HP Database and Middleware Automation

for the HP-UX, IBM AIX, Red Hat Enterprise Linux, Solaris, and Windows® operating systems

Software Version: 1.00

(Stratavia Data Palette version 6.0.11)

User Guide

Document Release Date: April 2011
Software Release Date: November 2010

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 Notices

Trademark Notices

iPod is a trademark of Apple Computer, Inc.

Google™ is a trademark of Google Inc.

HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

Intel and Itanium are trademarks of Intel Corporation in the U.S. and other countries.

Java® and all Java based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates in the U.S. and other countries.

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

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

UNIX® is a registered trademark of The Open Group.

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 edition.

Document Changes

Chapter	Version	Changes
All	1.00	First edition under HP name. Minor content updates throughout the manual.

Support

Visit the HP Software Support Online web site at:

www.hp.com/go/hpsoftwaresupport

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

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

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

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

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

To find more information about access levels, go to:

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

Contents

1 Product Overview

HP Database and Middleware Automation (HP DMA) automates complex, error-prone, manually-intensive yet repetitive data center automation tasks and processes. HP DMA also provides predictive analytics to detect and resolve issues before they can impact operations. HP DMA allows operations teams to automate the repair and maintenance of business critical applications, regardless of platform, version, or vendor.

HP DMA groups components into Solution Packs to address a specific customer need. Solution Packs are customizable to your unique data center environment. For detailed information on using Solution Packs within HP DMA, see Chapter 9, Solutions, on page 159 of this guide.

HP DMA also leverages existing best practices, providing centralization of knowledge base articles, scripts, and run books, allowing for standardization of standard operating procedures and the eventual automation of the processes that dominate IT today.

Being a true platform, operations teams across the enterprise are able to leverage automation to increase the efficiencies of tasks and processes including:

- Server provisioning
- Database administration
- Service requests & maintenance
- Event correlation and root cause analysis

Key Benefits of HP DMA

- Reduce operational expenditures associated with data center automation
- Automate routine IT lifecycle tasks and procedures
- Automate incident resolution for recurring alerts and performance challenges
- Predict system resource behavior, avoiding potentially costly downtime or reduced service levels

HP DMA Components

Each of the HP DMA components is separately deployable and runs on a variety of servers.

Web Server

The Web Server is the primary means of interacting with HP DMA. The Web Server is a web-based technology that fully supports all modern browsers and WAP-enabled cell phones. A streamlined interface is available for Apple iPhone and iPod touch. This architecture enables users to immediately start using HP DMA without having to deal with complicated local installations or setups.

Expert Engine

The Expert Engine controls all interactions between all product components.

Repository

The Repository stores all collected information, user settings, system configuration information, and metadata. The Repository uses Oracle as its underlying database platform.

Agent

The Agent runs on each managed server. It is responsible for collecting performance data and running automation routines.

2 Getting Started

This chapter includes the following topics:

- About This Guide on page 15
- System Requirements on page 16
- Introduction to the HP DMA Web Server on page 16
- Using the Mobile DMA Interface on page 18

About This Guide

This guide describes how to set up and use HP DMA Solution Packs within the HP DMA Platform to manage all of your database needs. It assumes that you have a basic understanding of how to use the Windows operating system.

Conventions

This guide uses several typographical conventions to help explain how to use HP DMA.

Table 1 Conventions

Convention	Definition
Bold	Words in bold show items to select or click, such as menu items or buttons.
`Courier New`	Files, paths, and commands in Courier New style show items that are file names, path names, or commands.
🔒	This symbol means that the following information is read only.
✓	This symbol appears within a green bar across the top of the screen and means that you have successfully completed a task.
🔴	This symbol appears within a red bar across the top of the screen and means that you have not successfully completed a task.
⊖	Click this red remove icon to delete information adjacent to this sign.

Table 1 Conventions (cont'd)

Convention	Definition
⊕	Click this green add icon to add information adjacent to this sign.
✖	Click this red "x" icon to perform a component delete and confirm your action.
?	Click this question mark icon to view a helpful description about the current screen you are viewing.
Filter	Where you see a filter box, you can perform a real-time filter. Type what you are searching for in the Filter field and see the filter results display as you type.

System Requirements

See the *HP DMA Installation Guide* Guide for information on system requirements.

Introduction to the HP DMA Web Server

Figure 1 shows the opening screen of the HP DMA Web Server.

Figure 1 HP DMA Dashboard

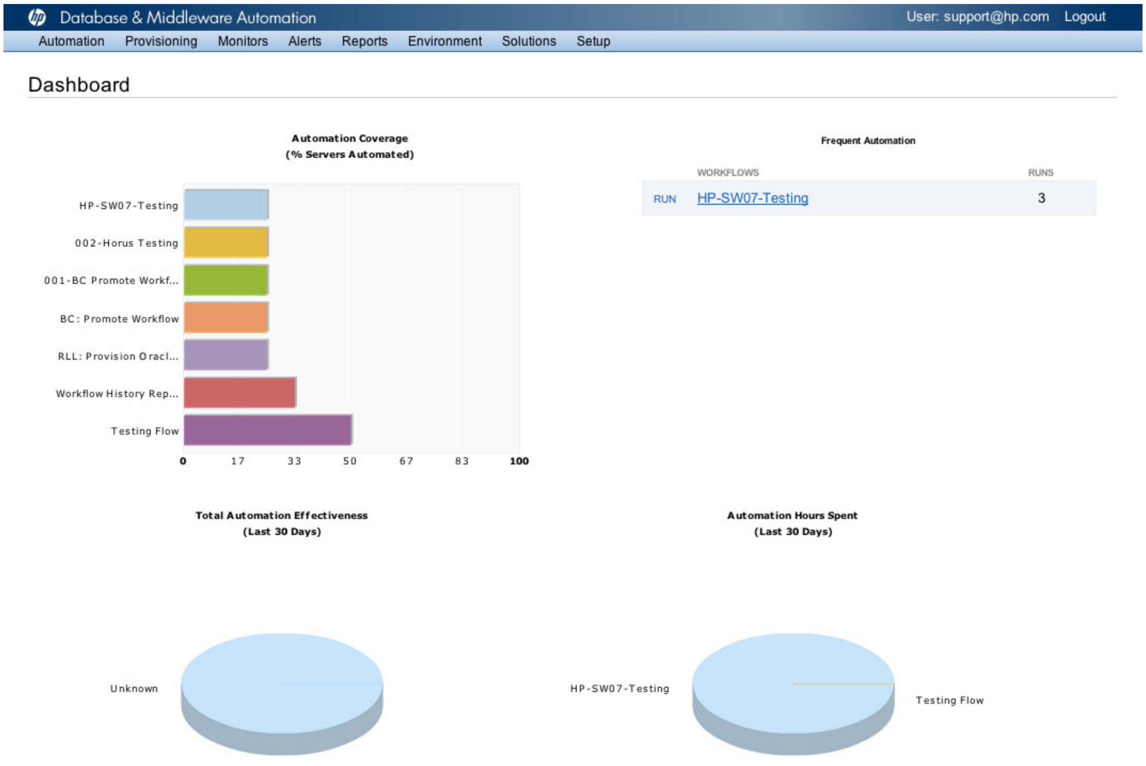

Logging On to the Web Server

1. Open a web browser.
2. Type in the URL for accessing the Web Server.

   For example: `http://server_name:8080`
3. Log in to the Web Server using your user name and password.

   The HP DMA Dashboard screen opens.

The Web Server is organized into the following functional areas:

- Automation: See Chapter 3, Automation, on page 21 of this guide.
- Provisioning: See Chapter 4, Provisioning, on page 57 of this guide.
- Monitors: See Chapter 5, Monitors, on page 63 of this guide.
- Alerts: See Chapter 6, Alerts, on page 67 of this guide.
- Reports: See Chapter 7, HP DMA Reports, on page 75 of this guide.
- Environment: See Chapter 8, Environment, on page 145 of this guide.
- Solutions: See Chapter 9, Solutions, on page 159 of this guide.
- Setup: See Chapter 10, HP DMA Administration (Setup), on page 165 of this guide.

Using the Mobile DMA Interface

The Mobile DMA Interface is a web application that provides access the HP DMA Web Server from your cellular phone.

Logging On to the Mobile DMA Interface

1. Navigate to the link provided to you by your system administrator.
2. Enter your user name and password.
3. Click **Login**.

   The Menu displays with the following links:

   - Agents
   - Statistics
   - Tickets
   - Workflows
   - Workflow History

Viewing Agents

1. Click the **Agent** link.

   The Agent screen opens. From here, you can view all your Agents and statistics about each Agent.

2. Click **Back** to return to the previous screen or click **Menu** to return to the main menu.

Monitoring Statistics

1. Click the **Statistics** link.

   The Statistics screen displays.

2. Select the statistics you want to monitor.
3. Click **Select Server**.
4. In the Server list, select the server you want to monitor.
5. Click **Get Stats**.

   The Statistics screen displays showing the selected statistics you chose to monitor.

6. Click **Back** to return to the previous screen or click **Menu** to return to the main menu.

Viewing and Assigning Tickets

1. Click the **Tickets** link.

   The Tickets screen displays.

2. Perform one of the following actions:

   - Type a Ticket ID in the Ticket ID field and click **Find**.
   - Click on a ticket in the list to open.

3. Perform one of the following actions:

   - View ticket details.
   - Click **Assign** to assign the ticket to a user. Assigning the ticket sets the ticket's status to "In Progress," which stops the escalation path.
   - View the Assignees area for user changes.

4. Click **Back** to return to the previous screen or click **Menu** to return to the main menu.

Running Workflows

1. Click the **Workflows** link.

   The Run Workflow screen displays.

2. Click the Workflow you want to run or navigate to a Workflow using the scroll bar.

   The Workflow screen displays showing the selected Workflow you want to open.

3. Click button associated with the selected Workflow.

   The Workflow starts running.

4. Click **Back** to return to the previous screen or click **Menu** to return to the main menu.

Viewing Workflow History

1. Click the **Workflow History** link.

   The Server screen displays.

2. Select the server for which you want to view Workflow History.

   The Workflow screen displays showing the selected Workflow you want to open.

3. Click **Get History**.

   The Workflow History screen displays.

4. Click **Back** to return to the previous screen or click **Menu** to return to the main menu.

3 Automation

This chapter includes the following topics:

Workflows

Introducing Workflow Components

Documentation

Documentation defines all of the information required to understand not only how a procedure is executed, but also how that procedure has been qualified and tested. Documentation encapsulates a company's best practices into a shareable document that can be exported for IT auditors, change control boards, and for training manuals for new data center administrators.

Workflow

The Workflow automates the process followed for an operational procedure. Steps 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.

Step

Steps contains the actual code used to perform a unit of work detailed in a Workflow. The step can be an executable script, among other things, such as alerts, integrations with third-party products, file watchers, and timers.

Understanding Workflow Execution Architecture

The Expert Engine controls a Workflow's flow or progression through steps, not the Agent. This process limits the amount of memory that an Agent requires. Agents periodically poll the Expert Engine searching for work.

This procedure explains conceptually how HP DMA runs a Workflow:

1. The Expert Engine finds the first Workflow Step to execute.
2. All metadata, parameters, and header variables are replaced for the Workflow step on the Expert Engine.
3. The Agent connects to the Expert Engine, periodically checks for work, and finds the script to execute.
4. The Agent executes the script and returns the output and errors while executing.
5. When the script finishes executing, the Agent sends the return code back to the Expert Engine.
6. Based on the value of the return code, the Expert Engine decides which Workflow step to execute next.
7. See step 2 on page 22 and repeat.

Workflows are complete automation packages that automate tasks, like database backups or perform self-healing actions when error conditions are detected. A Workflow is made up of a series of reusable steps. Workflows can be scheduled to run periodically, or can be triggered by the rules engine. For more information on the rules engine, see Rules on page 36.

Searching for a Workflow

You can perform a real-time filter on any Workflow by name, type, or tags. Type what you are searching for in the Workflows field and see the filter results display as you type.

Viewing/Opening a Pre-Existing Workflow

From the Automation > Workflows screen, you can view all existing Workflows as well as some of the Workflow's more important properties. You can also view Workflows grouped/sorted by type (database platform, OS, etc.) or other criteria.

1. In the Workflows pane, point to the Workflow Name. As you point to the Workflow, you can view the associated steps in the Steps pane.
2. Click the Workflow you want to view.

   The Documentation tab displays.

   You can view also a Workflow's graphical layout. Click the Workflow tab.

   The Workflow tab displays.

Creating a New Workflow

1. On the Automation > Workflows screen, click **New Workflow**.

   The Documentation tab opens.

2. Type the following:

   - Name: Workflow's name.
   - Tags: Use this field as a keyword field, to type descriptive words about a Workflow's function, language, compliance, etc., so that you can easily find or filter for this Workflow.
   - Type: Helps you decide where or at what level to run a Workflow.
   - Version: Indicates how many times a Workflow has been saved.
   - Target Level: This determines what Targets you can select when you associate this Workflow to a Deployment. Example: Server Level means that only server targets are available. Instance Level means that only Instance Level targets are available.
   - Documentation

3. You must add at least one step to your Workflow in order to save the Workflow.

   - You can create a new step in two ways:
     - — See Creating a New Step on page 29.
     - — On the Workflow tab, click the **New Step** link to use the Step Wizard to create a step on the fly. The Step Wizard adds the newly-created step to the step list, lists the step on the Workflow panel, and displays the step in the Workflow workspace.
   - You can add a pre-existing step. See step 4 on page 23.

4. Search for the step type you want to add.

   - As you click on a step, the step's description displays below the step list.

> ➤ A Download Software step is available only from the Automation > Workflows > Workflow tab. This step allows you to download a file to a specified location. To locate this step, on the Workflow tab, type "download" in the filter box. See Software on page 61 for additional information.

5. Double-click a step to move it to the desired location in the workspace.

   As you drop the step onto the workspace, the step's details display directly below the workspace. You can view the following details about the chosen step:

   - Name: Step's name. Once you add a step, you can click on a Step's name to open Automation > Steps and view detailed information about the step.
   - Next: Click this field to edit the step that comes before or after the current step. Click this field to edit the result.
   - ⛔: Click to delete a step from a Workflow.

6. Click **Save**.

   "Workflow Saved Successfully" displays in a green bar at the top of the Workflows screen. Click the **Workflow** tab.

➤ Workflows must have unique names. If you attempt to save a Workflow by a previously-used name, “Workflow name already used” displays in a red bar at the top of the Workflows screen. Choose a different name for your Workflow.

Copying a Workflow

Copy is available from all the tabs in the Automation > Workflow area. Creating a copy of a Workflow saves time by allowing you to reuse information in a Workflow by renaming it without having to re-type the Workflow’s information.

1. Click **Copy**.

   The Documentation tab displays and the Workflow name changes to “Copy of <Workflow name>.”

2. Make any changes to the copy.
3. Click **Save**.

Importing a Workflow

1. On the Automation > Workflows tab, click **Import Workflow**.

   The Import Workflow screen opens.

2. Click **Browse** to find the Workflow you want to import.
3. Click **Import**.

Exporting a Workflow

1. Navigate to Automation > Workflows tab.
2. Click the Workflow you want to export.
3. Click **Export**.
4. A screen displays and allows you to select a location to which you want to save the Workflow.
5. Click **OK**.
6. Click **Save**.

Adding an Attachment to a Workflow

You many only add attachments to unlocked Workflows.

1. Navigate to Automation > Workflows.
2. Open the Workflow you want to view.
3. Click **Attach a File**.
4. Click **Browse** to search for an attachment.
5. Click **Open** to add the attachment.
6. Click **Save**.

Viewing Attachments

1. Navigate to Automation > Workflows.
2. Open the Workflow you want to view.
3. Click the filename of the attachment you want to view.
4. A screen displays and asks you to select an action of Open or Save.
5. Click **OK**.

   Attachment opens for viewing.

