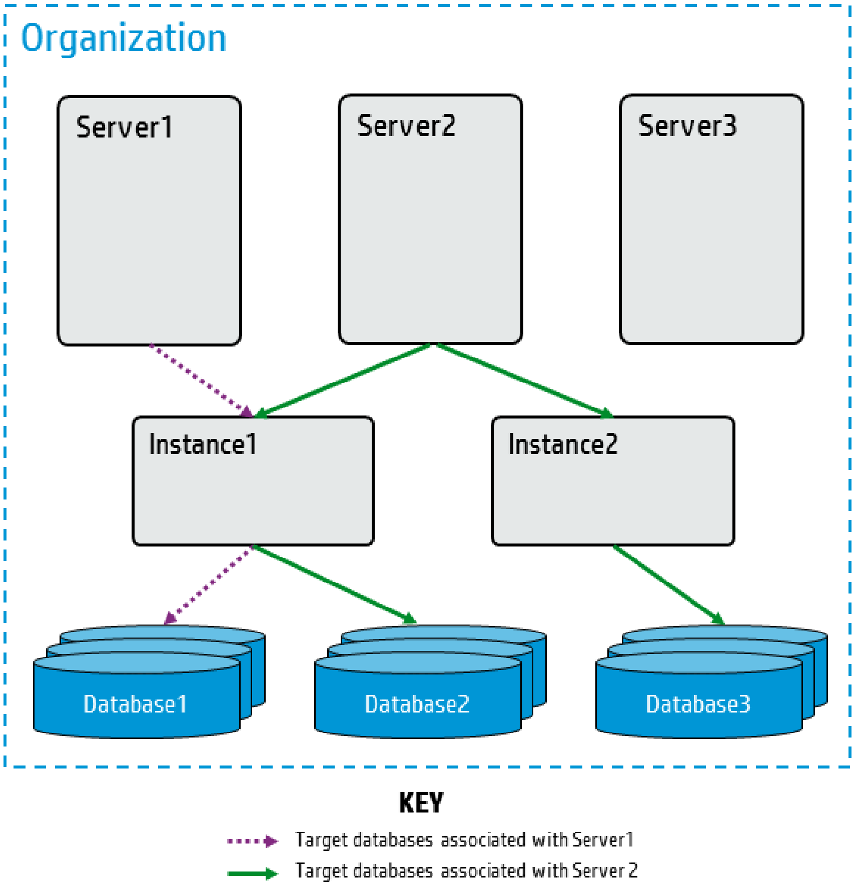
Organization
Server1
Server2
Server3
Instance1
Instance2
Database1
Database2
Database3
KEY
Target databases associated with Server1
Target databases associated with Server 2

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.

The following are the types of Custom Fields:

Custom Field Type	Description
Text	This type of Custom Field stores text. The size of the column in the database is 1000 bytes.
Multi-lineText	This type of Custom Field also stores text. In the user interface the user has a larger input box and can enter multiple lines.
Password	This 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	This 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 table 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 permissions assigned to a role 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 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:

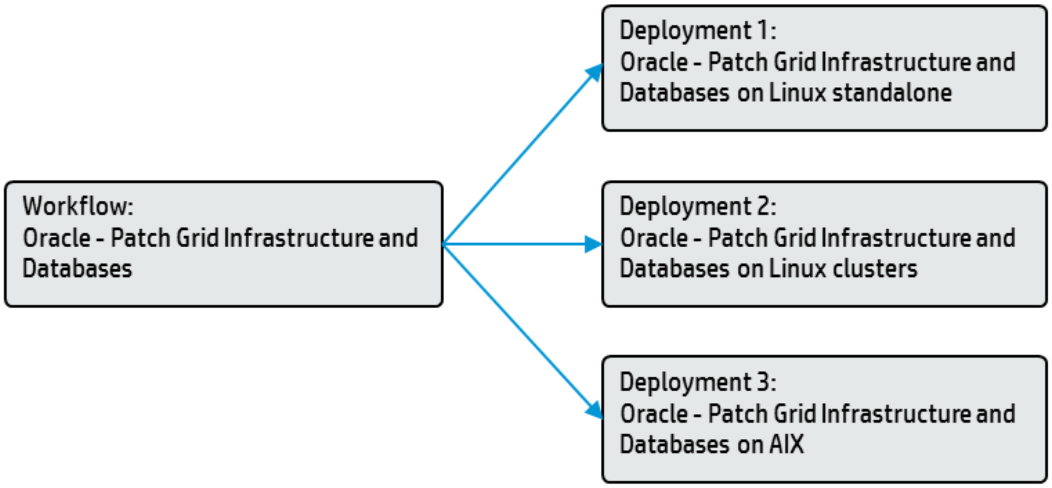
Workflow:
Oracle - Patch Grid Infrastructure and Databases
Deployment 1:
Oracle - Patch Grid Infrastructure and Databases on Linux standalone
Deployment 2:
Oracle - Patch Grid Infrastructure and Databases on Linux clusters
Deployment 3:
Oracle - Patch Grid Infrastructure and Databases on AIX

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:

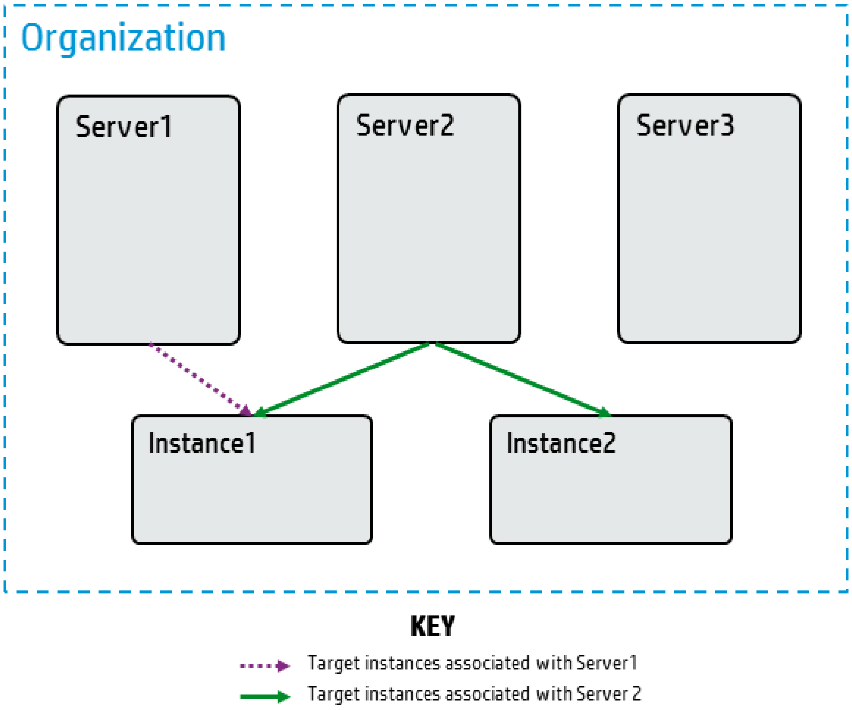
Organization
Server1
Server2
Server3
Instance1
Instance2
KEY
Target instances associated with Server1
Target instances associated with Server 2

You can only add specific targets if they have Write permission on the deployment and Deploy permission on a target's Environment (see "Permissions model").

You can only execute the deployment if you have Execute permission on the deployment and Read permission on the target's Environment. 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 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

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

Many 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 qXYZ.

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:

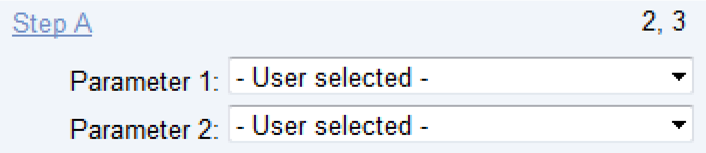

Parameter 1: - User selected -
Parameter 2: - User selected -

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.