Removing Attachments from a Workflow

1. Navigate to Automation > Workflows.
2. Open the Workflow you want to view.
3. Click ⊖ to remove the attachment.

Assigning Roles to a Workflow

1. In the Workflows pane, point to the Workflow Name.
2. Click the Workflow you want to view.

   The Documentation tab displays.
3. Click the **Roles** tab.

   The Roles tab displays.
4. Select or clear the **Read** or **Write** check boxes, depending on the permission you want to grant.
5. Click **Save**.

Deleting a Workflow

You can delete a Workflow unless its status is “Read Only.”

‼ If you delete a Workflow that has associated deployments, the associated deployments will be deleted automatically with the Workflow.

1. Navigate to Automation > Workflows.
2. Open the Workflow you want to delete.
3. Click and confirm delete.

Steps

Steps are reusable automation components. They are assembled into Workflows that automate a task or system healing action. Steps can accept input parameters for customization and provide output for subsequent steps to use.

Searching for Steps

You can perform a real-time filter on any step by name, type, or tags. Type what you are searching for in the Steps field and see the filter results display as you type.

Viewing a Step

➤ Steps are Read Only. You must copy the step before it can be used. See Copying a Step on page 30.

1. In the Steps pane, point to the Step Name.

   As you point to a step, you can view the Workflow that uses that particular step.

2. Click the step you want to view.

   The General tab displays.

General tab

The General tab displays information about the step you want to view. The step's name displays above the tabs. The fields available on the Action tab change depending on which Category is selected on the General tab.

In the Properties area, you can view and edit the following step information:

- Name: Step's name.
- Tags: Use this field as a keyword field to type descriptive words about a step's function, language, compliance, etc., so that you can easily find or filter for this step.
- Type: Helps you decide where or at what level to run a step.
- Category: Specifies the type of step. You cannot edit the Category field once this a step has been created and saved.
  - — Script Type Step: Executes the code on the Action tab.
  - — Alert Type Step: Sends an alert to the specified email address.
  - — Email Type Step: Sends an email to the specified email address.
  - — Prompt Type Step : Allows you to further customize a Workflow by accepting user input during a running Workflow.
  - — Variable Timer: Type the number of minutes you want to wait before the Workflow proceeds to the next step.
  - — File Watcher Type Step: Goes to the path that you specify and locates the file. Once it finds the file, the Workflow moves to the next step.

    - Fixed Timer Type Step: Type a specific time in 24-hour-time, and at the specified time, the Workflow moves to the next step.
- Version: Tells how many times a step has been saved.

In the Documentation area, you can edit, type, or view documentation that is related to the step that you are viewing.

Action tab

The Action tab displays what action the specified step takes when that step type is used in a Workflow. The fields available on the Action tab reflect the Step Category type that you designate on the General tab.

- Action tab: Inserting Functions

  Inserting a function appends data in the code field.

  a Click the **Insert Function** link to open the Insert Function box.

  b Select the function you want to insert.

  c Click **Insert function** or click **X** to cancel.

- Action tab: Importing Scripts

  Importing a script overwrites existing data in the code field.

  a Click the **Import Script** link to open the Insert Function box.

  b Click **Browse** to locate the script you want to import.

  c Click **Open** to import the script or **X** to cancel.

- Script Type Category
  - Call wrapper: Stores the interpreter location. Interpreter executes the script.

    For example: **/bin/ksh, /usr/bin/perl, cscript /E VBS)**

    Additionally, you can run the following built-in call wrappers:

    - jython: Runs the script using the built-in python interpreter shipped with every Agent.
    - SQL: Runs the code as SQL against the database (SQL only, no DML/DDL).
    - osql: Runs the code as TSQL against the database (MSSQL only).
  - Code: Code to run with the Call wrapper.
- Email Type Category
  - To: email address to whom you want to receive the message.
  - Subject: Subject of email.
  - Message: Contents of email.
- Alert Type Category
  - Severity level: Type a severity level, which tells a step when to create an alert.
  - Subject: Subject of email.
  - Message: Contents of email.

- Prompt Type Category
  - Prompts area: Displays new prompt questions as you create them. Click **Add Prompt** to customize your prompt. Type the prompt information in the Edit Prompt area.
  - Edit Prompt: Allows you to customize your prompt.
    - Question: Type a question to be presented to the user when this step is run.
    - Name: Type the prompt's name.
    - Type: Specify the type of prompt that displays.

      Text

      Password

      Dynamic list

      Static list. Click multi-select if you want the user to have multiple options from which to choose when this step runs.

      Note: Selecting Multi-select and specifying Options are available only if you choose Type = static list.
    - Options: Specify the answers to your questions.
- File Watcher Category
  - File: Type a path to a specific file. Once the designated file is located, the Workflow proceeds to the next step.
- Variable Timer Category
  - Delay (minutes): Type the number of minutes you want to wait before the next Workflow step is executed.
- Fixed Timer Category
  - Hour: Type a specific time in 24-hour-time, and at the specified time, the Workflow proceeds to the next step.
  - Minute: See above.

Parameters tab

Script steps are the only step types with both input and output parameters; all of the other step categories have only input parameters. You set the input to a value, and must set the output parameter within the code. For more information, see Understanding Parameters on page 30.

- Input Parameters

  The parameters tab defines the variables that a Workflow sets when running a step so that the step can run against different objects and still be reusable. For example, in a database backup, the directory where the backup should be placed would be a good candidate for an input parameter so that both development and production database backups could use the step without modification.
- Output Parameters

  Script type steps allow you to define input as well as output parameters. Steps use output parameters to provide information to be used in further steps. For example, if a step determines the location of the Oracle home directory on a machine, it can add that location to its set of output parameters for subsequent steps to use.

- Remove Link: Removing a Parameter from a Step

  You can click the **Remove** link, located next to a parameter, if you want to remove a parameter from a step. However, if that parameter is associated to a Workflow, you cannot remove it, and the Remove link does not display.

Creating a New Step

1. Navigate to Automation > Steps.
2. Click **New step**.

   The General tab displays information about the step you want to add. In the Properties area, you can add the following information:

   - Name: Step's name.
   - Tags: Use this field as a keyword field, to type descriptive words about a step's function, language, compliance, etc. so that you can easily find or filter for this step.
   - Type: Helps you decide where or at what level to run a step.
   - Category: Specifies the type of step.
   - Version: Tells how many times a step has been saved.
3. Add a step category.
   - Script Type Step: Executes the code on the Action tab.
   - Alert Type Step: Sends an alert to the specified email address.
   - Prompt Type Step: Allows you to further customize a Workflow by accepting user input during a running Workflow.
   - Email Type Step: Sends an email to the specified email address.
   - Variable Timer: Type the number of minutes you want the Workflow to wait before the next Workflow step is executed.
   - File Watcher Type Step: Goes to the path that you specify and locates the file. Once the step finds the file, it moves to the next step in the Workflow.
   - Fixed Timer Type Step: Type a specific time in 24-hour-time, and at the specified time, it moves to the next step in the Workflow.
4. You can see a Step Version: Version prepopulates this field.

   In the Documentation area, you can type or view documentation that is related to the step that you are creating.
5. Click **Save**.

   To type additional information about this step, see the following sections:

   - Action tab on page 27.
   - Parameters tab on page 28.
   - History tab on page 35.
   - Workflows tab on page 35.

Copying a Step

Since steps are read-only, you must copy a step before you can modify it.

1 Navigate to Automation > Steps.

2 Click a step you to modify.

3 Click **Copy**.

Understanding Parameters

You create parameters at the Step level and assign values at the Workflow or Deployment level. Use the following sections as examples to illustrate how parameters are created and then used at various levels.

- Creating Parameters on page 30
- Assigning Parameters on page 31
- Using Parameters on page 33

Creating Parameters

Figure 2 shows the screen for adding parameters to a step.

Figure 2 Automation > Steps > Parameters tab

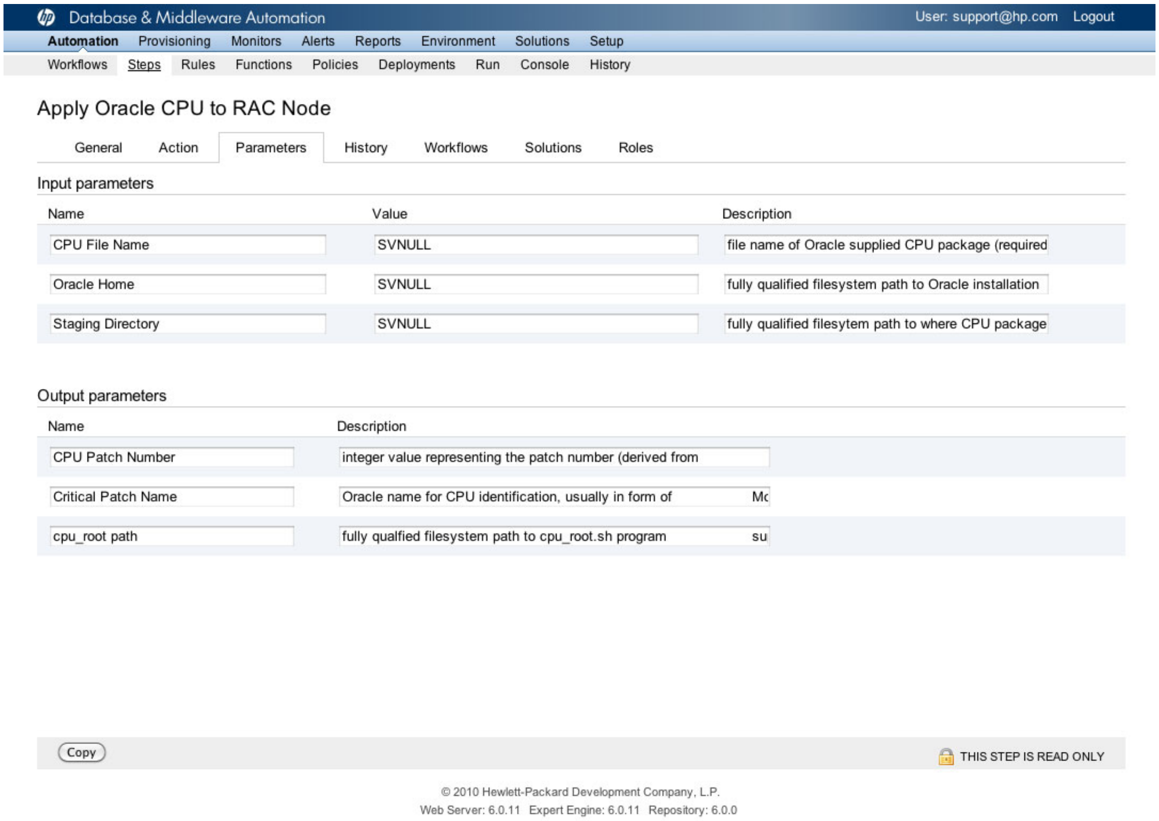

Apply Oracle CPU to RAC Node

Input parameters

Name	Value	Description
CPU File Name	SVNULL	file name of Oracle supplied CPU package (required
Oracle Home	SVNULL	fully qualified filesystem path to Oracle installation
Staging Directory	SVNULL	fully qualified filesytem path to where CPU package

Output parameters

Name	Description
CPU Patch Number	integer value representing the patch number (derived from
Critical Patch Name	Oracle name for CPU identification, usually in form of Mc
cpu_root path	fully qualfied filesystem path to cpu_root.sh program su

To create a new parameter, follow these steps:

1. Create a step by navigating to Automation > Steps.
2. Click **New step**.
3. Create a new step.
4. Add input parameters, and if applicable for a script type step, output parameters.
5. Click **Save**.

   To see your parameters in use, see Assigning Parameters on page 31.

Assigning Parameters

Parameters set at the Workflow level are typically constant values or assigned at execution time using custom fields. For example, these values may change based on your Organization, so the parameter remains consistent but the value assigned to that parameter changes. However, when working with parameters at the deployment level or with a Run action, the parameters are not as static as those parameters which you would set at the Workflow level. If you know that you have a value that changes often, do not assign it at the Workflow level--assign it at the Deployment level. For more information, see Using Parameters on page 33.

Assign the parameters you defined at the Step level in Automation > Workflows > Workflow tab.

Once you navigate to the Workflow tab, view the Workflow table, located below the Workflow workspace. If the step contains an Input parameter, there is an arrow next to each step that when clicked, displays the “values” that can be associated to the step's parameters.

Figure 3 Automation > Workflow tab

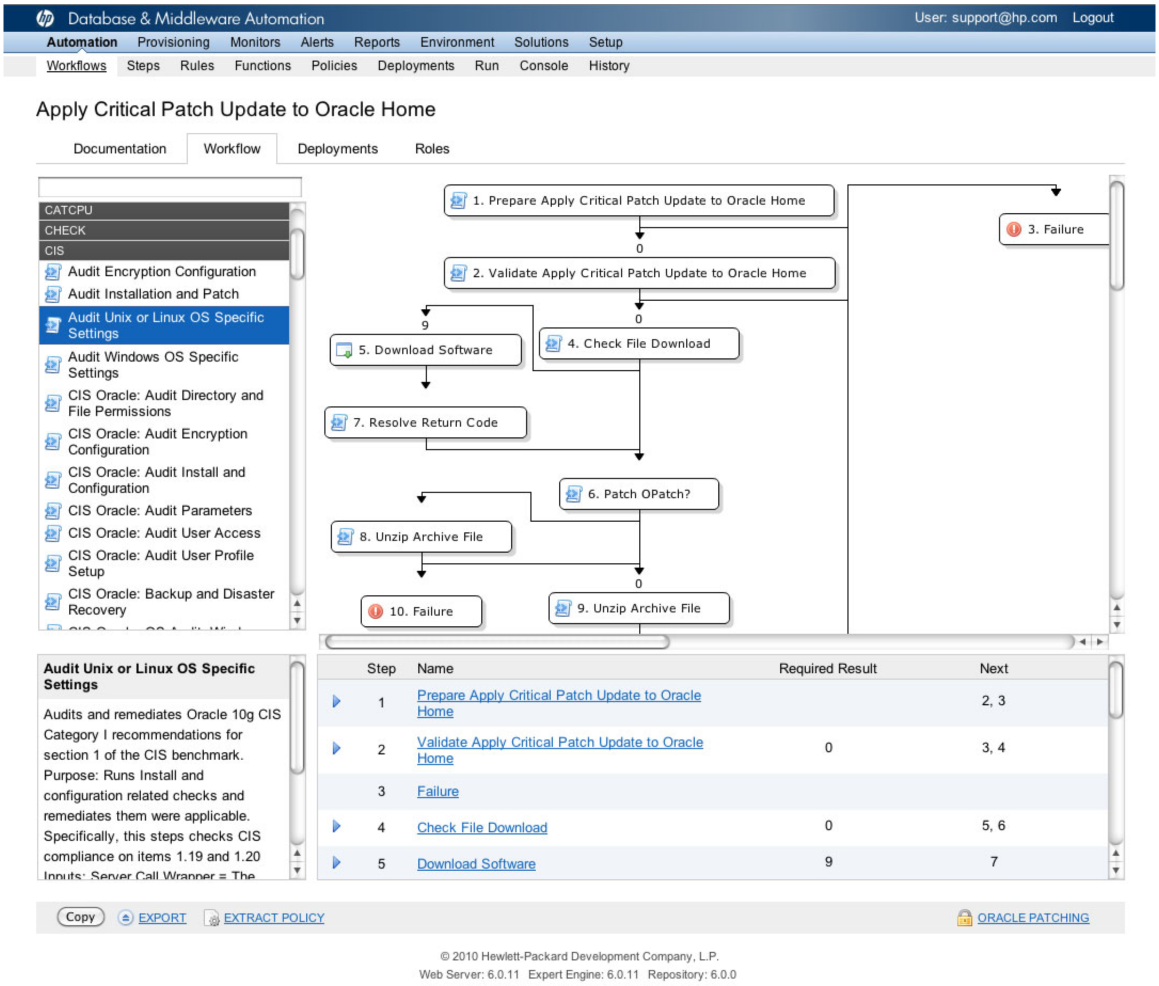
Database & Middleware Automation
User: support@hp.com Logout
Automation Provisioning Monitors Alerts Reports Environment Solutions Setup
Workflows Steps Rules Functions Policies Deployments Run Console History
Documentation Workflow Deployments Roles
CATCPU
CHECK
CIS
Audit Encryption Configuration
Audit Installation and Patch
Audit Unix or Linux OS Specific Settings
Audit Windows OS Specific Settings
CIS Oracle: Audit Directory and File Permissions
CIS Oracle: Audit Encryption Configuration
CIS Oracle: Audit Install and Configuration
CIS Oracle: Audit Parameters
CIS Oracle: Audit User Access
CIS Oracle: Audit User Profile Setup
CIS Oracle: Backup and Disaster Recovery
1. Prepare Apply Critical Patch Update to Oracle Home
3. Failure
0
2. Validate Apply Critical Patch Update to Oracle Home
9
0
5. Download Software
4. Check File Download
7. Resolve Return Code
6. Patch OPatch?
8. Unzip Archive File
0
10. Failure
9. Unzip Archive File
Copy
EXPORT
EXTRACT POLICY
ORACLE PATCHING
© 2010 Hewlett-Packard Development Company, L.P.
Web Server: 6.0.11 Expert Engine: 6.0.11 Repository: 6.0.0

Apply Critical Patch Update to Oracle Home

Audit Unix or Linux OS Specific Settings

Audits and remediates Oracle 10g CIS Category I recommendations for section 1 of the CIS benchmark. Purpose: Runs Install and configuration related checks and remediates them were applicable. Specifically, this steps checks CIS compliance on items 1.19 and 1.20 Inputs: Server Call Wrapper = The

Step	Name	Required Result	Next
1	Prepare Apply Critical Patch Update to Oracle Home		2, 3
2	Validate Apply Critical Patch Update to Oracle Home	0	3, 4
3	Failure		
4	Check File Download	0	5, 6
5	Download Software	9	7

The “Values” drop-down list contains built-in metadata and any user-defined custom fields, if any exists. Output parameters will only display in the “Values” drop-down list if a prior step in the Workflow contains an output parameter. If it is applicable to assign output parameters, then they will display in the list as well. In order to view an output parameter in the “Values” drop-down list, the step must be used in a Workflow.