DMA handles scheduled deployments as follows:

- The scheduled deployment runs in the time zone of the 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 enables you 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 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 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 DMA solution pack, but you can set their values as appropriate for your environment. An 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

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 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 component enables DMA to communicate with Server Automation. You must configure the Connector before you 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

Initially, the initial DMA administrator (see "Special User: DMA Initial Administrator") 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)
- SA username
- SA password to connect to SA

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

The following illustration shows how DMA connects to Server Automation:

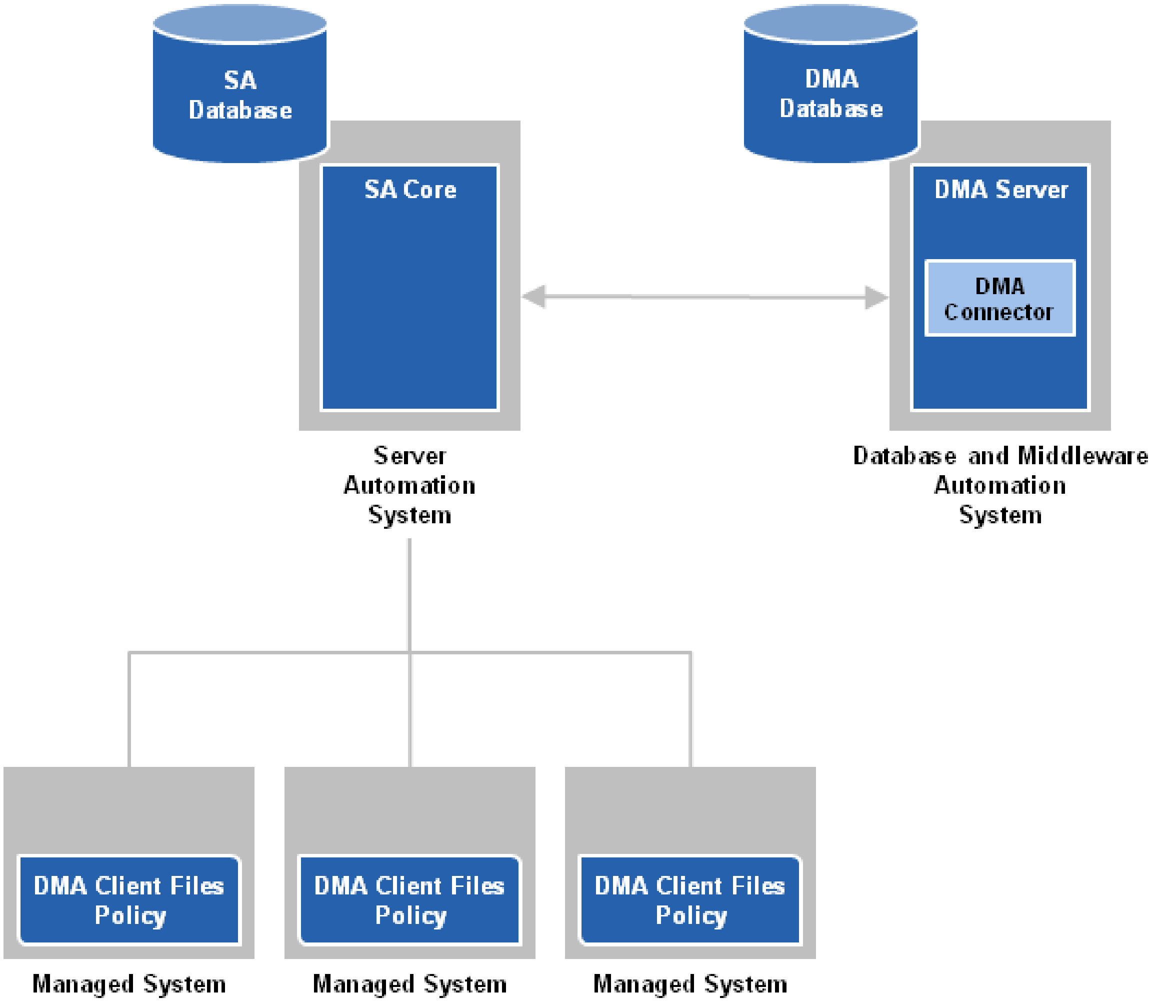

.

The Connector is added and configured when you install DMA (see *DMA Installation Guide*).

Caution: If you change the location or configuration of Server Automation, you may need to copy the JAR files and reconfigure the Connector.

If you switch the DMA Server to a different SA Core, the Connector needs to be reconfigured.

If the new SA Core is part of the same SA mesh, the same SA database is available. To complete the switch, follow the instructions in "DMA is Switched to Different SA Core" in the *DMA Troubleshooting Guide* .

It is not recommended to switch the DMA Server to an SA Core that is not part of the same SA mesh. The recommended solution is to install a new DMA Server. Follow the instructions in the "How to Install DMA" section in the *DMA Installation Guide*. To move your workflows from the old DMA Server to the new server, use the Promote workflows that are described in the *DMA Promote User Guide* .

Platform architecture

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

- "DMA Server"
- "SA Server"
- "Database"
- "Managed servers"

Architecture

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

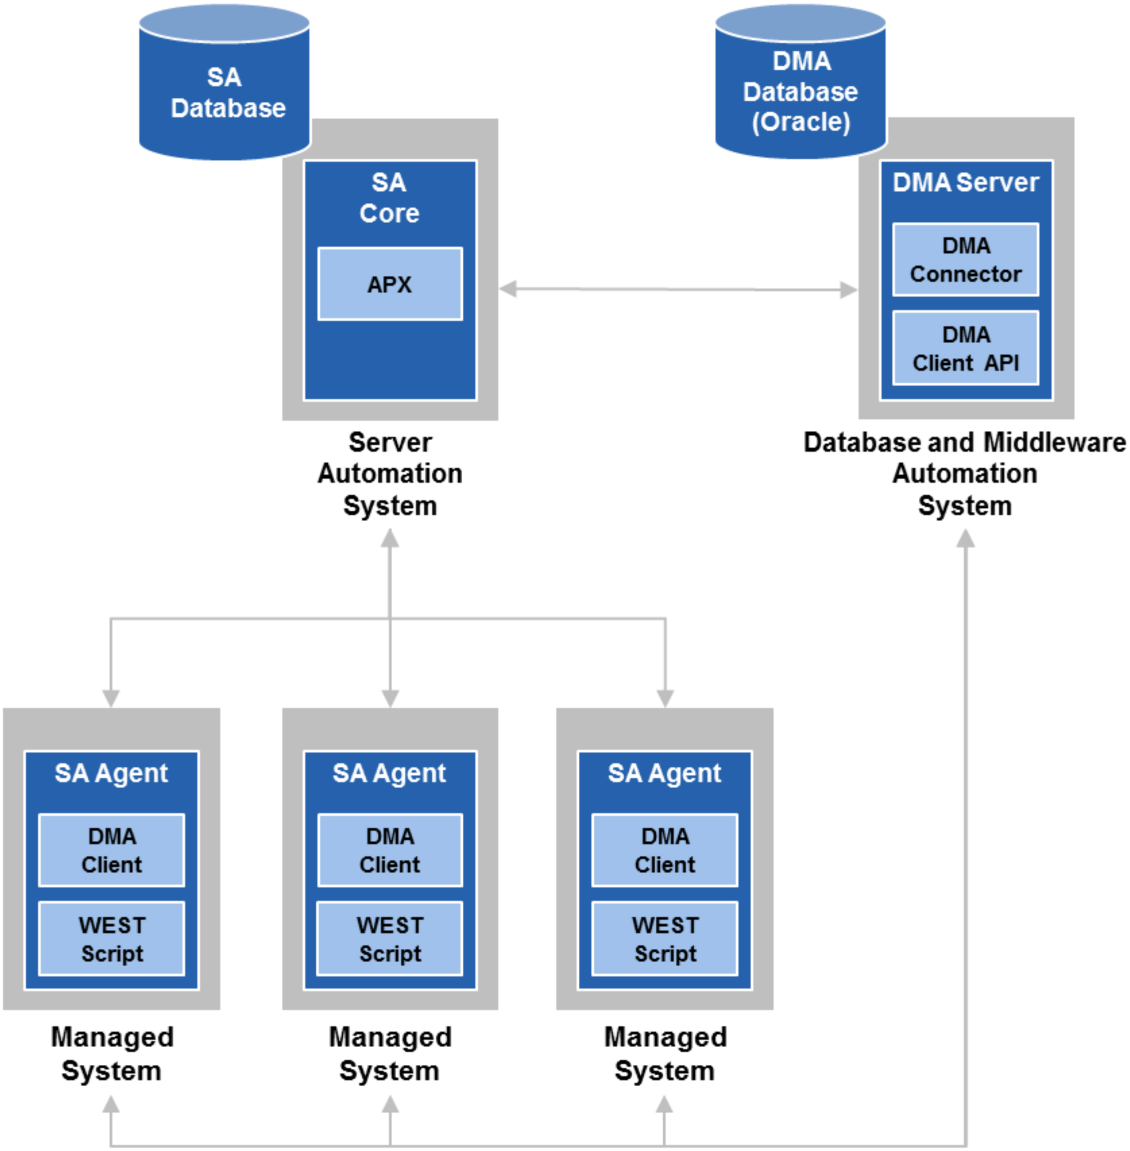

> **Tip:** For more information about how the DMA components interact, see "How the DMA Platform works".

DMA Server

The 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 *DMA Installation Guide*, available at: https://softwaresupport.hpe.com

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

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

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

DMA Servers can be clustered. You can connect multiple DMA Servers to a single back-end database (see "Database").

Tomcat

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 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 *DMA Installation Guide* available at: https://softwaresupport.hpe.com .

SA Server

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

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 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.

Integrating with Server Automation

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

> **Tip:** For more information, see "Integrating DMA with Server Automation" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com .

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

SA functionality used by DMA

DMA uses the following SA functionality:

- SA Application Programing Extension (see "Automation Platform Extension (APX)")
- DMA Client Files software policy, which includes the DMA Client (a script called by the SA Agent) and the run-time configuration necessary for DMA workflows

For information about how to integrate DMAwith Server Automation, see the Integrating DMA with SA topic in the Installation Guide.

Database

DMA uses an Oracle or a PostgreSQL as a back-end database.

The database contains:

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

> **Tip:** For more information, see "Pre-Installation Requirements" section in the *DMA Installation Guide*, available at: https://softwaresupport.hp.com/.

Managed servers

DMA requires managed servers to run. The DMA environment (see "Environment objects") contains information about the managed servers. The servers are targets for DMA deployments (see "Deployments").

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

Managed servers have the following SA attributes:

- Object ID
- Name
- Hostname

> **Tip:** For more information, see "Integrating DMA with Server Automation" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com.

How the DMA Platform works

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

- "Typical flow of DMA workflow execution"
- "Workflow Execution ScripT (WEST)"
- " REST APIs"
- "DMA baseline data"
- "Automation Platform Extension (APX)"

Typical flow of DMA workflow execution

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

> **Tip:** New terms are explained later in this section.

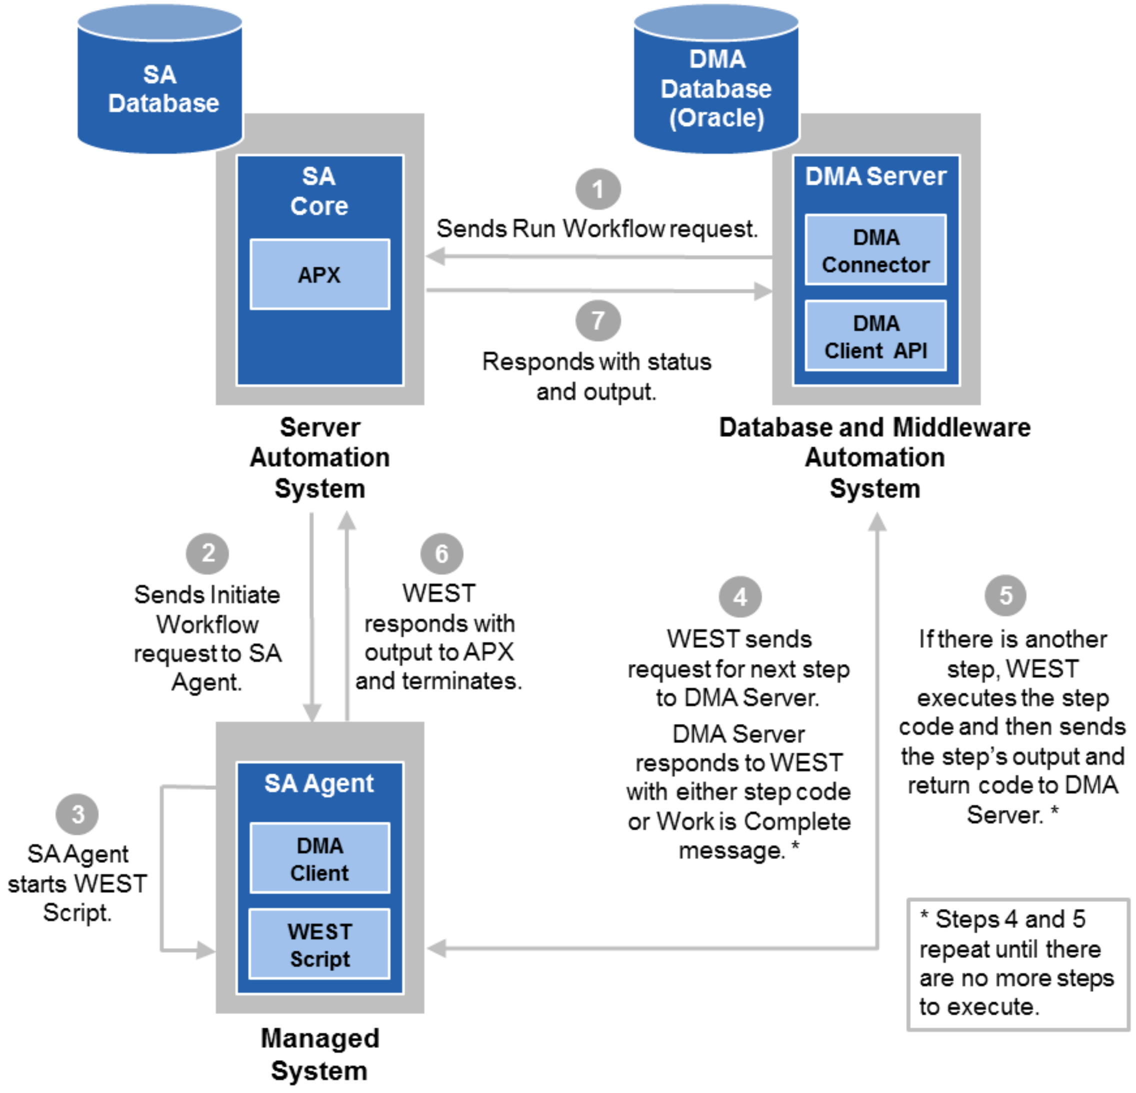
SA Database
DMA Database (Oracle)
SA Core
APX
DMA Server
DMA Connector
DMA Client API
1
Sends Run Workflow request.
7
Responds with status and output.
Server Automation System
Database and Middleware Automation System
2
Sends Initiate Workflow request to SA Agent.
6
WEST responds with output to APX and terminates.
4
WEST sends request for next step to DMA Server.
DMA Server responds to WEST with either step code or Work is Complete message. *
5
If there is another step, WEST executes the step code and then sends the step's output and return code to DMA Server. *
3
SA Agent starts WEST Script.
SA Agent
DMA Client
WEST Script
Managed System
* Steps 4 and 5 repeat until there are no more steps to execute.

> **Note:** This diagram is based on the DMA architecture diagram—"Architecture"—but with a single Managed System.

Process Description

1. The DMA Server (see"DMA Server") communicates with the SA Server (see"SA Server") to start the DMA APX (see "Automation Platform Extension (APX)").

2. The 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 DMA APX starts the DMA Client on each one.

3. The 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 DMA Server requesting the next workflow step (see "Steps").

   The 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 DMA Server responds with one of the following:

   - The 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 DMA Server responds telling WEST to stop.
   - If there currently is no more input for WEST, the 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 DMA Server. When the step has completed executing, WEST sends the step return code to the DMA Server. The information will be available in the Step Output, Step Errors, and Step Header tabs on the History page in the DMA Server UI.

6. The DMA APX gathers the output from WEST and then WEST terminates.

7. The DMA APX sends the workflow status and the output information to the DMA Server. This information will be available in the Connector Output and Connector Errors tabs on the History page in the DMA Server UI.

Workflow Execution ScripT (WEST)

Each DMA managed server (see "Managed servers") uses a program called Workflow Execution ScripT (WEST) to communicate with the DMA server (see"DMA Server"). WEST does the following:

- Executes workflow steps on the Managed Servers
- Communicates with the 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 "Integrating DMA with Server Automation" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com.

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 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 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 step types. The step is currently executing.

Deployment Execution States

State	Description
Aborted	The DMA aborted the deployment due to an unexpected or unrecoverable error condition.
Cancelled	The 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 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. The default value is 262144 characters.
`west_timeout`	Time out in seconds for HTTP messages. The default value is 60 seconds.
`west_retries`	Number of retries. The default value is unlimited.
`west_heartbeat`	Interval for heartbeat messages. The default value is 180 seconds.
`west_proxy_in_use`	Flag to indicate if a proxy is being used. The valid values are TRUE and FALSE. If specified, then `west_proxy_address` must also be specified.
`west_proxy_address`	Proxy address. If specified, `west_proxy_in_use` must also be specified. If using the SA Gateway Network, then set the value to `sa_auto_select`.
`west_verbose`	Flag to indicate whether or not to turn on debug logging in the 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 *DMA Administrator Guide* and the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com

REST APIs

The DMA REST application programming interfaces (APIs) are a collection of tools that you can use to operate 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 DMA APIs enable you to read and—in some cases—create, modify, and delete DMA objects: environment objects, automation objects, the set up, and solution packs.