Figure 4 Automation > Workflow tab: “Values” drop-down list

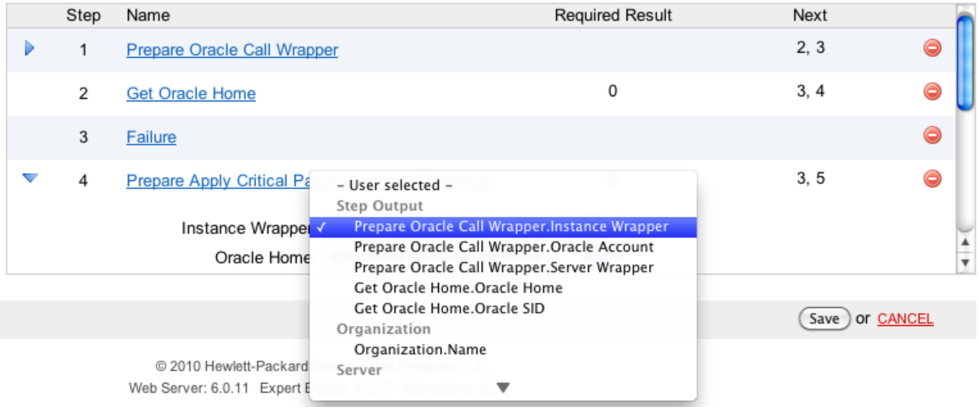
– User selected –
Step Output
✓ Prepare Oracle Call Wrapper.Instance Wrapper
Prepare Oracle Call Wrapper.Oracle Account
Prepare Oracle Call Wrapper.Server Wrapper
Get Oracle Home.Oracle Home
Get Oracle Home.Oracle SID
Organization
Organization.Name
Server
Save
or
CANCEL
© 2010 Hewlett-Packard
Web Server: 6.0.11 Expert E

Step	Name	Required Result	Next
1	Prepare Oracle Call Wrapper		2, 3
2	Get Oracle Home	0	3, 4
3	Failure		
4	Prepare Apply Critical Pa		3, 5

Instance Wrappe

Oracle Home

If you assign an output parameter to an input parameter at the Workflow level, the parameters do not display in the Deployment parameter list. All parameters that are not assigned to a custom field, or mapped at the Workflow level will be modifiable at the deployment level. Any parameter that is not set in the Workflow or deployment level will use the default value assigned for that step.

To assign parameters:

1 Navigate to Automation > Workflows.

2 Perform one of the following tasks:

- Click **New workflow** to create a Workflow.
- Select an existing Workflow.

3 Navigate to the Workflow tab. Design or edit a Workflow using the Workflow workspace.

The steps in the Workflow display in the Workflow table, below the Workflow workspace. The arrows to the left of each step expand the step to display any parameters you created, as shown in Figure 4.

- A numeric value in the Required Result column is the return code that must be received from at least one parent node in order for that step to run.
- You can use the Next field to reorder your Workflow's steps.

4 Click the arrow next to each step. The input or output parameter "values" display in a drop-down list. Assign a value from the "Values" list to the desired parameter.

➤ An arrow will not display next to a step that does not have input parameters.

5 Click **Save**.

The Workflow screen opens, and the message "Workflow saved successfully. Would you like to deploy the workflow now?" displays. See Chapter 3, Deployments, on page 50 of this guide.

Using Parameters

There are several levels at which you can set parameters: Workflow, Deployment, and Run. Parameters should be assigned at the Deployment level when the value is specific to the targets that are part of the deployment. For example, you may wish to use the same Workflow with production and development servers, but need to use a different parameter value. This can be accomplished by creating one deployment of that Workflow for production server, and a second deployment for development servers. This allows you to set the same parameter differently for the different set of targets.

➤ Once you assign a parameter at the Workflow level, it does not display at the Deployment level so it cannot be overridden.

1 Click link **...deploy workflow now**. If a deployment already exists, the prompt says "Would you like to run the workflow now?"

2 Create a new Deployment. On the Attributes tab, type:

- Name: Type a deployment name.
- Workflow: Workflow name is pre-populated. Click **View Workflow** if you need to see the Workflow for which you are creating a deployment.

- Schedule (optional): Select a deployment schedule from the drop-down list.

On the Deployments > Parameters tab, if you select “enter at runtime” and try to save a deployment with a schedule, you cannot save until you deselect the check box or unschedule the deployment.

3. Add available targets.
4. Navigate to the Parameters tab.

   You can assign built-in custom fields, user-defined custom fields, and policy attributes in this list. You will not see any output parameters; mapping output to input parameters is only possible in the Workflow editor. If you create a deployment and then add parameters to your step, newly-created parameters display in the deployment parameter list, and contain the default value assigned in the step editor.

   Deployment parameters can contain three different types of values.

   - Static text.
   - Custom fields (built-in or user created)
   - Policy Attributes, for policies that have all the targets for this deployment.

   As you type in the value you want, HP DMA uses pattern matching to try and find custom fields and policy attributes that match the value that you entered into the box. You can select the desired value from the drop-down, or type the desired value completely.

In order to obtain a complete list of all custom fields and policy attributes that are usable with this deployment, type “.” in the text-box.

Click **Restore Defaults** to replace all parameter values with their default values assigned at the Step level.

5. Click **Save**.

   The Deployment screen opens, and the message “Deployment saved successfully. Would you like to run the workflow now?” displays. For more information, see Running a Workflow on page 54.

Creating Prompt Email Messages Using Runtime Parameters

If a Workflow contains a prompt type step, you can configure the prompt step to email specified people, who then need to respond to the prompt’s question.

Using Metadata and Policies from a Workflow Step

You can use metadata from any Workflow Step type by using the **${Object.Attribute}** syntax.

For example, the ${DBServer.Password} metadata variable would be replaced at runtime with the actual password for the instance this Workflow Step was executed on.

For example, if the password for a given instance were `password`, the script

var password = "${DBServer.Password}"

would be replaced with

```
var password = "password"
```

at runtime.

You can also replace user-defined metadata using these conventions.

Using Parameters from a Workflow Step

You can call a parameter from a Workflow Step using the **${ParameterName}** convention. Use a parameter to change the way this Workflow Step runs prior to runtime. For example, if you have a Parameter called `InstallDirectory` with a default value of `/opt/app/oracle`, the script

var installDir = “${InstallDirectory}”

would be replaced with

```
var installDir = "/opt/app/oracle"
```

at runtime.

Using Headers from a Workflow Step

You can use input from other Workflow Steps using the **${Header.Variable}** convention. Use an input header when passing information from one step to another, such as the **ORACLE_HOME** or other variables. For example, the following Workflow contains two Workflow Steps. The first step finds the **ORACLE_HOME** variable, and the second step performs an Oracle Refresh. The Oracle Refresh Workflow Step requires the **ORACLE_HOME** variable in order to execute.

History tab

The History tab allows you to view historical step information. You can click on the desired step version and then view the Details for each version listed. You can also make a copy of the step.

- Version: The step’s version.
- Date: Version’s creation date.
- Person: Person who created the step’s version.
- Message: Details about the step’s version.

To view historical step information, click on the desired step’s version. The Details area populates with the step’s Call wrapper and Code.

Workflows tab

The Workflows tab tells you which Workflows are using your step. It also allows you to access the Workflows area in HP DMA. Click a Workflow to go to the Automation > Workflows area, where you can view and edit the Workflow.

Solutions tab

The Solutions tab tells you in which Solution Packs the current step is being used. For more information on Solutions, see Chapter 9, Solutions, on page 159 of this guide.

Roles tab

Permissions settings for baseline steps (ones that ship in HP DMA) cannot be changed, even by an administrator. To change permissions for a step that was created outside of the baseline steps, select or clear the check boxes to grant or revoke write access. None of the roles have write access to any steps if there is a "----" in the Write column. For more information on changing permissions as an administrator, see Roles on page 167.

Rules

Using the Rule Controller to Run Workflows

The rules engine is a component inside of the Expert Engine that has access to the stream of data coming into the Repository. When you write a rule and activate it for a target server, the rules engine starts looking for data collected from the Agents that is relevant to your rule on that server. Rules are not associated with servers per se - they have deployments who have targets.

When the rule evaluates to true, based on data from one of the targets, it runs a Workflow, not the entire deployment, on the affected target server. Rules cannot assign input parameters to the Workflows they run. You still have to use the ${Header.MyVariable} syntax in the Workflow to access rule data.

Each time new data comes into the Expert Engine from the target server, your rule is evaluated or executed. The rule checks the state of the data and determines whether or not to execute a Workflow. The final result of your rule must be a boolean (true/false) value. When a rule evaluates to true, it runs a Workflow. The Workflow can be something as simple as firing an alert or it can be a self-healing action that runs when a complex error condition is discovered.

Defining a Rule

The rule engine uses JavaScript as the scripting language syntax. The code in the examples that follow is JavaScript code. Documenting the entire JavaScript language is outside the scope of this guide. However, excellent documentation can be found at the Mozilla Development Center (**http://developer.mozilla.org/**).

Functions

The rule engine provides several functions that give you access to all the collected data in the repository. Using these functions you can write rules to detect errors in your environment and run Workflows to alert on or fix the problem. Table 2 lists the available functions. The following sections document each function.

Table 2 Functions

Function	Description	Example
org (propertyName)	The propertyName value you pass in to the function can be a built-in property or a user defined field you created.	var orgName = org('Name')
os (propertyName)	The propertyName value you pass in to the function can be a built-in property or a user defined field you created.	var osName = os('Name')
dbs (propertyName)	The propertyName value you pass in to the function can be a built-in property or a user defined field you created.	var instanceName = dbs('Name')
db (propertyName)	The propertyName value you pass in to the function can be a built-in property or a user defined field you created.	var dbCustom = db('My Custom Field')
stats (type, name)	The stats function returns an array of Stat objects (stats.length). You pass in the type of statistic and the statistic name you want to retrieve.	var found = stats('OS', 'User CPU Usage')
red (type, name)	Retrieves the red threshold for the statistic. Thresholds are configured on the Alerts > Thresholds screen.	var red = red('OS', 'Disk Busy')
yellow (type, name)	Retrieves the yellow threshold for the statistic. Thresholds are configured on the Alerts > Thresholds screen.	var yellow = yellow('OS', 'Disk Busy')
repository (dbType, tableName, params)	Returns an array of Row objects containing all of the rows from the most recent collection snapshot (array.length). Each Row object has columns in it containing the data you need to use in the rule. The rows do not return in any guaranteed order.	var sessions = repository('Oracle', 'Sessions')

Table 2 Functions (cont'd)

Function	Description	Example
header (name, value)	The header() function allows you to pass data in the automation workflow that runs when the rule evaluates to true.	header('Message', 'CPU is 100% used.')
filter ( )	Runs a function on every item in the array and returns an array of all items for which the function returns true.	var alerts = cpuStats.filter(function(stat) { return stat.value > 90 })
map ( )	Runs a function on every item in the array and returns the results in an array.	var msgs = diskStats.map(function(stat) { return stat.resource + ' is ' + stat.value + ' % used.' })

Environment Settings

There are four functions that return metadata about the servers and databases in your monitored environment:

- `org(propertyName)`
- `os(propertyName)`
- `dbs(propertyName)`
- `db(propertyName)`

The propertyName value you pass in to the function can be a built-in property or a user defined field you created.

Example

Let's say you've defined some objects in your environment like this:

```
Production (Organization)
–- pluto (server)
——P01 (instance)
————-P01_DB (database)
```

The following rule snippet:

```
var orgName = org('Name')
var osName = os('Name')
var instanceName = dbs('Name')
var dbCustom = db('My Custom Field')
```

When the rule runs, the orgName variable will be set to "Production," osName will be "pluto," etc. The dbCustom variable will be set to the value of the My Custom Field for the P01_DB. This last example uses a field that you've defined on the database object rather than a built-in property that ships with HP DMA. Using these custom fields you can tune rules and automation for your environment.

All built-in properties as well as your custom fields are listed in the Rule editor user interface. Select the object you want inspect (i.e. Organization) and the properties display in the list box to the right. Double-clicking a property inserts it into your rule.

Retrieving Collected Data

For example, you are looking for an operating system statistic called User CPU Usage and alerting when it is more than 80 percent.

The value property contains the most recently collected value of the statistic. The date property contains the timestamp of when that collection occurred.

The `value2/date2` and `value3/date3` pairs work the same way but represent increasingly older collected values. So, value is the most recent and value3 is the oldest.

Having access to the three most recent values allows you to write rules that evaluate the trend of a statistic rather than just one data point. In the CPU example, you probably do not want to alert when CPU usage spikes but you do want an alert when it is consistently high.

```
var found = stats('OS', 'User CPU Usage')
var cpu = found[0]
var threshold = 80
cpu.value > threshold && cpu.value2 > threshold && cpu.value3 > threshold
```

Table 3 Stat Object Properties

Function	Description	Example
value	The value property contains the most recently collected value of the statistic.	`cpu.value > threshold`
date	The date property contains the timestamp of when that collection occurred.	var msg = "The error occurred at " + cpu.date
value2	The value2 property contains the second most recently collected value of the statistic.	`cpu.value2 > threshold`
date2	The date2 property contains the timestamp of when that collection occurred.	var msg = "The error occurred at " + cpu.date2
value3	The value3 property contains the third most recently collected value of the statistic.	`cpu.value3 > threshold`
date3	The date3 property contains the timestamp of when that collection occurred.	var msg = "The error occurred at " + cpu.date3
resource	Contains the resource name associated with this statistic. For disk stats this will be the mountpoint name (`/tmp`, `/home`). For tablespace stats it will be the tablespace name (TEMP, USER).	var msgs = busy.map(function(disk) { return disk.resource + " = " + disk.value})
rate	The rate() function built into every Stat object calculates how fast the value is changing in minutes by looking at the most recently collected data.	disk.rate() > 1024

Discovering How Fast a Statistic is Changing

Once you have been collecting data for a few minutes, you can find how fast a statistic is changing.

```
var found = stats('OS', 'Disk Usage (KB)')
var disk = found[0]
disk.rate() > 1024
```

The rate() function built into every Stat object calculates how fast the value is changing in minutes by looking at the most recently collected data. In the above rule, you are going to alert when the disk space used is growing by at least 1 MB per minute.