> **Tip:** The *DMA API Reference WebHelp* is available on all DMA Servers at:
>
> `https://<DMA_SERVER>:8443/dma/api`
>
> Here, <*DMA_SERVER*> is the fully qualified host name of your DMA server.
>
> The documentation gives examples of the XML and lists valid methods and query parameters.

DMA baseline data

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

DMA Installation

To install 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 "Installing the DMA Server" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com.

DMA Upgrade

To upgrade DMA to a newer DMA release, you must also run `dmaBaselineData.sh`. To upgrade, 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 "Upgrading DMA" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com.

Troubleshooting

> **Tip:** For information about errors that you may encounter related to baselining, see "Common Baseline Errors" section in the *DMA Troubleshooting Guide*, available at: https://softwaresupport.hpe.com.

DMA Baseline Options

> **Tip:** For a complete list of baseline options, see "DMA Baseline Options" section in the *DMA Installation Guide* available at: https://softwaresupport.hpe.com.

Automation Platform Extension (APX)

DMA needs Server Automation APXs (Automation Platfrom Extensions) to start WEST (see "Workflow Execution ScripT (WEST)") and to run a workflow (see "Deployment Runs"). The APX layer allows DMA to add users and roles to SA so that SA can communicate with 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 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 "Importing the DMA APX" section in the *DMA Installation Guide*. The guide describes how to load the APXs depending on the version of 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" section in the *DMA Troubleshooting Guide*.

For 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: https://softwaresupport.hp.com/.

Permissions model

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

- "Roles"
- "Capabilities"
- "DMA user"
- "Permissions"

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

Obtained from SA

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

Commonly defined roles are:

- DMA Admins
- DMA Users
- DMA Workflow Developers

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

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

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

Made Available in DMA

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

The available roles are saved in the DMA database (see "Database").

A role can have multiple capabilities (see "Capabilities").

> **Tip:** For more information, see "Settting SA Groups and Users" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com.

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"DMA objects").

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

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

The following are the capabilities that determine whether you can access DMA and what you can do within the DMA UI:

DMA Capabilities

Capability Name	Description
Login Access	This permission enables you to login in to DMA. With this permission you can: • View organizations for which you have Read access. • Edit organizations and associated target objects for which you have Write access. • Run workflows against targets in organizations for which you have Deploy access. There are additional permissions for specific automation items (see Permissions).
Workflow Creator	This permission enables you to create or copy DMA workflows (see "Workflows"). Each workflow also has its own Read and Write permissions.
Administrator	This permission enables you to act as the DMA administrator. The Administrator capability is (in most companies) synonymous with the DMA Admins role, yet multiple roles can be Administrators. With this permission you can: • Access the Setup page in the DMA UI (see the *DMA Administrator Guide* for more information). • Create or modify any DMA organization.

DMA Capabilities, continued

Capability Name	Description
	- Grant users (roles) access to specific workflows, steps,deployments, policies, and organizations. - Configure the Outgoing Email settings. - Create workflows.

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

> **Tip:** For more information, see the "Settting SA Groups and Users", "Starting DMA" and "Setup DMA" sections in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com.

DMA user

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

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

DMA calls the Connector (see "Connector object") regularly to re-evaluate 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 DMA

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

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

> **Tip:** For more information, see the "Set Up the SA Groups and Users" section in the DMA *Installation Guide*.

Special User: DMA Initial Administrator

The DMA initial administrator accomplishes the initial configuration of the Connector (see "Connector object"). Since the Connector has to be configured by a user with the DMA Admin role (see "Roles") and the connector determines the DMA Admins user for DMA. DMA requires a specific user to log in and perform the initial configuration when DMA is started after the initial install. Effectively, this is the only DMA user. This user is always `dma_initial_admin`.

The DMA initial administrator performs the following tasks:

- Sets the initial password
- Configures the Connector (see "Connector object")
- Registers the DMA Roles (see "Roles")
- Assigns the DMA Capabilities (see "Capabilities")

Other than performing the above tasks, the DMA initial administrator can be used in the following situations:

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

> **Tip:** For more information, see the "Start DMA" section in the DMA *Installation Guide*.

Special User: DMA Connector User

A single SA user—the DMA connector user—is required to configure the DMA Connector (see "Connector object") to SA (see"SA Server"). This user is used by DMA to connect to SA whenever a specific personalized SA account cannot be used—for example, to verify

a login is allowed.

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

> **Tip:** For more information, see the "Set up DMA" section in the *DMA Installation Guide*, available at: https://softwaresupport.hpe.com

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 on 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 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—inherit 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. 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 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	*Not delivered in solution packs.*	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 XYZ 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.

Special types of workflows

This section describes two unique types of 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. For re-targetable steps, the Targetable checkbox is selected to indicate they are 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 is different from running a regular DMA workflow. At runtime, you can specify the different targets, when prompted. Internally, when you run the workflow, DMA starts WEST (see "Workflow Execution ScripT (WEST)") on all of the target servers. WEST on each target server repeatedly "asks" the DMA server (see "DMA Server") if there is anything to do. The DMA server orchestrates the steps so that only one step runs at a time. For more information see "Typical flow of DMA workflow execution".

If you watch the DMA Console while the bridged execution workflow is executing, you will see 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 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 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:

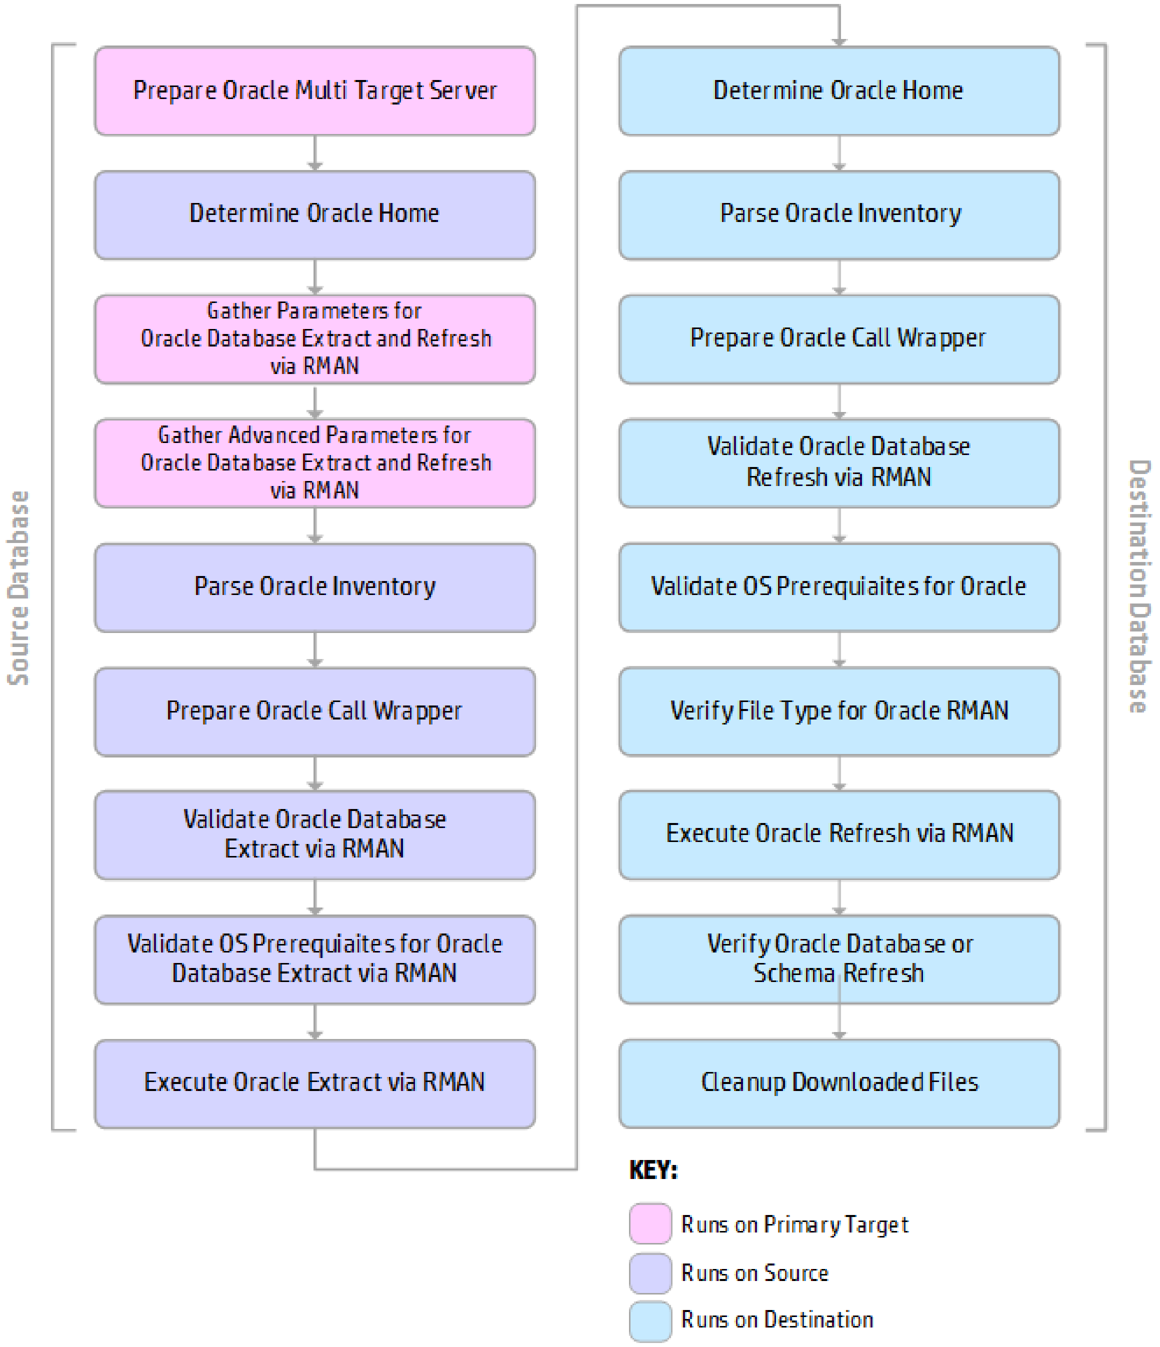
Source Database
Prepare Oracle Multi Target Server
Determine Oracle Home
Gather Parameters for Oracle Database Extract and Refresh via RMAN
Gather Advanced Parameters for Oracle Database Extract and Refresh via RMAN
Parse Oracle Inventory
Prepare Oracle Call Wrapper
Validate Oracle Database Extract via RMAN
Validate OS Prerequiaites for Oracle Database Extract via RMAN
Execute Oracle Extract via RMAN
Destination Database
Determine Oracle Home
Parse Oracle Inventory
Prepare Oracle Call Wrapper
Validate Oracle Database Refresh via RMAN
Validate OS Prerequiaites for Oracle
Verify File Type for Oracle RMAN
Execute Oracle Refresh via RMAN
Verify Oracle Database or Schema Refresh
Cleanup Downloaded Files
KEY:
Runs on Primary Target
Runs on Source
Runs on Destination

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.

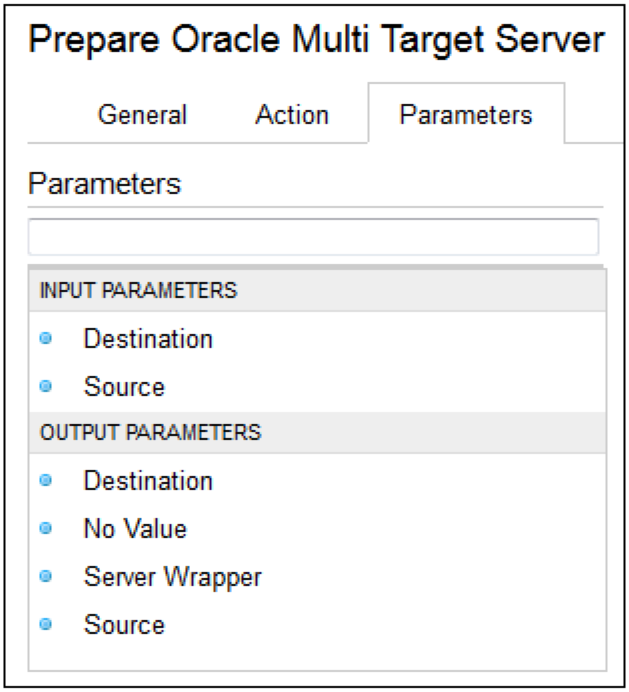

Prepare Oracle Multi Target Server

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:

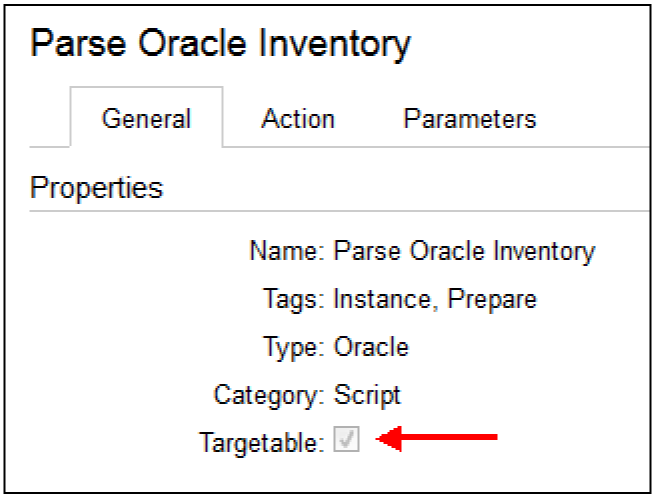

Parse Oracle Inventory

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:

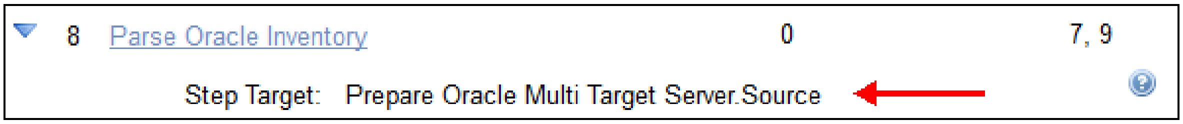

8 Parse Oracle Inventory 0 7, 9

Step Target: Prepare Oracle Multi Target Server.Source

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

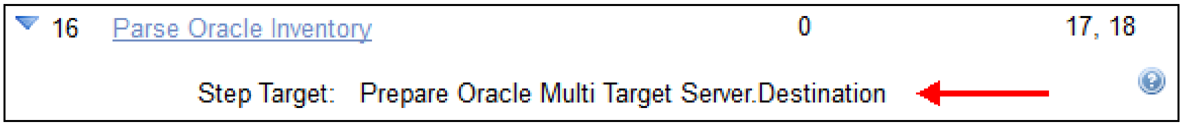

16 Parse Oracle Inventory 0 17, 18

Step Target: Prepare Oracle Multi Target Server.Destination

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

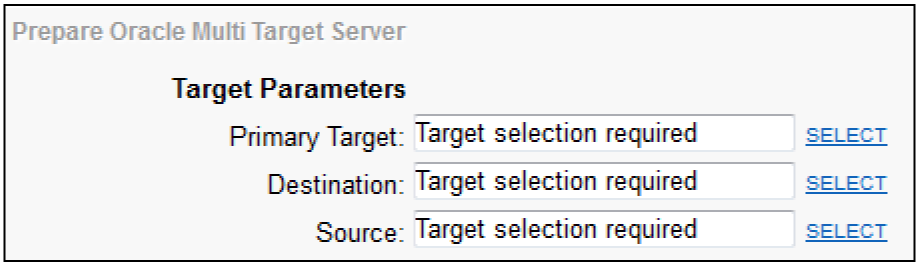

Prepare Oracle Multi Target Server

Target Parameters

Primary Target: Target selection required SELECT

Destination: Target selection required SELECT

Source: Target selection required SELECT

Parameter Name	Description
Source	Instance that contains the database whose contents will be exported.
Destination	Instance where the database will be imported.

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

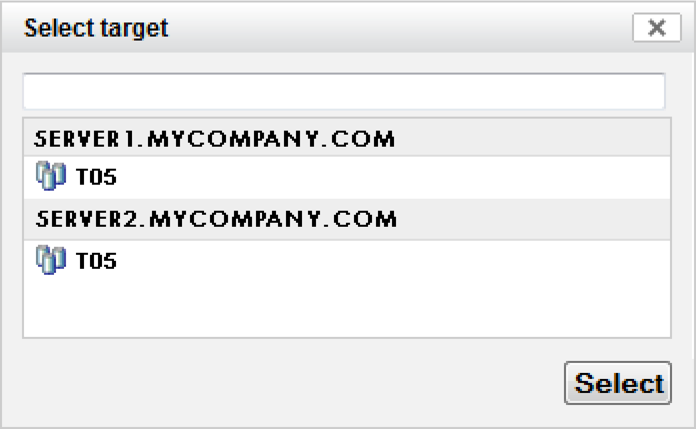

Select target

SERVER1.MYCOMPANY.COM

T05

SERVER2.MYCOMPANY.COM

T05

DMA Bridged Execution Workflows

The following bridged execution workflows are delivered in 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

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.

When the master workflow is executed, the deployment for the master workflow starts first. 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 in the following figure, 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.