Statistics Tied to Resource Name

In the CPU example there was only one Stat object returned by stats() because CPUs do not have names. However, some stats are tied to resource names. For example, you have a Windows machine with 3 disk drives: C, D, and E and you want to alert whenever the drives are extremely busy.

```
var found = stats('OS', 'Disk Busy')
var busy = []
for (var i = 0; i < found.length; i++) {
  if (busy.value > 80) {
    busy[busy.length] = found
  }
}
busy.length > 0
```

This rule loops through all of the disk statistics and finds drives that are more than 80 percent busy. Since all rules must end in a boolean statement to tell the engine whether or not to fire the Workflow, check the length of the busy array to see if there are any offending drives.

This looks pretty straightforward but a bit verbose. Generating a helpful error message containing the names of drives and their busy statistic requires several more loops and dozens of lines of code. There is a more efficient method.

First, change the rule to remove that for loop:

```
var found = stats('OS', 'Disk Busy')
var busy = found.filter(function(disk) { return disk.value > 80 })
busy.length > 0
```

This introduces a powerful Array method called filter. The filter method returns an array of objects that satisfy some criteria. In this case, the criteria is any disk drive being used 80 percent of the time or higher. The method takes a single function argument. The function is called one time for each item in the found array. The current item is passed into the function’s disk argument. The function must return a boolean value. If it returns true, the current item is added to the new array. If it returns false, the current item is filtered out of the resulting array.

You reduced six lines of code down to a single statement. That statement better represents the rule’s intentions. The rule is filtering out any disks that are in normal usage ranges and alerting on the busy disks. It is not concerned about the mechanics of looping and building arrays.

Now that the rule is cleaned, add a helpful message to the resulting alert Workflow.

First, change the rule to generate the message from the list of busy disks.

```
var found = stats('OS', 'Disk Busy')
var busy = found.filter(function(disk) { return disk.value > 80 })
var msgs = busy.map(function(disk) { return disk.resource + " = " +
disk.value})
busy.length > 0
```

This introduces another Array method called map. The map method loops through the array and returns a new array of values. You are mapping the original array into a new array after performing some kind of transformation on each element.

In this case, you are creating a message String for each disk in the busy array.

The msgs variable is an Array containing a message like: C: = 82 for each busy disk.

There is a list of messages that need to be sent to the Workflow

The header() function allows you to store data in the Workflow header for use in alerts, scripts, e-mails, etc.

```
var found = stats('OS', 'Disk Busy')
var busy = found.filter(function(disk) { return disk.value > 80 })
var msgs = busy.map(function(disk) { return disk.resource + " = " +
disk.value})
header('Message', msgs.join('\n'))
busy.length > 0
```

In the above code you completed the following steps:

1 Joining every element in the msgs array into a single String separated by the newline character. This turns an array of [“C: = 82”, “D: = 91”, “E: = 99”] into a String of:

```
"C: = 82
D: = 91
E: = 99"
```

2 Storing that resulting String into the header under the variable name “Message.”

When the rule returns true and the Workflow is executed, it will have the message available to it by using ${Header.Message}.

You can improve the rule by removing the hard-coded threshold of 80 percent by using the red() and yellow() functions. These functions return the thresholds defined in the Server editor into the current rule.

```
var found = stats('OS', 'Disk Busy')
var red = red('OS', 'Disk Busy')
var busy = found.filter(function(disk) { return disk.value > red })
var msgs = busy.map(function(disk) { return disk.resource + " = " +
disk.value})
header('Message', msgs.join('\n'))
busy.length > 0
```

By asking the red() function for the threshold rather than coding an 80 into the rule, you have made this rule reusable across hundreds of servers in your environment. You do not need to create a customized rule for each server. You can simply update the thresholds on the Server editor and let the rule grab that value out of the red() or yellow() function.

Accessing Other Repository Metadata

Sometimes you need more in your rule than just statistics. What if you want to trigger a Workflow to kill Toad sessions that are holding table locks or causing performance problems?

In this rule you need access to two tables: Sessions and Locks. The Sessions table contains the name of the client application and the SID. The Locks table contains information about which SIDs are holding locks. So, you can retrieve the list of sessions, find the Toad SIDs, and use those SIDs to find any locks they are holding. You can then pass the SIDs to the Workflow to go off and kill those sessions.

```
var sessions = repository('Oracle', 'Sessions')
```

The repository() function returns an array of Row objects containing all of the rows from the most recent collection snapshot. Each Row object has columns in it containing the data you need to use in the rule.

```
var sessions = repository('Oracle', 'Sessions')
var toads = sessions.filter(function(session) { return session['OS User']
== 'toad' })
var sids = toads.map(function(session) { return session['SID'] })
```

At this point, you found the most recent sessions connected to the database, filtered out all but the Toad sessions, and created a list of SIDs for the Toad sessions. But you still do not know if any of these sessions are holding locks and causing trouble.

```
var sessions = repository('Oracle', 'Sessions')
var toads = sessions.filter(function(session) { return session['OS User']
== 'toad' })
var sids = toads.map(function(session) { return session['SID'] })
var locks = repository('Oracle', 'Locks', { 'SID': sids })
```

The final line in the rule queries the repository Locks table for SIDs in the list. The last argument to repository() is {'SID': sids}. This argument is essentially an SQL where clause. The generated SQL query might look like:

```
select * from locks where SID in (12,13,14)
```

So, you found the SIDs for Toad sessions and also found which of those Toad sessions are holding locks. All that remains is to pass the locked SIDs to the Workflow so it can terminate the processes.

```
var sessions = repository('Oracle', 'Sessions')
var toads = sessions.filter(function(session) { return session['OS User']
== 'toad' })
var sids = toads.map(function(session) { return session['SID'] })
var locks = repository('Oracle', 'Locks', { 'SID': sids })
var lockedSIDs = locks.map(function(lock) { return lock['SID'] })
header('LockedSIDs', lockedSIDs.join(','))
locks.length > 0
```

By passing a comma separated list of SIDs into the header() function, you can use the SIDs in the resulting Workflow with this syntax: ${Header.LockedSIDs}. The Workflow can do further investigation and kill off any of the SIDs that are causing performance problems.

Searching for a Rule

You can search for all rules or for deployed rules (drop-down box in the upper left corner of the screen). You can filter by rule name or Workflow type when viewing by “All,” or you can filter by server name when viewing by “Deployed.”

Viewing All Rules

1 In the Rules pane, point to the Rule Name. As you point to the rule, its associated deployments display.

2 Click the rule you want to view.

The Rule tab displays.

Creating a New Rule

1 On the Rules screen, click **New rule**.

The Rule tab opens.

2 Type the following:

- Name
- Reset Immediately: Select or clear this check box. See About Reset Immediately on page 44 for more information.

3 Configure a rule to evaluate against data coming into the server. In the lower portion of the screen, use the three builder panes to sequentially select the following:

- Objects
- Attributes
- Functions

Once you select a desired function, it displays in the coding area above the builder panes. You can view your rule’s code.

4 Click the **Workflow** tab.

5 Select a Workflow.

6 Select Workflow type. A list of Workflows displays.

7 Select the Workflow to display a set of deployments.

8 Select the **Active** check box for all deployments to be run if the rule is evaluated to true.

9 Click **Save**. The rule will pass or be invalid. To continue, you must create a valid rule. HP DMA dynamically captures and displays any coding errors as you type.

➤ To select a deployment, you must have execute privileges for at least one of the roles to which the deployment is assigned.

About Reset Immediately

When a rule is evaluated by the rule engine it returns either a true or false value.

- True means the rule will fire its assigned Workflow.
- False means that no action should be taken.

However, these true and false values have an additional behavior.

- Once a rule evaluates true, it runs its Workflow and remembers that the last evaluation was successful.
- The rule will not run its Workflow again until it returns false at least one time.
- As long as the rule continues to return true, the Workflow will not run again.

This behavior is designed to prevent self-healing Workflows from running multiple times to attempt to fix the same problem. The healing Workflow will run the first time the rule turns true but needs to be given a chance to complete its healing action. If the Workflow were to run at the same time, it might harm the healing action.

However, some Workflows need to be excluded from this behavior. An example of this is when a disk reaches a certain percent filled. There is no reason to run a Workflow (which most likely opens a ticket) every time a collection occurs - so you do not set “Reset immediately” for a rule that is triggered on disk full. Alerting Workflows, for example, should run each time the rule evaluates to true, not just the first time. The **Reset immediately** check box on the Rule editor screen accommodates these Workflows. Select the box for Workflows that need to be run each time the rule evaluates to true.

Deleting a Rule

You cannot delete read-only rules.

1. Navigate to Automation > Rules.
2. Click the rule you want to delete.
3. Click and confirm delete.
4. Click **Save**.

Editing a Rule

1. Navigate to Automation > Rules.
2. Select rule that you want to edit and make the desired edits.
3. Click **Save**.

Assigning Rules to Roles

Any user can create a new rule. You must have execute privileges for a deployment in order to view the Workflow in the drop-down list. Deployment privileges override Workflow privileges.

For more information on assigning permissions as an administrator, see Roles on page 167.

1. Navigate to Automation > Rules.
2. In the Rules pane, click the rule you want to view.

   The Rule tab displays.
3. Click the **Workflow** tab.
4. Assign a deployment.

➤ If you have execute privileges for deployments, the **Active** check box displays next to the deployment.

5. Click the **Roles** tab.
6. Assign privileges for roles by selecting or clearing the **Read** or **Write** check boxes.

➤ For administrators, role names display as links to the Role tab.

7. Click **Save**.

Enabling a Rule to Run Using Deployments

You can select one, none, or many deployments. If you do not select any deployments, the Workflow will not be run, even if the associated rule evaluates to true.

➤ Any workflow with runtime parameters cannot be scheduled or triggered by a rule; These Workflows must be run manually using Automation > Run. See Running a Workflow on page 54. For more information on Parameters, see Understanding Parameters on page 30.

1. Navigate to Automation > Rules.
2. Select rule that you want to edit.
3. Click the **Workflow** tab.
4. Select a Workflow.
5. If there are deployments associated with the selected Workflow, select the **Active** check box to activate a deployment.

➤ To select a deployment, you must have execute privileges for at least one of the roles to which the deployment is assigned. All deployments are available but only the ones to which you have execute privileges are able to be deployed.

6. Click the **Roles** tab.
7. Add desired roles.
8. Click **Save**.

Deactivating Deployments

1 Navigate to Automation > Rules.

2 Select the rule you want to edit.

3 Click the **Workflow** tab.

4 Deactivate the deployment you want to disassociate.

5 Click **Save**.

Functions

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, as well as the operating system with which they work.

Searching for a Function

You can perform a real-time filter on any function by name or by tags. Type what you are searching for in the Functions field and see the filter results display as you type.

Viewing/Opening a Function

From the Automation > Functions screen, you can view all existing functions as well as preview a function's code.

➤ Some functions are Read Only.

1 Select the function you want to view.

The General tab displays.

Creating a Step Function

1 Navigate to Automation > Functions.

2 Click **New function**.

The General tab displays.

3 Type the function's desired name.

➤ Function Names must be unique.

4 Type any tags or documentation in the appropriate areas. The Tags and Documentation fields are not required.

5 Click the **Code** tab.

6 Type new code in the Code area.

Functions must contain script code. You cannot save a function without script code.

7 Click **Save**.

Copying a Step Function

1 Navigate to Automation > Functions.

2 Select function.

3 Click **Copy**.

4 Type new name for function.

5 Click **Save**.

Modifying a Step Function

Updating a Function's code does not automatically update the code already imported into Steps from the Function. Use the Copy feature to import functions.

You cannot modify read-only functions.

1 Navigate to Automation > Functions.

2 Select the function you want to modify.

3 Edit the desired information.

4 Click **Save**.

Deleting a Step Function

1 Navigate to Workflow > Functions

2 Select function to delete.

3 Click and confirm delete.

Policies

Policies are reusable sets of attributes that can be assigned to servers, instances, and databases. Automation Workflows can reference policy attributes in their code to change the automation behavior. For example, you might define a Production Database Backup Policy that contains a Backup Location attribute. That policy can then be shared by all production databases and referenced in the Database Backup automation Workflow to store the backup files in the appropriate place.

Policies allow HP DMA to manage groups of hundreds or thousands of servers at a time without the need for configuration of each individual server. For example, you could create a Web Server policy that defines what every web server in your data center looks like. This policy might contain the following attributes:

- Software List
- Root Password
- Apache User ID
- Apache User Password

The Software List attribute would be a list of software that must be installed on every web server. This might include openssl, apache, perl, etc.

Policy Attribute Types

Policies have four different types of attributes:

- Text: This is a simple text value that users can view while deploying and running automation.
- Password: This is a simple text value. However, the value is masked when displayed so users cannot see the value.
- List: This is a free-form text field that can contain comma-delimited lists of values or other large text data not suitable for a Text type attribute.
- Software: This defines a list of software from the Software Library. Files in the library are selected as part of the policy attribute. Workflow steps can download and install this list of software during a provisioning task. See Software on page 61 for more information on downloading software during provisioning.

Deploying a Policy

Once a policy is defined with its list of attributes, it must be assigned to managed objects within the environment. Policies can be deployed at one of four levels:

- Organization
- Server
- Instance
- Database

When you choose the Server level, you see a list of servers to which to assign the policy. Now you can select the web servers in your environment.

Now that managed objects have the policy assigned, those policy attributes are available to every Workflow deployed to those objects. For example, if you have Data Backup and Security Compliance Workflows deployed to your web servers, both of those Workflows now have access to the Web Server policy attributes (root password, apache user ID, etc.)

Creating a New Policy

All policies are custom, meaning that you specify all policies for your Organization. HP DMA does not contain any default policies.

1. Navigate to Automation > Policies.
2. Click **New Policy**.
3. Type a Policy Name.
4. Select a policy Attribute. Assign a name to it, and click **Add** to add it to the list.
   - Text
   - List
   - Password
   - Software: Takes you to provisioning software list “select software dialog” select software by checking box, click **OK**.
5. Click the **Usage** tab.
6. Assign the policy to an object.

Associating a Policy with an Object

1. Navigate to Automation > Policies.
2. Select a Policy.
3. Click the **Usage** tab.
4. Assign the policy to an object.
5. In the Available area, click **Add** next to the object to which you want to assign a policy.
6. Click the **Deployments** tab to see where the policy is in use. You cannot delete a policy if it is being used.

Disassociating a Policy from an Object

1. Navigate to Automation > Policies.
2. Select a Policy.
3. Click the **Usage** tab.
4. In the Assigned area, click **Remove** to remove the policy from the object.

Extracting a Policy

Clicking the Extract Policy link at the bottom of a screen allows you to create a reusable policy that you can apply to multiple objects. Rather than manually creating the same policy multiple times, use Extract Policy to remove an existing policy from a Workflow, to customize the policy, and then to apply the policy to the desired object.

Where is a Policy in Use?

1. Navigate to Automation > Policies.
2. Select a Policy.
3. Click the **Deployments** tab to see where the policy is in use.

Deleting a Policy

1. Navigate to Automation > Policies.
2. Select a Policy.
3. Click the **Deployments** tab to see where the policy is in use. You cannot delete a policy if it is being used.
4. Select the deployment to which the policy is associated.
5. To delete a policy, remove it from the deployment. The **Delete** button on the Usage tab will be activated once you disassociate the policy from the deployment.

Assigning Policies to Roles

To change permissions for a policy, select or clear the check boxes to grant or revoke read and write access. For more information on changing permissions as an administrator, see Roles on page 167.

1. Navigate to Automation > Policies.
2. In the Policies pane, click the policy you want to view.

   The Attributes tab displays.
3. Click the **Roles** tab.
4. Select or clear the **Read** or **Write** check boxes appropriately.
5. Click **Save**.