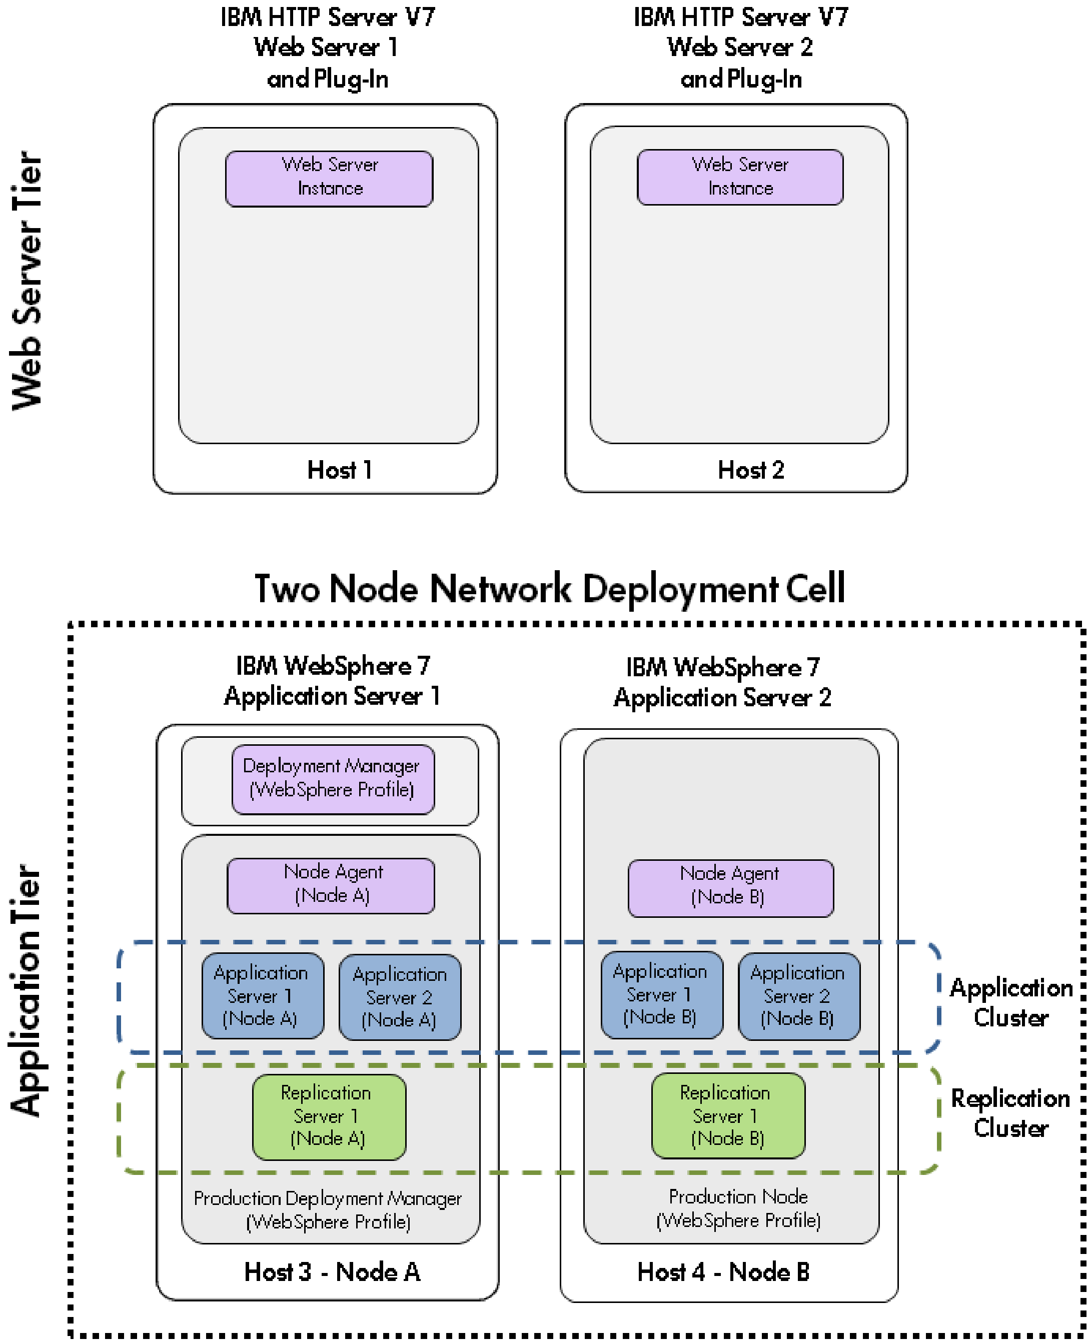

You can accomplish this architecture by running five DMA workflows (found in the DMA Application Server Provisioning Solution Pack) in a orderly, repeatable manner:

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 DMA Application Server Provisioning Solution Pack) creates the entire architecture. The following diagram shows all the steps in the workflow. The subflow steps—steps that invoke the component workflows—are in blue:

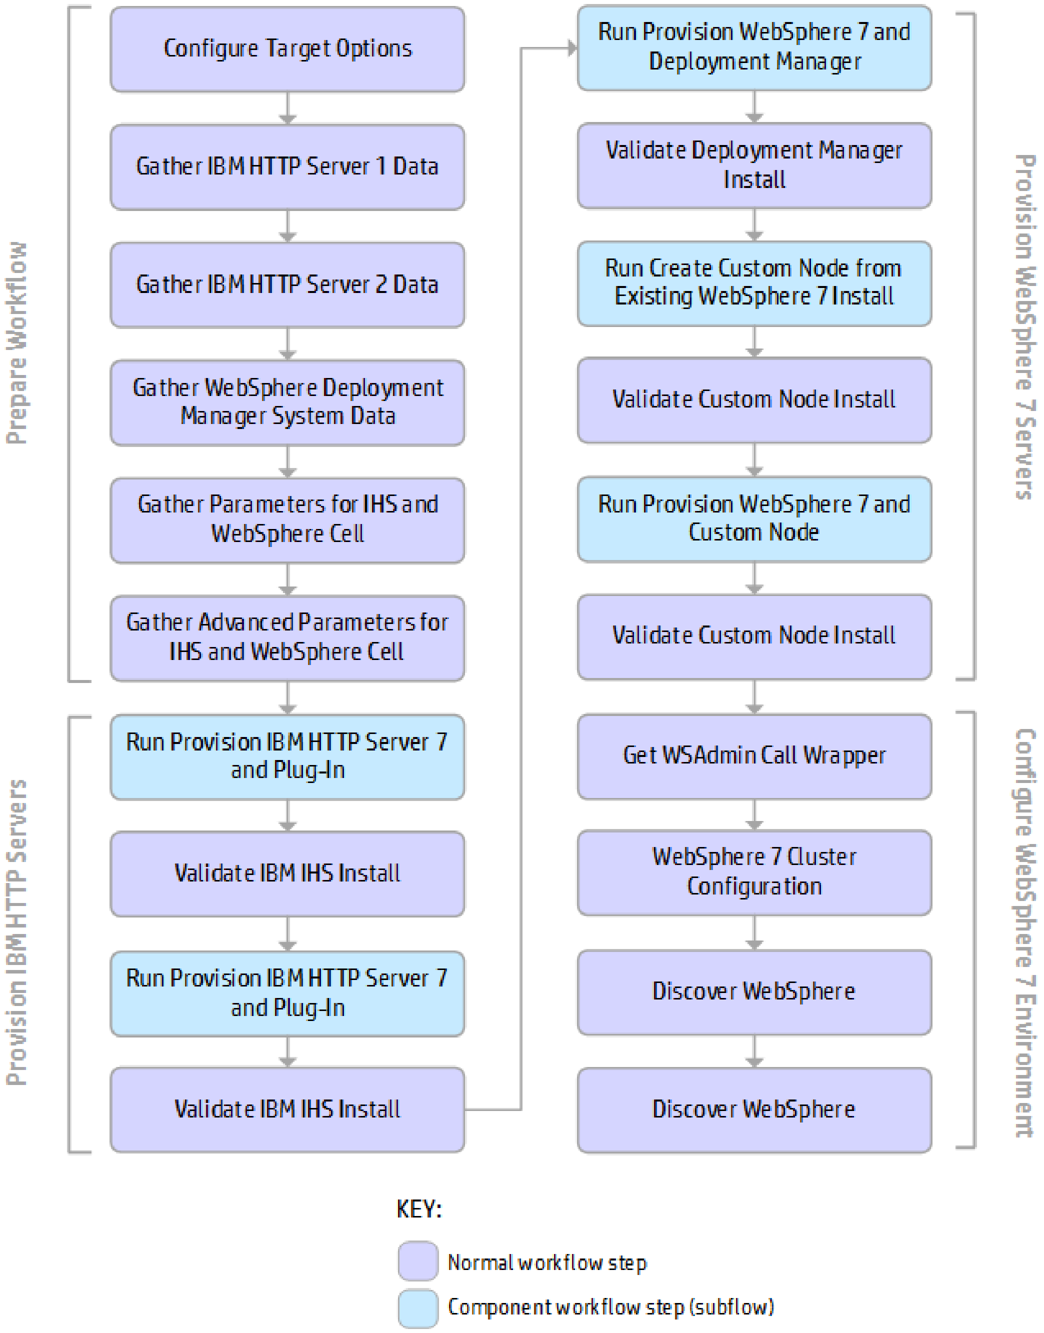
Prepare Workflow
Configure Target Options
Gather IBM HTTP Server 1 Data
Gather IBM HTTP Server 2 Data
Gather WebSphere Deployment Manager System Data
Gather Parameters for IHS and WebSphere Cell
Gather Advanced Parameters for IHS and WebSphere Cell
Provision IBM HTTP Servers
Run Provision IBM HTTP Server 7 and Plug-In
Validate IBM IHS Install
Run Provision IBM HTTP Server 7 and Plug-In
Validate IBM IHS Install
Provision WebSphere 7 Servers
Run Provision WebSphere 7 and Deployment Manager
Validate Deployment Manager Install
Run Create Custom Node from Existing WebSphere 7 Install
Validate Custom Node Install
Run Provision WebSphere 7 and Custom Node
Validate Custom Node Install
Configure WebSphere 7 Environment
Get WSAdmin Call Wrapper
WebSphere 7 Cluster Configuration
Discover WebSphere
Discover WebSphere
KEY:
Normal workflow step
Component workflow step (subflow)