Deployments

Deployments associate a Workflow with a target environment in which a Workflow runs. Servers, instances, and databases can be managed in groups of hundreds per deployment. Using custom fields and policies, you can customize the Workflow's behavior for groups of targets in the deployment. For example, you can create a CIS Compliance Workflow that validates whether or not servers are configured securely. Then you can create not only a Development deployment to manage your development machines but also a Production deployment to manage the production servers.

A deployment appropriately customizes a Workflow's behavior for each set of servers on which it runs. When you assign a Workflow, the selected Workflow’s available targets are determined by the selected Workflow’s target level and type.

A Workflow consists of steps. Each step in a Workflow contains a set of parameters. A deployment can set the parameters of all steps that have not been mapped at a Workflow level.

Deployments are associated with rules. See Rules on page 36 for more information.

Searching for Deployments

There are two views for the Deployments screen. You can search for Deployment by Workflow or by Targets (drop-down box in the upper- right corner of the screen). You can also filter by name.

Viewing Deployments

From the Automation > Deployments screen, you can view all existing Workflows. In the Deployments pane, you can view each Workflow's associated deployments.

1. Navigate to Automation > Deployments.
2. Select a Workflow, and then click the associated deployment you want to view by clicking **View**.

Creating a New Deployment

1. Navigate to Automation > Deployments.
2. Click **New deployment**.
3. Type the desired attributes.
    - Name: Type the deployment's name.
    - Workflow: Associate the deployment to a Workflow.

      Click **View Workflow** to take you to the Workflow that is associated with the deployment that you are viewing, if you created and associated a Workflow.
    - Schedule: Select default schedule or create a custom schedule.
    - Add or remove targets individually or as a group.
4. Click **Save**.

➤ You can run the Workflow from the current deployment by clicking on the "Would you like to run the Workflow now?" link.

Editing Deployment Attributes

1. Navigate to Automation > Deployments.
2. Select a Workflow, and then click the associated deployment you want to view by clicking **View**.
3. Edit the desired attributes.
    - Change the deployment's name.
    - Change the deployment's schedule.
    - Add or remove targets individually or as a group.
4. Click **Save**.

➤ You can run the Workflow from the current deployment by clicking on the “Would you like to run the Workflow now?” link.

Viewing Parameters Associated with a Deployment

You can view parameters in multiple areas:

- Automation > Steps: Select a step and click the Parameters tab.
- Automation > Workflow: Select a Workflow and click the Parameters tab. You can see the parameters in each step within the Workflow. “User selected” is the default. If you map to an output parameter or to built-in metadata at the Workflow level, the parameter is not available for editing at the deployment or run level.
- Automation > Deployments: Select a deployment and click the Parameters tab.
- Automation > Run: The parameters are read-only unless the runtime is selected at the Deployment level.
- Solutions > Installed Solutions > Steps > Parameters

Editing Parameters Associated with a Deployment

1. Navigate to Automation > Deployments.
2. Select a Workflow, and then click the associated deployment you want to view by clicking **View**.
3. Click **Parameters** tab.
4. Perform the following:
    - Free-form text entry: Type text as necessary.
    - Select a value: Type a “.” in the field to view a list of all built-in metadata, custom metadata, and policies available.
    - Enter at runtime: Select this check box if you want to enter a parameter when you run the Workflow.
    - Click **Restore Defaults** if you want to restore the values you typed at the step level. If you modify a value at the step level that is associated with a deployment, to update the parameter value at the deployment level, you must click **Restore Defaults**.

➤ If any deployment parameters are set to be runtime, you cannot give the deployment a schedule.

5. Click **Save**.

➤ You can run the Workflow from the current deployment by clicking on the “Would you like to run the Workflow now?” link.

Viewing Rules Associated with a Deployment

1. Navigate to Automation > Deployments.
2. Select a Workflow, and then click the associated deployment you want to view by clicking **View**.
3. Click **Rules** tab.

Deleting a Deployment

If you delete a Workflow that has associated deployments, the associated deployments will be deleted automatically with the Workflow.

1. Navigate to Automation > Deployments.
2. Select a Workflow, and then click the associated deployment you want to view by clicking **View**.
3. Click and confirm delete.

Assigning Deployments to Roles

To change permissions for a deployment, select or clear the check boxes to grant or revoke read, write, or execute access. For more information on changing permissions as an administrator, see Roles on page 167.

1. Navigate to Automation > Deployments.
2. In the Policies pane, click the policy you want to view.

   The Attributes tab displays.
3. Click the **Roles** tab.
4. Select or clear the **Read**, **Write**, or **Execute** check boxes appropriately.
5. Click **Save**.

Run

Running Workflows provides the following benefits:

- Drives standardization in asset management across a data center.
- Empowers Data Center Administrators to define and share best practices.
- Simplifies and automates common administrative tasks. Workflows separate tasks into reusable steps that offer a simple, holistic view of how data center administration tasks are performed. Workflows have the unique ability to run in any environment without any modifications to the underlying automation code.
- Reduces human error when performing data center administration tasks. When common tasks are automated, the risk of human error is reduced, and consistency is driven across an enterprise.

Run automation by selecting a deployment and target. If a deployment contains runtime-specified parameters, they can be entered, otherwise the previously configured parameter values will be displayed.

Running a Workflow

There are several ways you can run a Workflow:

- From within a Workflow or a Deployment, after you create, edit, and then save a Workflow or a deployment, click the “Would you like to run the Workflow now?” link.
- From the Automation > Run screen.

1 Navigate to Automation > Run.

2 Select a Workflow, a deployment name, and the target on which you want to run the Workflow. Set any runtime parameters. See Setting Parameter Values at Runtime on page 54.

Setting Parameter Values at Runtime

If there are runtime parameters, they should be assigned now. All other parameter values will be displayed for you to review before executing the Workflow. It is not possible to change non-runtime parameters at this time.

Any Workflow with runtime parameters cannot be scheduled, or triggered by a rule; Workflows with runtime parameters must be run manually using Automation > Run. You will be prompted to set only the runtime parameters, and will then click **Run** to start the Workflow.

1 Double-click the target to display the associated Runtime Parameters and to display the **Run Workflow** button.

2 Click **Run Workflow**.

3 Navigate to Console or History to view more details about the Workflow’s progress.

Console

The Workflow console provides a real-time view of what automation is currently running on servers in your environment. In the output area, you can see “started” and “completed” as each Workflow step completes. Output and error messages from the running steps can be used to debug problems in the automation code.

Viewing Workflows

1 Navigate to Automation > Console.

2 Search for the desired Workflow using the Filter box.

3 Click to select a Workflow.

4 View the Workflow in the Output area as it is running. When the Workflow is finished running, a green check mark displays. After some time, completed Workflows will be removed from the Console view, unless it has been selected to view details. All completed Workflows can be found in the History view.

History

The Workflow run history view provides an audit trail of who ran automation on servers in your environment. It also provides a useful debugging tool for creating automation Workflow steps. Output and error messages are captured from all steps in the running Workflows. History information is added only once a step is complete.

Figure 5 Workflow History

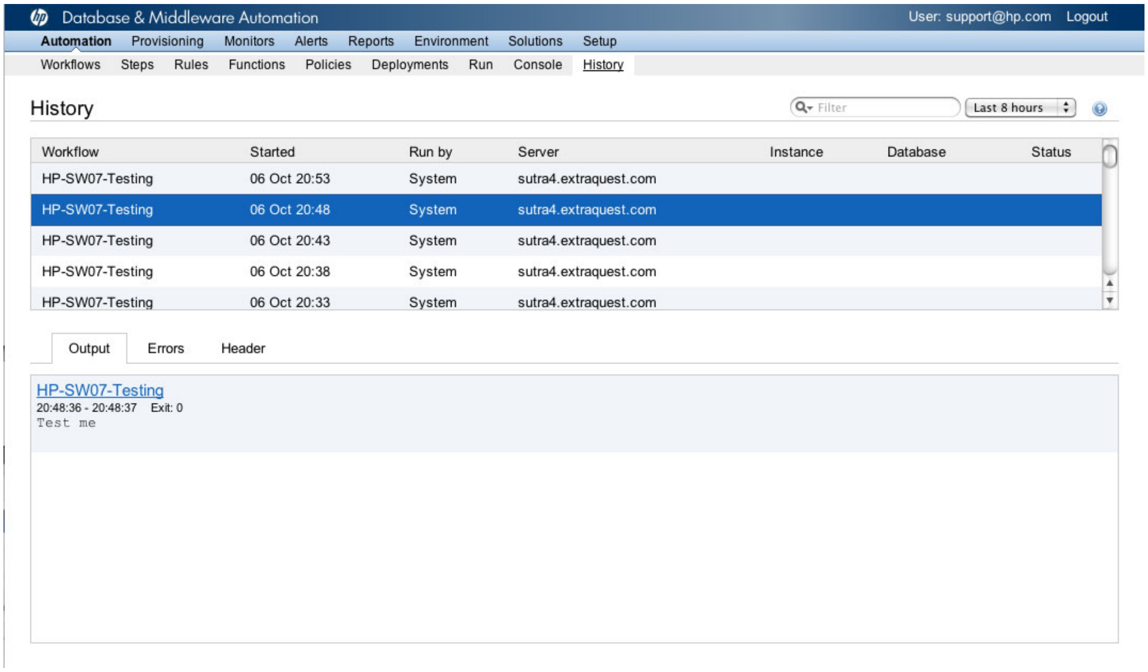

History

Workflow	Started	Run by	Server	Instance	Database	Status
HP-SW07-Testing	06 Oct 20:53	System	sutra4.extraquest.com			
HP-SW07-Testing	06 Oct 20:48	System	sutra4.extraquest.com			
HP-SW07-Testing	06 Oct 20:43	System	sutra4.extraquest.com			
HP-SW07-Testing	06 Oct 20:38	System	sutra4.extraquest.com			
HP-SW07-Testing	06 Oct 20:33	System	sutra4.extraquest.com			

```
HP-SW07-Testing
20:48:36 - 20:48:37   Exit: 0
Test me
```

Viewing Workflow Execution History

1 Navigate to Automation > History.

2 Search for the desired Workflow using the Filter box.

3 Click a Workflow to view the history on the Output, Errors, or Header tabs.

➤ Click on any step name within the these tabs to view specific step details.

Viewing the Output tab

The Output tab tells you which step you are at within a running Workflow. If the Output tab states: “No steps have finished running” this means that the Workflow has started but no steps have completed.

Viewing the Errors tab

The Errors tab displays any errors that occur during a Workflow execution.

Viewing the Header tab

The Header tab is only applicable for Script steps, where it displays data using the function.

4 Provisioning

This chapter contains the following topics:

Overview

Server provisioning is the process of turning a bare metal machine into a functional, managed server. The process consists of four steps: discovery, template assignment, installation, and post provision Workflow.

The primary driver of provisioning is an industry-standard provisioning protocol called PXE (Pre Execution Environment).

Users must have write access in order to add a server to an organization. See Roles on page 167 for more information on roles and permissions.

Discovery

In order to provision a bare-metal machine, it must first be discovered on the network. PXE uses DHCP to hand out the location of the TFTP server and the file to download from the TFTP server. The DHCP server must be configured appropriately in order to tell the new machine what server to talk to. The following example is from a `dhcpd.conf` file.

A complete DHCP subnet configuration:

```
subnet 10.10.1.0 netmask 255.255.255.0 {
    range 10.10.1.200 10.10.1.254;
    option broadcast-address 10.10.1.255;
    option domain-name-servers 10.10.1.243;
    option domain-name "localdomain";
    option routers 10.10.1.1;
    next-server 10.10.1.221;
    filename "pxeprovision.0";
}
```

Of special interest are the lines:

```
next-server 10.10.1.221;
filename "pxeprovision.0";
```

The next-server field is the IP address of the Expert Engine, and the filename is the actual file that PXE will attempt to download. The TFTP server must be listening on UDP 69. On many systems port 69 is a privileged port so the Expert Engine must run as root or have appropriate port forwarding configured. Port forwarding is not required on Windows operating systems.

The following example will add an iptables rule to forward from port UDP 69 to UDP 6969 (The port on which the Expert Engine will start the TFTP server if it cannot start on port 69).

```
iptables -t nat -A PREROUTING -i eth0 -p udp --dport 69 -j REDIRECT
--to-port 6969
```

Template Assignment

A server template consists of several components: the image location within the software library, the post provision Workflow, and the parameters associated with that Workflow.

The creation of the OS image is not currently within the scope of HP DMA, however there are many third party tools that can be used to create these images. Kickstart can be used for RHEL and WAIK (windows automated installation kit) can be used for Windows operating systems.

Once an image is created, it is uploaded to the software directory. This can be done through the web interface if the files are smaller than 100MB. If the files are larger than 100 MB, they can be copied directly into the software directory of HP DMA. The default path is `$DATA_PALETTE_HOME/Web/tomcat/Webapps/ROOT/software/`. Files loaded into this directory are considered part of the software library and can be used with provisioning.

A template can be assigned to a machine in one of two ways: through the machine view or the template view. Both are similar in functionality, but the template view allows a single template to be assigned to many selected machines, while the machine view only works on a single machine at a time.

Clicking on a machine that has no template assigned allows the user to select a template and enter the parameters required for that template to be applied. If a template has a Workflow, a target Organization is required in order to execute the Workflow and transition automatically from a bare metal machine to a fully provisioned and managed server. The parameters form supports input from Organization and server level policies.

Installation

Once a template is applied to a machine, the provision process will begin the next time the machine boots. In most environments, this will be within five minutes of template assignment, although this can be lowered if the DHCP server supports additional timing parameters. The time to provision a machine depends on the size of the template and the capacity of the network. The status of a particular machine’s provision process can be viewed by clicking on the machine. If a template has been applied to a machine, instead of showing the parameters it will show the status and files that are being transferred to the machine.

Post Provision Workflow

Once the machine has been fully provisioned, it will execute its configured post provision Workflow. In many cases this Workflow will set the IP address, join a domain, or install a set of applications. This post provision Workflow may also be used to force a machine into compliance and create the necessary exceptions for the role of the server. Once the Workflow has started execution, the machine will no longer be visible from the provisioning screen as it has transitioned to a managed server object and can now be visible in the environment. See History on page 55 for more information.

Provision History

Although machines will no longer be visible from the provisioning screen, the history of each provisioned machine is stored in the repository and can be viewed from the history screen. Historical provision information can be viewed by template or by machine. Clicking on a single historical provision will display the states that the machine has passed through and the time it took to provision the machine from discovery to managed object.

Machines

Provisioning allows you to view bare metal machines on your network and assign an operating system template to them. The machine will then install the assigned operating system. The DHCP server configuration must have the next-server field set to the address of the Expert Engine and the filename field set to pxeprovision.

1. Navigate to Provisioning > Machines.
2. Select a machine.
3. Select a template with which to provision the machine.
4. Select **Target Organization** and other parameters as needed.
5. Click **Save parameters**.

   The Provision button becomes active.
6. Click **Provision**.

Templates

Server templates are network based configurations for provisioning bare-metal machines (including virtual machines). The template configuration directory contains the operating system image that will be used to provision a given machine. After the operating system is installed an optional automation Workflow can run to configure the server. Typical post-install configuration includes setting a network address, login credentials, and installing additional software.

Creating a New Template

1. Navigate to Provisioning > Templates.
2. Click **New template**.
3. Type the following:
   - Name
   - Image directory
   - Workflow
4. Type any notes in the Documentation area.
5. Click **Save**.

➤ The Verified column changes to “yes” for templates that have configured at least one machine. Once the template has been modified, the Verified value changes to “no.”

Editing an Existing Template

1. Navigate to Provisioning > Templates.
2. Select a template to edit.
3. Click **Edit template**.
4. Type the following:
   - Name
   - Image directory
   - Workflow
5. Edit or add notes in the Documentation area.
6. Click **Save**.