DMA Master Workflows

The following master workflows are delivered in 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

Requirements

This section provides information about the supported hardware and software that you must have in order to successfully install and run DMA.

Platform requirements

The DMA version 10.50.001.000 server requires the following server platform infrastructure:

Server Platform Infrastructure	Product	Version
DMA Server platform	Red Hat Enterprise Linux	5.8, 6.1, 7 (or later), 64-bit
	SUSE Enterprise Linux	11 (or later) 64-bit
Server Management Tool	Server Automation	Server Automation Ultimate Edition 10.1x and 10.2x (SA 10.1x and 10.2x)
		Server Automation Enterprise Edition 10.1x and 10.2x (SA 10.0x and 10.2x)
		Server Automation 10.0x Standard Edition (SAVA 10.0x)
	DCA Virtual Appliance	DCA Virtual Appliance 2016.01
DMA Backend Database Tool	Oracle Database Enterprise or Standard Edition	12.1.0.2 (container DB)
		11gR2
	PostgreSQL	9.3.5
Oracle Java	Java Runtime Environment (JRE)	Version 8 Update 92 (1.8_u92) on Linux, Windows, Solaris Version 8 Update 5 (1.8_u5) on HPUX

Server Platform Infrastructure	Product	Version
		Version 8 Update 130 (1.8_u130) on AIX

> **Note:** Although DMA works on other Linux operating systems, supports these certified versions.

Hardware requirements

See the "Performance and sizing" on page 72.

> **Note:** DMA is fully supported to be installed and run on VMware versions 5 and 5.1 virtual machines.

Software requirements

- Server Automation, any of the following versions:

  - Ultimate Edition 10.1x, 10.2x (SA 10.1x, 10.2x)

  - Standard Edition 10.10 (SAVA 10.10)

  - Enterprise Edition 10.0x, 10.2x (SA 10.0x, 10.2x)

  > **Note:** You must purchase this license separately.

- Oracle Database Enterprise Edition/PostgreSQL 9.3.5

  > **Note:** does not provide the Oracle Database and the PostgreSQL license to run DMA.

  > **Tip:** If you plan to co-locate DMA with SA 10.10 (or later)—which uses Oracle 12c—set up DMA to also use Oracle 12c (to only require a single version of Oracle).

- Java Runtime Environment (JRE) Version 7 Update 80 (1.7u80) and Version 7 Update 85 (1.7u85). The JRE 1.7u85 is supported only in AIX and HPUX platforms.

Server Automation (SA)

Server Automation needs to be up and running.

The person who integrates DMA with SA—probably your SA administrator—needs the following:

- Root access to the SA server
- Ability to create users, groups, and permissions
- OGSH (SA Global Shell) access

This person should have the highest possible administrative rights. Although these rights may not be needed for all steps, they will help the process go smoothly.

Servers

DMA and SA can run on the same server (OS instance).

DMA and SA can use the same Oracle/PostgreSQL installation and database, but each product needs to be configured in separate schemas.

Ports

The following table provides a list of default ports used in DMA. You can configure different ports as required:

Server/Database	Port (default)
DMA	8443
Oracle Database	1521
PostgreSQL Database	5432
SA	443

Firewalls

The firewalls need to have the following ports open:

- Incoming on the port configured for DMA, 8443 is the default port—or a proxy server can be used.
- Outgoing on the ports configured for Oracle Database/PostgreSQL, and SA.

> **Tip:** For more information about how to set up a proxy server with DMA, see Using a proxy server

> section in Administration Guide.

If you are configuring DMA server with an IPv6 address, ensure that the firewall (for IPv6 traffic) is turned off or configured to do the following:

- Allow bi-directional communication on port 443 between DMA server and SA server
- Allow incoming communication to DMA server from all the DMA target servers that connect directly, on port 8443 (or whatever port DMA is server is configured to run on)
- Allow bidirectional communication between the DMA target servers and their respective SA Satellite proxy servers, on port 443.

Note: DNS resolution must be enabled across the infrastructure for IPv6 address resolution.

Note: If the default hostname does not resolve an IPv6 address, use the *-dmah* option while using the dmaBaseline.sh command, after installing DMA 10.40.

Privileges

To install packages on all UNIX® machines you must log on as a user that has root access.

Supported browsers for DMA

The DMA web UI supports the following browsers:

- Chrome
- Firefox
- Internet Explorer 10 and 11

Supported DMA target platforms

DMA supports managed target servers that use the following operating systems:

- Linux
- Solaris

- AIX
- Windows
- HP-UX

For details regarding the operating systems and versions supported by each DMA workflow, see the DMA Support Matrix available at https://softwaresupport.hpe.com/. See Documentation Updates for information about accessing this website.

Support matrix

This section provides information about the supported hardware and software that you must have in order to successfully install and run DMA.

- "Requirements" on page 66
- Workflow support matrix
- "Performance and sizing" on page 72

Performance and sizing

This topic provides the sizing recommendations for the DMA hardware and infrastructure and also for the DMA Client.

Hardware and Infrastructure Sizing

This section suggests deployment sizing guidelines to help you decide the hardware and infrastructure that you need to deploy DMA in your environment. This section lsist the minimum recommended CPU count, RAM, and disk space for the DMA server and the DMA database server—the server that houses your Oracle/PostgreSQL database.

> **Tip:** This topic does not give sizing recommendations for Server Automation (SA). The assumption is that SA is already up and running in your environment.

DMA Deployment Modes

DMA supports the following deployment options:

- Single Server: Install both the DMA server and the DMA database on a single server
- Dual Server: Install DMA on one server and create the DMA database on a separate server

Deployment sizing categories

Category	Number of DMA Clients
Small	<100
Medium	<500
Large	1,500+

> **Note:** The number of clients is not an exact measure for sizing. Sizing depends greatly on what you do with the operational system.

Recommended Sizing

The following table describes sizing suggestions for deploying the DMA server:

Sizing recommendations for deploying the DMA server

Category	Number of CPUs (2.66 GHz)	RAM	Disk Space
Small	1	4 GB	25 GB
Medium	2	8 GB	50 GB
Large	4	16 GB	100 GB

> **Note:** The recommendations are minimum requirements for the installation. These recommendations are based on dual core installation.
>
> If you install DMA on a virtual machine you must ensure that the actual available CPUs and RAM for the DMA server virtual machine meets the same requirements.

The following table describes sizing suggestions for deploying the DMA DMA database component:

Sizing recommendations for deploying the DMA database server

Category	Number of CPUs (2.66 GHz)	RAM	Disk Space
Small	4	4 GB	50 GB
Medium	4	8 GB	100 GB
Large	4	16 GB	250 GB

> **Note:** When considering sizing for these types of deployments, each sizing recommendation should be considered independently of whether or not the components are installed on the same server or on different servers. In other words, these sizing recommendations are additive.
>
> If you install the DMA database on a virtual machine you must ensure that the actual available CPUs and RAM for the DMA database virtual machine meets the same requirements.

DMA Client Sizing

This section suggests sizing guidelines for the DMA Client. The DMA Client is installed on each DMA Managed Server. The DMA Client consists of the software modules used by DMA to initiate and control workflow executions on the managed server, as well as the runtime software required for the DMA workflows.

The disk space required for the DMA Client depends on the number of workflow executions that are planned and whether the managed server will be used as an DMA development target. Thus, the required disk space is not fixed.

The following table outlines what you should consider to size the DMA Client's disk space correctly:

Directory	Description	Size
`/opt/hp/dma`	Contains the software modules used to initiate and run DMA workflows. This directory only contains static content.	For the current release, the required disk space is about 0.4 GB. The actual size varies slightly depending on the operating system of the managed server.
`/var/opt/hp/dma`	Not used in the current release. Will be used in future releases.	
`/var/tmp/dma`	Contains temporary files needed during workflow execution, such as step and function code.	The disk space required for a workflow's execution depends on the workflow and the debug level that is used. Typically, a workflow's execution requires less than 15 MB of disk space, even with the maximum debug level. Unless specifically configured to keep the temporary files, DMA will delete the temporary files upon workflow completion. The disk space for this directory can be calculated as the number of workflows running in parallel multiplied by 15 MB. Development systems—where files may be kept for debugging—require additional disk space. The additional disk space depends on the number of workflows that run in parallel and number of workflow artifacts saved for debugging.
Temporary directories	User-specified directories on the managed server that hold temporary files. Some directories may contain installation binaries or patches that are either stored on the managed server or downloaded from Server Automation. Other directories may contain extracted ZIP files. Common parameter names are: Staging Directory, Download Location, Extract Location, Archive Location, and Download Target Destination.	Adequate disk space must be available in the temporary directories to avoid workflow failures. The size depends upon which workflows will be executed on a target and whether or not temporary files are deleted upon workflow completion. Refer to the workflow documentation for disk space requirements.

Roles, Capabilities, and Permissions

This section describes the permissions you need to use and administer DMA.

> **Note**: See the *SA Administration Guide* for additional information about setting global permissions.

Access Control

DMA provides very finely grained role-based access control over the following:

- Who can access DMA
- Who can view, modify, or deploy to a specific organization
- Who can view or modify a specific workflow
- Who can create a workflow
- Who can modify a specific step
- Who can view, modify, or execute a specific deployment
- Who can view or modify a specific policy
- Who can administer DMA, including setting permissions for all these items

Roles, capabilities and permissions are the mechanisms used to establish this control. Roles are simply groups of users who have the same levels of access. Capabilities determine which DMA operations each user can perform. Permissions help you precisely manage access to automation items (workflows, steps, policies, and deployments) and organizations.

Access Control Mechanisms

There are three mechanisms that affect what you can see and access in DMA:

Roles

Each DMA user has one or more roles. Roles are used to grant users permission to log in to DMA, to determine who can create new workflows, to grant users access to specific automation items, and to determine which users have administrative privileges. Roles are defined in Server Automation. There, for example, a role is an SA group to which a user belongs. Roles must be registered in DMA before they can be used. This is done by the DMA administrator on the Role Registration page.

Capabilities

Capabilities determine whether you can access DMA, whether you can create workflows, and whether you have DMA administrator privileges. Capabilities are set by the DMA administrator. The following capabilities determine whether you can access DMA and what you can do within the DMA UI. These capabilities are assigned by the DMA administrator.

DMA Capabilities

Capability Name	Description
Login Access	This permission enables you to login in to DMA. With this permission you can: • View organizations for which you have Read access. • Edit organizations and associated target objects for which you have Write access. • Run workflows against targets in organizations for which you have Deploy access. There are additional permissions for specific automation items (see Permissions).
Workflow Creator	This permission enables you to create or copy DMA workflows (see "Workflows"). Each workflow also has its own Read and Write permissions.
Administrator	This permission enables you to act as the DMA administrator. The Administrator capability is (in most companies) synonymous with the DMA Admins role, yet multiple roles can be Administrators.

DMA Capabilities, continued

Capability Name	Description
	With this permission you can: • Access the Setup page in the DMA UI (see the *DMA Administrator Guide* for more information). • Create or modify any DMA organization. • Grant users (roles) access to specific workflows, steps,deployments, policies, and organizations. • Configure the Outgoing Email settings. • Create workflows.

Permissions

Permissions determine whether you can view, create, or modify automation items and organizations. Permissions for automation items can be set by the user who created the item or any user who has Write permission for that item. They can also be set by the DMA administrator. Permissions for organizations can only be set by the DMA administrator.

Five items have role-based permissions in DMA:

DMA Role-Based Permissions

Item	Read	Write	Execute	Deploy
Workflows	yes	yes	n/a	n/a
Deployments	yes	yes	yes	n/a
Steps	n/a	yes	n/a	n/a
Policies	yes	yes	n/a	n/a
Organizations	yes	yes	n/a	yes

> **Note**: In DMA, you only see servers that reside in organizations for which you have Read permission. In order to add a server to an organization, you must have Write permission for that organization and Login Access capability.

Permissions for each automation item (workflow, step, policy, or deployment) are set by the user who creates the item—or by any user who has Write permission for the item. They can also be set by the DMA administrator.

Permissions for organizations are set by the DMA administrator.

If you want other users to be able to access a particular item that you create, you must explicitly grant them permission to do so. You can do this on the Roles tab for that item.The following figure, for example, shows the Roles tab for a workflow.

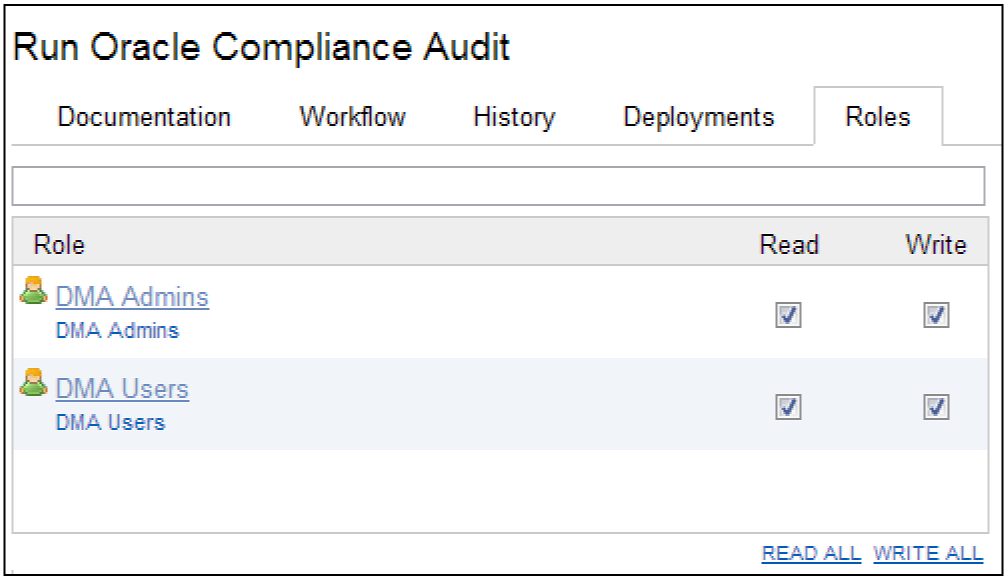

Run Oracle Compliance Audit

Role	Read	Write
DMA Admins DMA Admins	☑	☑
DMA Users DMA Users	☑	☑

Only those roles that have Login Access capability appear in the list (see "Capabilities" on page 76).

The following instructions show you how to set the permissions for a workflow. The procedure for the other types of automation items is similar.

A description of the minimum permissions needed to accomplish common DMA tasks is provided in this section. For a more comprehensive discussion of permissions, see the *DMA Administrator Guide*.

To view the roles and capabilities associated with your DMA user name, click your user name in the upper right corner:

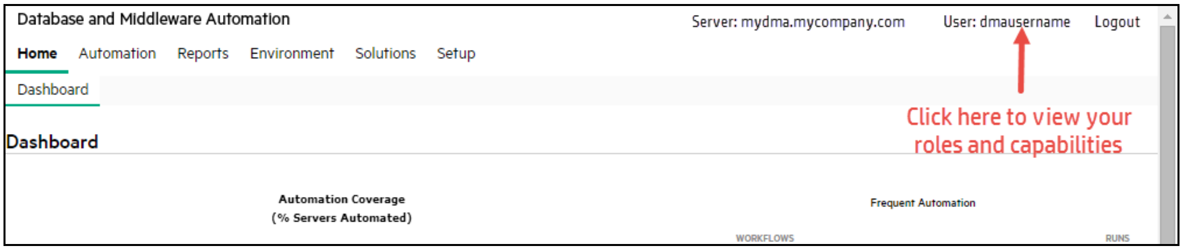

Database and Middleware Automation

Server: mydma.mycompany.com

User: dmausername

Click here to view your roles and capabilities

Send documentation feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Planning Guide (Database and Middleware Automation 10.50)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to hpe_dma_docs@hpe.com.

We appreciate your feedback!