Configuring Runtime Parameters

1. Navigate to Provisioning > Templates.
2. Select a template to edit.
3. Select a machine by selecting a check box. To provision multiple machines select the check box for each machine to provision with a specific template.
4. Type the following:
   - Target Organization: You only need to assign an Organization if the template has an associated Post Provision Workflow. (If there is no associated Post Provision Workflow, the Target Organization field does not display.) If the Post Provision Workflow has configurable parameters, you may overwrite those values with either text or Policy values. (Policy values are accessible by typing a "." in the field.)

5. Click **Save parameters**.

   The Configured column changes to “yes,” and the Provision button is active.

6. Click **Provision**.

Software

The software library contains packages that can be installed as part of a provisioning Workflow. Example library files include: Linux rpms, Windows msi executables, and simple configuration files. Software under 100 MB can be uploaded to the library from the Web Server. Larger files can be placed directly into the web server file system.

➤ Only administrators have the right to delete and upload files. Non-administrators can only download files.

Searching for Software

You can perform a real-time filter on any software. Type what you are searching for in the filter field and see the results display as you type.

Uploading Software

1. Navigate to Provisioning > Software.
2. Click **Browse**.
3. Locate the software you want to upload.
4. Click **Upload**.

Creating New Directories within the Software Folder

You can also create new directories within the software folder.

1. Navigate to Provisioning > Software.
2. Click **New** (folder icon at the top of the screen).
3. Type folder name.
4. Click **Create** or click **X** to cancel.

Deleting Previously Uploaded Software

1. Navigate to Provisioning > Software.
2. Locate the software you want to upload.
3. Click and confirm delete.

Download Software

See Policy Attribute Types on page 48 for more information on the Software Library.

1. Navigate to Provisioning > Software.
2. Locate the software you want to download.
3. Click the software.

History

Each machine builds a provision history during the installation process. The history is updated each time the status changes. Some fields may be empty if the data has not yet been gathered during the provisioning process. The template view groups provisioned machines by the operating system that was installed.

Verifying Successful Provisioning

- If provisioning was successful, "Agent" will display a link to the Environment > Agent screen.
- Status should be "online." Navigate to Automation > History. You can see if your Workflow has been run, if you selected a Workflow.
- You can also check in Environment > Dashboard, and you can view your new environment.

5 Monitors

This chapter contains the following topics:

- Graphs on page 63
- Database on page 65

Graphs

Monitors are graphs of statistical information gathered from the Agents running on remote servers. Within a graph, selecting a category and time frame will graph the available data for the selected resource (server, instance, or database). If there is no data for a statistic available for the selected time frame, the graph will not be drawn. Red thresholds are drawn as a dashed red line and only appear when either a yellow or red threshold is active.

Viewing Selected Monitor Statistics

1. Navigate to Monitors > Graphs.
2. Double-click a resource to view key performance indicators. Once you double-click a resource, use the drop-down lists to customize various criteria for each resource. The following information is available for each resource:
   - Server Level
   - CPU
   - RAM
   - Swap
   - I/O
   - Disk Space
   - Network
   - Top Processes
   - Custom
   - Instance Level
   - Database Server
   - Custom
   - Database Level
   - Database
   - Custom

3. Select a time frame:
   - Last 24 hours
   - Last 7 days
   - Last 30 days
   - Custom

Creating Custom Statistics

1. Navigate to Monitors > Graph.
2. Select an object, such as server, instance, or database.
3. Select **Custom** from the **Graph** list.
4. Select multiple statistics graph.
   - Use Shift to select one or more statistics to graph together.
5. Click **Graph**.

Creating Custom Time Frames

1. Navigate to Monitors > Graphs.
2. Select an object.
3. Select **Custom** from the **Time frame** list.
4. Select a **Start date**.
5. Select and **End date**.

➤ The End date cannot occur before that Start date. If you attempt to type an End date prior to the Start date, the Start date will be changed automatically to one day before the selected end date.

Saving Graphs

Using Save Image As

1. Navigate to Monitors > Graphs.
2. Navigate to the graph that you want to save.
3. Right-click the graph and then click **Save image as**.

Using Drag and Drop

1. Navigate to Monitors > Graph.
2. Navigate to the graph that you want to save.
3. Select the graph, and then drag it to the desktop.

Printing Graphs

Once you have saved the graph, use your browser's print command to print the desired graph.

Database

Database monitors are a series of snapshots of information about an instance or database gathered by the Agent. They can be used to perform root cause analysis of problems detected in your database environment. For example, when a server is experiencing high load, you can view the SQL statements that ran on a particular server. You can then look at any locks or waits that seemed to affect the SQL. Tuning the SQL can then decrease the load on the affected server.

Filtering Database Monitors

You can perform a real-time filter on database monitor results. Type what you are searching for in the Filter field and see the filter results display as you type.

Viewing Database Monitors

1. Navigate to Monitors > Database.
2. Navigate to the desired database you wish to monitor.
3. Perform one of the following tasks:
    - Click on the specific graph you want to view.
    - Navigate to the desired database you wish to monitor. Select a date using the calendar or instance.

6 Alerts

The alerting and ticketing modules work together to provide a foundation for the analysis and automation capabilities of HP DMA. Tickets are organizationally helpful in and of themselves. The HP DMA functional areas are meant to work together to make ticket escalation an automated process that ensures all issues follow a chain of command, which is custom-defined by you for your Organization's needs.

When HP DMA generates an alert, it is based on monitor thresholds that you define, or via a Workflow/Deployment (alert step), instance logs, or various internal alerts. Then HP DMA generates a ticket that identifies the ticket's subject, displays a severity level depending what you defined, and sets the ticket's status to "new." The generation of a ticket triggers the escalation path that you defined, and the alerts follow that specific chain of command until someone on the escalation path list responds to the trouble ticket.

This chapter contains the following topics:

- Tickets on page 67
- Escalation on page 69
- Instance Log on page 70
- Thresholds on page 72

Tickets

Searching for Tickets

You can search for a specific ticket by the ticket's ID number.

1. Navigate to Alerts > Tickets.
2. Type the desired ticket's ID number in the search box.
3. Click **Find ticket**.

Viewing Specified Tickets

1. Navigate to Alerts > Tickets.
2. Use the drop-down boxes to specify the following criteria:
   - Time frame
     - — Days: Returns all days.
     - — Last 24 hours
     - — Last 3 days

- Last 7 days
- Last 30 days

- Status
  - Status: Returns all statuses.
  - New
  - In Progress
  - Closed
- Resource: Displays a list of all Organizations and Instances. You can also select **Resource** to return all resources.
- Filter: You can perform a real-time filter on any ticket or on any column in the Ticket screen. Type what you are searching for in the **Filter** field and see the filter results display as you type.

Viewing Ticket Details

1. Navigate to Alerts > Tickets.
2. Click Alert's **ID** or Alert's **Subject** to view details.
3. For additional Alert details, navigate to the History tab.

Editing Tickets

1. Navigate to Alerts > Tickets.
2. Select the ticket you wish to update.
   - You can update tickets by modifying the severity
   - Select a Status from the Status list.
     - New
     - In Progress
     - Closed
   - Navigate to the History tab to add a new note.
3. Click **Save**.

Updating Bulk Ticket Status

1. Navigate to Alerts > Tickets.
2. Select the check box(es) of the tickets you wish to update.
3. Click **Update tickets**.
4. Select a Status from the Status list.
   - New
   - In Progress
   - Closed

5. Type any desired comments.
6. Click **Save**.

Escalation

HP DMA uses escalation paths as the mechanism that informs users of tickets that have been created. For every step in an escalation path a user, or set of users, are contacted. If, after a specified time, the ticket has not been addressed by changing the status to “In progress” or “Closed,” the escalation proceeds to the next step.

Before you define your Organization’s escalation path, you must define severity levels for your Organization. After defining severity levels, you need to specify the hierarchy of people who will receive alerts within your Organization by defining an escalation path for each severity level.

Viewing an Organization’s Severity Levels

1. Navigate to Alerts > Escalation.
2. Hover over an Organization to view the severity levels and contacts associated with each Organization.

Setting Severity Levels

Severity level 9 is always active as a default severity level, but you can modify the name.

1. Navigate to Alerts > Escalation.
2. Select an Organization.
3. Select the check box next to severity levels 1-9 to activate the desired level.
4. Type a severity.
5. Select **Warning** if this is a warning severity.
6. Click **Save**.

Defining Escalation Paths

1. Navigate to Alerts > Escalation.
2. Select an Organization.
3. Click **Escalation** of the severity level you wish to define.
4. Click **Add Step**.
5. Specify a time in minutes to wait to escalate the ticket to the next level.
6. Click **Add Contact**.

The **Add Contact** link will only display if there are contacts that are defined that can be selected.

7. Click **Save**.

Deleting an Escalation Step

1. Navigate to Alerts > Escalation.
2. Select an Organization.
3. Click **Escalation** of the severity level you wish to open.
4. Click ⛔ to remove an escalation step from the escalation path.
5. Click **Save**.

Deleting an Escalation Contact

1. Navigate to Alerts > Escalation.
2. Select an Organization.
3. Click **Escalation** of the severity level you wish to open.
4. Click ⛔ to remove an escalation contact from the escalation path.
5. Click **Save**.

Instance Log

The HP DMA Agents collect errors from database log files and return them to the Expert Engine. When an expression defined here matches an error message, an alert is generated or ignored depending on the current alert settings. For example, if the error message is "ORA-1234 Tablespace full," an alert expression of "ORA-1234.*" would match the error and generate an alert.

Viewing Instance Expressions

1. Navigate to Alerts > Instance Log.
2. Select an Organization, Server, and then double-click on the instance.
3. View the Ignore and Alert Expressions for the desired instance.

Creating a New Ignore Expression

1. Navigate to Alerts > Instance Log.
2. Select an Organization, Server, and then double-click on the instance.
3. Click **New Ignore Expression**.
4. Type the new ignore expression.
5. Click **Save**.

Creating a New Alert Expression

1 Navigate to Alerts > Instance Log.

2 Select an Organization, Server, and then double-click on the instance.

3 Click **New Alert Expression**.

4 Type the new alert expression.

5 Select a severity level.

6 Click **Save**.

Copying Expressions

When copying an expression from one Organization's instance to another that does not have the same severity levels (severity levels for alert expressions) defined, the lowest severity will be selected (the default severity level of 9 will always be active, so something will always be selected).

1 Navigate to Alerts > Instance Log.

2 Select an Organization, Server, and then double-click on the instance.

3 Click **Copy Expressions**.

4 Select the check box next to the instances or the server (which selects all instances for that server) to which you want to copy expressions.

5 Click **OK**.

Testing Expressions

This feature is helpful when you want to verify that an expression using the ".*" wildcard is correct.

1 Navigate to Alerts > Instance Log.

2 Select an Organization, Server, and then double-click on the instance.

3 Click **Test Expressions**.

4 Type the Expression you want to test.

5 Type text.

6 Click **Test**.

7 View the result.

Removing an Expression

1 Navigate to Alerts > Instance Log.

2 Select an Organization, Server, and then double-click on the instance.

3 Click ⛔ to remove an expression.

4 Click **Save**.

Thresholds

Monitor thresholds are set at each target level. Before you define your monitor thresholds, ensure that you have already defined severity levels for your Organization. Secondly, ensure that you have defined an escalation path for each severity level.

Now you should define your monitor thresholds, indicated by two colors: yellow and red.

Yellow is the warning level that indicates a possible problem so you can investigate the issue before it becomes a major problem. Red is the critical error level where there is definitely a problem that needs to be investigated. When the statistic value exceeds the defined threshold, an alert is sent. Generally, the yellow level will generate a Severity 3 alert, while the red might generate a Severity 1 alert.

You can set monitor thresholds at the server, instance, or database level, based on which statistics you need to monitor for each resource. Depending on which resource you select, different statistics are available. For example, only Oracle statistics will appear for Oracle databases.

Once you have defined your monitor thresholds, you have to activate them by selecting the severity level to which they should be associated. The selected severity level will trigger an alert that follows the escalation path previously defined.

Defining Monitor Alert Thresholds

1. Navigate to Alerts > Thresholds.
2. Navigate to the desired Server, Instance, or Database.
3. Type the desired setting for each statistic you want to define.
4. Select the severity for each statistic that is defined from the previous step. If the severity level is “Not active,” then even if your specified threshold is met and exceeded, you will not receive a ticket or notification of any form.
    - Yellow Threshold and Severity
    - Red Threshold and Severity

➤ There are some default thresholds already set within HP DMA. You can edit the default thresholds.

Modifying Monitor Alert Thresholds

1. Navigate to Alerts > Thresholds.
2. Navigate to the desired Server, Instance, or Database.
3. Edit the desired setting for each statistic you want to modify.
4. Select the severity for each statistic that is defined from the previous step. If the severity level is “Not active,” then even if your specified threshold is met and exceeded, you will not receive a ticket or notification of any form.
    - Yellow Threshold and Severity
    - Red Threshold and Severity

➤ There are some default thresholds already set within HP DMA. You can edit the default thresholds.

5 To view a Ticket's Severity Level, navigate to Alerts > Tickets and view the desired ticket.

Copying Thresholds

1 Navigate to Alerts > Thresholds.

2 Select an Organization, Server, and then double-click on the instance.

3 Click **Copy Thresholds**.

4 Select the target resource to which you want to copy the thresholds.

➤ You can select all servers under an Organization by selecting the Organization (all instances by selecting the server, all databases by selecting the instance). Only Organizations can copy thresholds to other Organizations. For Instances and Databases, only Instances and Databases of the same type will be available to copy to thresholds (for example, Oracle to Oracle).

5 Click **Copy**.

7 HP DMA Reports

This chapter contains the following topics:

Introduction

HP DMA provides a variety of useful reports to analyze your database or general IT environment. These reports are especially useful as an aid in identifying areas in which to apply automation. The HP DMA reports are broken into four distinct categories: Operating System (OS), Microsoft® SQL Server, Oracle, and DB2. Each report category begins with a categorical report summary, which is followed by detailed information for each report within each category. Prior to the summary and detailed report specifications, some usage guidelines will help you understand some of the general concepts behind all of the reports in HP DMA.

General Usage Guidelines

As HP DMA monitors your server and database environment, it collects a vast amount of information. This information is stored in its repository and is the basis for all of the HP DMA reports. Therefore, HP DMA reports do not run real-time against your current data environment. Instead, they run and return results from the data that was previously collected from that environment that is stored in the HP DMA repository. Reporting this way has some enormous advantages, including the ability to understand what has happened in the past as well as understanding upcoming trends.

General Report Format

Reports in HP DMA can take a variety of formats, from tabular, to graphical, to a mixture of each. Each report's definition specifies a unique format as well as how to interpret the results. Moreover, since a given report may be specified to return results covering the period of a single day or the period of several days, months, or years, the report mechanism in HP DMA uses an innovative way to display results over diverse time periods and also scales appropriately the time axis for the graph.

Report Availability

One consequence of running against collected, non real-time information is that a report may be unavailable (and will not appear in the reports list) if the HP DMA Agent(s) on which it depends is not running or not collecting data. The detailed information for each report specifies which Agent modules and which routines within those Agent modules must be running in order to be able to run the report.

Common Parameters: Time span

- Last 24 hours
- Last 3 days
- Last 7 days
- Last 30 days
- Last 90 days
- Custom

Report-Specific Parameters

In addition to the common user-defined report criteria, such as start time and end time, each report may have parameters that are specific to that report. Such parameters will have a default value and will be described by each report.

The HP DMA Graphing Mechanism

HP DMA runs collections at user-defined intervals - usually every five minutes. That implies that there will be 12 data points per hour for each collected item of information (of which there are thousands). To retain the accuracy of its collected information while providing reasonable reporting format and performance, the following mechanism is used: All of the individual data points (of a common collection type) within a one hour time span are averaged yielding exactly one value. This value is represented on a graph (at the right edge of each hour increment rather than in the middle of each increment). In the case of a line graph, all of these points are then linked together.

Interpreting Graphical Information

Given this mechanism, it should be noted that:

- Exceedingly high or low values within the one-hour period will adversely affect the averaged value for that period.
- Missing information within the one-hour period will result in that period being omitted and the graph being drawn from the last real value to the next real value. That means that there may be unexpected jumps or spikes in the graph if missing data is encountered.
- Far right and far left graph endpoints may be misleading because the graph is being drawn from "nothing" to the first real data point or from the last real data point to "nothing."

Running a Report from the HP DMA Web Server

Figure 6 shows the HP DMA reporting dashboard.

Figure 6 Report Dashboard

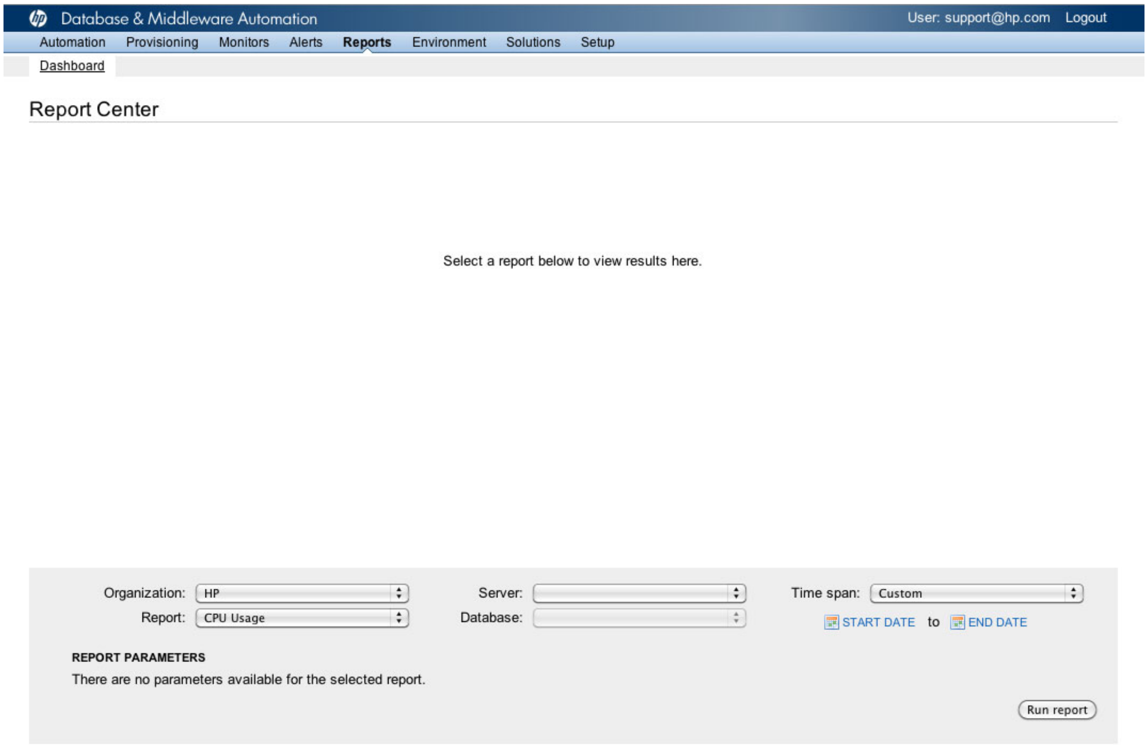

Report Center

Select a report below to view results here.

Organization: HP

Report: CPU Usage

Server:

Database:

Time span: Custom

START DATE to END DATE

REPORT PARAMETERS

There are no parameters available for the selected report.

To run a report, follow these steps:

1. Log in to the Web Server.
2. Navigate to Reports.

   When you generate a report, you must define the Organization, the report name, and some additional report-specific criteria.
3. Select the server.
4. Select the database.
5. Select the time span.
6. Click **Run report**.

OS Reports

HP DMA provides the following OS reports:

- OS: Network Usage Report on page 79

  Shows the network bandwidth usage over a period of time. Spikes or consistently high usage may indicate a database performance problem caused by network performance problems and not the database itself.

- OS: Storage Space Projections Report on page 80

  Projects your current disk usage habits into the future to determine when you will run out of space on a physical or logical storage device.

- OS: Load Average Report on page 83

  Shows the 1 minute, 5 minute, and 15 minute load on the OS over a period of time. This allows you to understand the overall system load of your servers and aids in analyzing performance problems.

- OS: Memory Usage Report on page 84

  Shows the amount of RAM used over a period of time. Memory usage is one key element in overall system performance. Knowing how your server's memory is being utilized will help you analyze and tune the performance of your systems.

- OS: CPU Usage Report on page 85

  Shows the amount of CPU being used over a specified time period. CPU usage is another indicator of a system's performance and knowledge of it will help you analyze and tune your servers.

- OS: Swap Usage Report on page 86

  Shows the amount of swap consumed over a period of time. Swap usage is also indicative of a system's performance. Knowledge of swap utilization will help you analyze and tune your servers.

- OS: Disk Busy Report on page 87

  Shows the percentage of time the hard disk is performing work (a read or write) rather than sitting idle. You can use this information to spread database datafiles over several disks to increase performance.

OS: Network Usage Report

The Network Usage Report provides network throughput information across all network interfaces. Specifically, this graph indicates the maximum network activity, for inbound and outbound traffic, over the time period specified. Large, sustained amounts of network traffic could indicate a potential network bottleneck.

➤ The MTU (Maximum Transmission Unit) is the size of the largest datagram that can be sent over a network (the value can vary depending on the Network Type, for example the default MTU value for Ethernet Networks is 1500). This report uses an MTU value of 1500.

Required agent modules and routines:

- OS Monitors (frequent)/Network Information

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

- X-axis

  Time over which the report was run (Time span). For time periods of one day, the x-axis increments are in hours and the information presented covers that day. For periods greater than a day, the x-axis increments may be in hours, days, months, etc. depending upon the most appropriate way to represent the data clearly.

- Y-axis

  Network traffic, in Kilobytes/second for each network interface (both inbound and outbound).

OS: Storage Space Projections Report

The Storage Space Projections Report analyzes your data usage patterns to predict when you will run out of space on physical and logical storage devices. Specifically, it allows you to:

- Discover which mountpoints are running out of space and when they will be full.
- Discover how fast mountpoints/databases/tablespaces are growing.
- Discover how each database is affecting a given mountpoint.
- Discover how each mountpoint's storage is being used by databases.
- Discover tablespace/data file growth rates and predicted time until full.

The Storage Space Projections Report is somewhat unique in that it allows you to see information about mountpoints, databases, tablespaces (Oracle), and data files (SQL Server). Each of these different views shows information in Primary Report Information.

- Mountpoints

  The mountpoints view shows all of the mountpoints for the given Organization and server together with a variety of storage information (described under Primary Report Information). To see all of the databases on this server regardless of mountpoint, select Databases from the drop-down list at the top of the report.

- Databases

  The databases view shows all of the databases for the given Organization and server, regardless of what mountpoint(s) they may occupy. Clicking on any one of these databases will display information about its tablespaces (Oracle) or data files (SQL Server). To see all of the mountpoints or databases again, choose Mountpoints or Databases respectively from the drop-down list at the top of the report.

Required agent modules and routines:

- DB2 Container Mount Point / DB2 Analytics (once a day)
- DB2 Tables / DB2 System Catalog-5
- DB2 Tablespaces / DB2 System Catalog-5
- MS Alt Files / MS Master DB Configuration Info
- Oracle DBADataFiles / Oracle DDL5
- Oracle DBAFreeSpace / Oracle DDL5
- Oracle File$ / Oracle DDL2
- Oracle TS$ / Oracle DDL3
- Oracle V$DataFile / Oracle DDL5

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

The usage statistics are available for mountpoints as well as any predictive trends are available for mountpoints, databases, tablespaces (Oracle), and data files (SQL Server). The following describes this information:

- Mountpoint/Tablespace/Data File Icon

  Selecting this icon displays a graph showing the historical usage trend on the mountpoint, tablespace, and datafile. For mountpoints, several lines are displayed showing database usage as well as total usage (database + non-database files).

- Mountpoint/Database/Tablespace/Data File Name

  Name of the mountpoint (Example: `/` or `/var` or `C:`), database (Example: `/dbinst2/sample`), tablespace (Example: `USERSPACE`) or data file (Example: `PAT01\SW1/sql_12`).

- Capacity (MB)

  The total capacity of the mountpoint/tablespace/data file (depending upon the view that you are looking at) in megabytes.

- Time Left (Months)

  Time Left (Months) is the amount of time remaining before the mountpoint/tablespace/data file (depending on your view) is filled to capacity. For mountpoints, it assumes that all (if any) tablespaces/data files will extend when filled. This value is computed by extending both the Data Usage and Total Usage regression lines to Capacity, and choosing the earlier of the two future points. This approach ensures that the "Time Left" prediction is the most conservative (earliest) projection possible.

- Overall Usage MB (mountpoint only)

  The amount of space in megabytes that is being used on the mountpoint. This overall space is a combination of both database usage (if any) and non-database usage.

- Overall Usage % (mountpoint only)

  The percentage of space being used on the mountpoint, by both database and non-database files.

- Overall Usage MB / Day (mountpoint only)

  The overall growth (database and non-database files), in megabytes per day, of the mountpoint.

- Overall Usage % / Day (mountpoint only)

  The percentage overall growth of an empty mountpoint that starts growing at 1% per day relative to its capacity.

- Databases Using (mountpoint only)

  A comma separated list of databases that are currently using space on the mountpoint. Clicking on any one of them will take you to a list of all databases using that mountpoint, along with their respective usage statistics relative to the mountpoint.

- Database Usage MB (mountpoint only)

  Value indicates the total amount, in megabytes, being used by databases on the mountpoint.This value will always be less than or equal to total usage, since non-database files may also be present on the mountpoint.

- Database Usage % (mountpoint only)

  The percentage of the mountpoint being used by databases.

- Database Usage MB / Day (mountpoint only)

  The collective amount, in megabytes, that all database usage is growing per day on the mountpoint. For example, 5MB would indicate that the mountpoint's space is being consumed an average of 5MB per day by database related files.

- Database Usage % / Day (mountpoint only)

  The overall daily percentage growth, relative to its capacity, associated with databases on the mountpoint. For example, 7% would indicate that space is being consumed by database files at a rate of 7% per day.

- Allocated Space MB (tablespace only)

  The amount of space, in megabytes, that has been allocated to an Oracle tablespace's datafiles. This value is always less than its capacity and represents an allocation, not necessarily usage, of space by that tablespace.

- Allocated Space % (tablespace only)

  The percentage of space currently allocated to the tablespace relative to its capacity.

- Allocated Space MB / Day (tablespace only)

  The growth in megabytes per day, of the tablespace's allocated space.

- Allocated Space % / Day (tablespace only)

  The daily percentage of growth, relative to its capacity, of the tablespace's allocated space.

- Used Space MB (tablespace only)

  The amount of space, in megabytes, that is actually being used within the datafiles of an Oracle tablespace.

- Used Space % (tablespace only)

  The percentage of space, relative to the tablespace capacity, currently being used by the tablespace's datafiles.

- Used Space MB / Day (tablespace only)

  The growth, in megabytes per day, of data being used within a tablespace.

- Used Space % / Day (tablespace only)

  The daily percentage of growth, relative to its capacity, of data being used within the tablespace.

- Used MB (SQL Server data file only)

  The amount of space, in megabytes, that is being used by the SQL Server data file.

- Used Space % (SQL Server data file only)

  The percentage of space, relative to its capacity, currently being used by the SQL Server data file.

- MB / Day (SQL Server data file only)

  The growth, in megabytes per day, of the SQL Server datafile.

- % / Day (SQL Server data file only)

  The percentage of growth, per day, of the SQL Server data file.

OS: Load Average Report

The Load Average Report is a graphical report that shows data for 3 different sampling intervals: 1 minute system load, 5 minutes system load, and 15 minutes system load over a specified period of time. Each average is based upon a calculation involving I/O, CPU utilization, memory utilization, swap utilization and other operating system performance indicators.

Ways to interpret this graph:

- If you see a high 15 minute load average but a low 1 minute load average, it means that the server being monitored has experienced some CPU intensive processing over the last 15 minutes that just recently stopped and the system has gone back to normal.
- If all three stats are similar (low, medium or high), the server is under a consistent load.
- Spiked graphs generally show expensive operations being performed on the server (Example: an expensive query).

Required agent modules and routines:

- OS Monitors (frequent)/Load Average

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

- X-axis

  Time over which the report was run (Time span). For time periods of one day, the x-axis increments are in hours and the information presented covers that day. For periods greater than a day, the x-axis increments may be in hours, days, months, etc. depending upon the most appropriate way to represent the data clearly.

- Y-axis

  The actual value of the load average for the 1 minute, 5 minute and 15 minute samples. Each data point represents a distinctly color-coded load average taken over a 1 minute interval, 5 minute interval, and a 15 minute interval.

OS: Memory Usage Report

Memory Usage displays the average percentage of RAM utilized (0 to 100%) over a specified period of time. It also shows maximum percentage of RAM utilized over that same time period. Use this report to understand when, in the past, an instance has experienced shortages. Such information will then help you to more accurately predict the RAM demands for your database and potential adjustments you may need to make relative to hardware or process scheduling.

Required agent modules and routines:

- OS Monitors (frequent)/Memory, Swap, Page, CPU

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

- X-axis

  Time over which the report was run (Time span). For time periods of one day, the x-axis increments are in hours and the information presented covers that day. For periods greater than a day, the x-axis increments may be in hours, days, months, etc. depending upon the most appropriate way to represent the data clearly.

- Y-axis

  Average Memory Utilization as a percentage or Maximum Memory Utilization as a percentage respectively, depending upon the graph.

OS: CPU Usage Report

The CPU Usage Report displays two graphs: the average percentage the CPU is utilized (0 to 100%) over a specified period of time and the maximum percentage the CPU utilized over that same time period. CPU utilization is displayed for user (CPU utilization for user initiated tasks), system (CPU utilization for system tasks) and idle (CPU idle time). Such information helps identify and resolve performance bottlenecks, an invaluable aid when trying to understand performance problems. Typically, performance problems are associated with high CPU usage but interestingly, excessive idle time could also indicate that contention for one or more resources is occurring with the associated performance degradation for those competing processes.

➤ User + System + Idle CPU utilization does not always equal 100 since the values are averaged over time.

Required agent modules and routines:

- OS Monitors (frequent)/Memory, Swap, Page, CPU

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

- X-axis

  Time over which the report was run (Time span). For time periods of one day, the x-axis increments are in hours and the information presented covers that day. For periods greater than a day, the x-axis increments may be in hours, days, months, etc. depending upon the most appropriate way to represent the data clearly.

- Y-axis

  Average CPU Utilization as a percentage or Maximum CPU Utilization as a percentage respectively, depending upon the graph.

OS: Swap Usage Report

The Swap Usage Report displays two graphs: the average percentage of swap space utilized (0 to 100%) over a specified period of time and the maximum percentage of swap space utilized over the same time period. Knowing a server's swap utilization helps to identify excess and frequent swap activity, indicating that memory is over committed or the server is over utilized. Such situations are often associated with performance problems. When swap usage consistently approaches 80% utilized, you should consider reducing memory usage (concurrent programs), or adding more memory (or increasing your swap space).

Required agent modules and routines:

- OS Monitors (frequent)/Memory, Swap, Page, CPU

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

- X-axis

  Time over which the report was run (Time span). For time periods of one day, the x-axis increments are in hours and the information presented covers that day. For periods greater than a day, the x-axis increments may be in hours, days, months, etc. depending upon the most appropriate way to represent the data clearly.

- Y-axis

  Average Swap Usage or Maximum Swap Usage respectively, as a percentage from 0% to 100%.

OS: Disk Busy Report

The Disk Busy Report displays each physical disk device or each mountpoint's activity (read/write access) as a percentage from 0% - 100%. Zero percent means there is no activity on that device/mountpoint while 100% means that the device/mountpoint is fully utilized. Below this graph is a table that shows the mapping of physical device names to system mountpoints (there could be one or more mountpoints per physical device name).

Required agent modules and routines:

- OS Monitors (frequent)/Disk Busy
- OS Monitors (infrequent)/Disk to Mountpoint

Required report criteria:

- Organization
- Server
- Time span

Parameters: None

Primary report information:

- X-axis

  Time over which the report was run (Time span). For time periods of one day, the x-axis increments are in hours and the information presented covers that day. For periods greater than a day, the x-axis increments may be in hours, days, months, etc. depending upon the most appropriate way to represent the data clearly.

- Y-axis

  Disk activity, as a percentage from 0% to 100%.

Microsoft SQL Reports

HP DMA provides the following Microsoft SQL reports:

- Microsoft SQL: Top SQL Report on page 89

  Shows the most expensive SQL statements running against the database.

- Microsoft SQL: Top Wait Processes Report on page 91

  Shows statements in SQL Server that are waiting to run but are being blocked by other processes already running.

- Microsoft SQL: Database Backup Report on page 93

  Shows information about the execution history of scheduled jobs by SQL Server Agent and database backup jobs. This report helps one understand specific information for the various types of backups occurring on a given database environment.

- Microsoft SQL: CPU Load Report on page 95

  Shows how a particular SQL Server instance is handling its process threads on each CPU of a server.

- Microsoft SQL: DB Configuration Check on page 96

  Shows a database's overall health as measured by numerous KPIs (Key Performance Indicators).

- Microsoft SQL: Job Failure Report on page 100

  Shows all of the SQL Server scheduled jobs on a server and indicates which servers are experiencing the most failures. Allows one to drill-down into individual jobs to determine the frequency of failure on each step.

Microsoft SQL: Top SQL Report

The Top SQL Report provides Information regarding the worst performing (i.e., most expensive) queries for any given database Instance. This information is vital in analyzing and correcting the worst performing processes (queries) in order to optimize database performance.

Two tables may be displayed. The primary table displays information about each SQL Server process. By clicking on the process ID, you can get a secondary table that displays additional information about that process.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- MS Trace Current Activity/MS Trace Current Activity

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters:

- Sample Size:
  - — Sample size limits the output results to the specified number of rows.
  - — The default value is `0`.

Primary report information:

- SPID

  SQL Server process ID associated with the SQL statement. By clicking on the data within the SPID column, you will see additional information about that particular process.

- Duration

  Amount of time in milliseconds spent by this process.

- Start Time

  Time at which the SQL statement finished executing.

- Reads

  Number of logical disk reads performed by the server on behalf of this process.

- Writes

  Number of logical disk writes performed by the server on behalf of this process.

- CPU Time

  Amount of CPU time in milliseconds spent by this process. CPU time can occasionally be 0 indicating that this process is in sleeping status.

- Statement Text

  The text of the actual SQL statement associated with this process.

Additional report detail:

- NT User Name

  The Windows user (login) name.

- Application Name

  Name of the client application that created the connection to SQL Server.

- Login Name

  Name of the user, either SQL Server login or the Windows login name in the form domain\username.

Microsoft SQL: Top Wait Processes Report

In Microsoft SQL Server, intensive queries, which are dependent on resources that are blocked or unable to process at that given time, must queue or wait for the required resources to be freed up. Transactions containing the waiting query may hold locks while the query waits for memory or any other resources. In rare situations, it is possible for an undetectable deadlock to occur. Decreasing the query wait time lowers the probability of such deadlocks. This report displays the top waiting processes and will help identify such bottlenecks.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- MS Trace Current Activity/MS Trace Current Activity

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

- SPID

  SQL Server process ID associated with the waiting process.

- Collection Date

  This is the date and time at which the wait process was discovered and populated into the HP DMA repository. Note that the process may have been initiated much earlier than this date but only recently entered a wait state (and this date reflects when you discovered that wait - not when the process was initially created).

- SQL Text

  SQL statement associated with the waiting process. Rarely, this column may be blank. This will only occur for short running SQL Server processes that have already been removed from SQL Server at the time HP DMA asks for its SQL text. For processes of any complexity, this column will be populated with non-blank SQL text.

- Blocked By SPID

  The server process ID of the process blocking this process (if this process has been blocked). Blocking may occur from applications running on the same server as this process' application or from applications on different client computers or even from the same application. In addition, an application may block itself (in this case, this ID may be the same as the SPID). Common blocking scenarios include:

  - Queries having long execution times.
  - Canceling queries that were not committed or rolled back.
  - Applications that are not processing all results to completion.
  - Distributed client/server deadlocks.

- Wait Time

  Time in milliseconds that this process has been waiting. This duration is relative to the time the waiting process was initially discovered and collected (and therefore may be less than the total time this process has been waiting). If this value is 0, the process is not waiting.

- CPU Cycles

  Cumulative CPU time in milliseconds for this process.

- I/O Reads/Writes

  Cumulative logical disk reads and writes associated with this process.

- Memory Usage

  Number of pages in the procedure cache that are currently allocated to this process. A negative number indicates that the process is freeing memory allocated by another process.

- Number of Threads

  Number of assigned SQL Server threads associated with this process.

Microsoft SQL: Database Backup Report

This report contains information about the execution history of scheduled jobs by SQL Server Agent and shows successful database backup jobs. This report is important to understand specific information for the various types of backups occurring on a given database environment.

Only successful backups are shown in this report.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- MS Backup Restore Info / MS Backup Set
- MS Backup Restore Info / MS Backup Mediaset

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters:

- Databases:
  - The name of a specific database on which to run the report.
  - The default value is `Master`.

Primary report information:

- Database Name

  Name of the database for which the backup occurred.
- Backup Set ID

  SQL Server backup set identification number that uniquely identifies the backup.
- Database Creation Date

  Date and time the database was originally created.
- Backup Time Span

  Date and time the backup operation started.
- Backup Finish Date

  Date and time the backup operation finished.
- Backup Time

  The time taken for backup to complete (displayed in Hours:Minutes:Seconds). This value may display as zero for very short running backup operations.
- Size (MB)

  The size of the resulting backup file in megabytes.

- Backup Type

  The following backup types are available in this report:

  - Database: Complete backup for a given database.
  - Database Differential (also known as delta or differential backup): Backs up only those items that have changed since the last backup operation.
  - Log: Every SQL Server database has a transaction log that records all transactions and the database modifications made by each transaction. This record of transactions and their modifications supports three operations: recovery of individual transactions, recovery of all incomplete transactions when SQL Server is started, and rolling a restored database forward to the point of failure.
  - File or File Group: Backs up the database on an individual file or group of files. This approach provides the flexibility to improve the efficiency of backups.

- Device Type

  Backup device type from among the following:

  - Disk: Temporary, Permanent.
  - Tape: Temporary, Permanent.
  - Pipe: Temporary, Permanent.

- Physical Device Name

  The physical name of the backup device. For example:
  `C:\MSSQL\BACKUP\RD1_77200615430.bak`

Microsoft SQL: CPU Load Report

A Microsoft SQL Server instance can be run on servers having one or more CPUs. SQL Server has its own scheduler, independent of the operating system, to schedule its own process threads. This report indicates how a particular SQL Server instance is handling its process threads on each CPU of a server. There is one graph per CPU, each showing its SQL Server thread information (as described in Primary Report Information). The CPU Load Report gives important about a SQL Server instance's load balance on its server.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- MS Reports Collection / MS CPU Load

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

- Graph Heading

  Each graph refers to a single CPU in the system. The CPUs are labeled as Processor 0 - Processor N. This heading is "Processor null" for servers with only a single processor.

- Graph Legend

  Each graph has a legend that uses color to differentiate items displayed on the graph. If an item has a zero value, it does not display on the graph.

- Total Tasks

  The total number of SQL Server threads for that particular processor.

- Active Worker Threads

  The number of active SQL Server threads for that particular processor.

- Idle Worker Threads

  The number of idle SQL Server threads for that particular processor.

- Total Worker Threads

  The total number of SQL Server threads possible in the SQL Server thread pool.

- Work Queued

  The number of SQL Server threads having queued work.

- Scheduler Queue Length

  The number of tasks queued to run on the processor. This value is a relative rating of how busy the processor is. A high number indicates that the processor is overloaded.

Microsoft SQL: DB Configuration Check

This report provides numerous database health check indicators. Each indicator is marked as GREEN (indicator is within acceptable limits), YELLOW (indicator may need to be addressed), or RED (indicator needs immediate attention). Each row of this report represents a separate key indicator with a description of the indicator and additional comments pertaining to the indicator's usage, the indicator's parameters, and so on.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- MS Health Check / MS Server Configure
- MS Health Check / MS Index Defrag Info.
- MS Health Check / MS DB Option Info
- MS Health Check / MS DB Auto Grow Info.

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

The key indicators are divided logically by Configuration Indicators, Fragmentation Indicators, DB Option Indicators, and Auto Growth Indicators.

Regarding Fragmentation, validation is done for Index fragmentation for the data and indexes for all tables. This validation determines whether the table is heavily fragmented. Table fragmentation occurs through the process of data modifications (INSERT, UPDATE, and DELETE statements) made against the table. Because these modifications are not usually distributed equally among the rows of the table, the fullness of each page can vary over time. For queries that scan part or all of a table, such table fragmentation can cause additional page reads, which hinders parallel scanning of data.

- Configuration Indicator: Allow updates

  Allow Updates indicates whether direct updates can be made to system tables. If this value is 0, then green; otherwise yellow.

- Configuration Indicator: Cross DB Ownership Chaining

  Cross DB Ownership Chaining indicates if the database owner is referenced across multiple databases. If this value is 0, then green; otherwise yellow.

- Configuration Indicator: Default Language

  Default Language specifies whether the default language option is used to specify the default language for all newly created logins. If this value is 0, then green; otherwise yellow.

- Configuration Indicator: Max Text Repl Size

  The Max Text Repl Size indicates the maximum size (in bytes) of text and image data that can be added to a replicated column in a single INSERT, UPDATE, WRITETEXT, or UPDATETEXT statement. If the value equals 65,536, then green; otherwise yellow.

- Configuration Indicator: Nested Triggers

  Triggers are nested when a trigger performs an action that initiates another trigger (which can initiate another trigger, and so on). Triggers can be nested up to 32 levels. The Nested Triggers server configuration option controls whether triggers can be nested. If the value equals 1, green; otherwise yellow.

- Configuration Indicator: Remote Access

  The Remote Access option controls logins from remote servers running instances of Microsoft SQL Server. Remote access is required for remote stored procedures. Set remote access to 1 (default) to allow logins from remote servers. Set the option to 0 to secure a local server and prevent access from a remote server. If the value is 1, green; otherwise yellow.

- Configuration Indicator: Remote Login Timeout(s)

  The remote login timeout option specifies the number of seconds to wait before returning from a failed remote login attempt. For example, if you are attempting to log in to a remote server and that server is down, remote login timeout specifies your wait time until a failed login is returned. If the value is 20, green; otherwise yellow.

- Configuration Indicator: Remote Proc Trans

  The Remote Proc Trans option protects the actions of a server-to-server procedure through a Microsoft Distributed Transaction Coordinator (MS DTC) transaction. Set remote proc trans to 1 to provide an MS DTC-coordinated distributed transaction that protects the ACID properties of transactions. Sessions begun after setting this option to 1 inherit the configuration setting as their default. If the value is 0, green; otherwise yellow.

- Configuration Indicator: Remote Query Timeout(s)

  The remote query timeout option specifies the number of seconds that must elapse when processing a remote operation before Microsoft SQL Server assumes the command failed or took too much time to perform (times out). The default is 600, which allows a ten minute wait. A value of 600 is green; otherwise, yellow.

- Configuration Indicator: Show Advanced Options

  The Show Advanced Options option displays the sp_configure system stored procedure advanced options. When you set show advanced options to 1, you can list the advanced options by using sp_configure. The default is 0. A value of 1 is green; otherwise yellow.

- Configuration Indicator: User Options

  The User Options specifies global defaults for all users. A value of 0 is green; otherwise yellow.

Database options:

This report audits database options and recommends best practices. Table 4 lists the validated settings.

Table 4 Database Options

Option	Description	Database Settings per Threshold		
		Green	Yellow	Red
Auto create statistics	When “ON,” any missing statistics needed by a query for optimization are built automatically during optimization.	= ON	< > ON	N/A
Auto update statistics	When “ON,” any out-of-date statistics needed by a query for optimization is automatically built during optimization.	= ON	< > ON	N/A
Autoclose	When “ON,” the database is shutdown cleanly and its resources are freed after the last user logs off.	= OFF	< > OFF	N/A
Autoshrink	When “ON,” the database files are candidates for automatic periodic shrinking.	= ON	< > ON	N/A
Dbo use only	When “ON,” only the database owner can use the database.	= OFF	< > OFF	N/A
Offline	When “ON,” the database is offline.	= OFF	< > OFF	N/A
Read only	When “ON,” users can only read data in the database, not modify it. The database cannot be in use when a new value for the read only option is specified.	= OFF	< > OFF	N/A
Recursive triggers	When “ON,” enables recursive firing of triggers. When false, prevents direct recursion only.	= OFF	< > OFF	N/A
Single user	When “ON,” only one user at a time can access the database.	= OFF	< > OFF	N/A
Torn page detection	When “ON,” incomplete pages can be detected.	= ON	< > ON	N/A
Trunc. log on chkpt	When “ON,” a checkpoint truncates the inactive part of the log when the database is in log truncate mode. This is the only option you can set for the master database.	= ON	< > ON	N/A

Database auto grow info:

This report audits the Auto Grow Statistics and provides an appropriate suggestions to review the database file growth.

Table 5 Database Auto Grow Statistics

Option	**Database Settings per Threshold**		
	Green	**Yellow**	**Red**
Auto grow statistics	= Auto Grow	< > Auto Grow	N/A

Microsoft SQL: Job Failure Report

This report summarizes all SQL Server scheduled job executions. It provides a high-level view across servers, as well as the ability to drill down into specific jobs and job-steps.

This report is useful in determining which servers require attention. It also allows the DBA to quickly identify not only which jobs are failing, but also the frequency of failure for each job-step.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- MS SQL Agent Info/MS Job History
- MS SQL Agent Info/MS Jobs

Required report criteria:

- Organization
- Time span

Parameters:

- Start Time:
  - The inclusive start time period of the report.
- End Time:
  - The inclusive end time period of the report.

If the same date is provided for both start and end, that 24 hour period is reported on. 48 hours for a start and that differ by 1 day, and so forth.

Primary report information:

- Summary Table
  - **Server Name**
  - **Instance Name**
  - **Successful**

    Number of successful jobs completed on this server/instance during the specified time interval.
  - **Failures**

    Number of failed job executions on this server/instance during the specified time interval.
  - **Failure Rate(%)**

    Failures / (Successes + Failures)
  - **Failure Time**

    Cumulative time for all failed jobs during the specified time interval.

- Detail Table (Drilldown: Single Instance):

  (Three types of rows in detail table: job, step, error messages)

  - **Job Name -or- Step Names**

    Name of the specific job or step

  - **Success**

    # of times this jobs/steps completed successfully

  - **Failure**

    # of times this job/step failed

  - **Percent Failure**

  - **Average Success Runtime**

    Average time this job/step takes to complete successfully.

  - **Average Failure Runtime**

    Average time spent on this job/step if it fails.

  - **Total Runtime (Failures)**

    Cumulative time for this job for all failed executions.

- Error Messages (Drilldown: Single Step)

  - **Count**

    Number of times this error message occurred.

  - **Error Messages**

    Full text of the error message.

Oracle Reports

HP DMA provides the following Oracle reports:

- Oracle: Health Check Report on page 103

  Shows a database's overall health as measured by numerous KPIs (Key Performance Indicators).

- Oracle: Top SQL Report on page 107

  Shows the most expensive or worst performing SQL statements running against the database over a period of time.

- Oracle: Latch Contention Report on page 109

  Shows the number and type of latches over a period of time.

- Oracle: System-Wide Statistics Report on page 111

  Shows overall system statistics summary over a period of time.

Oracle: Health Check Report

This report shows the key indicators of database health. Each indicator is marked as GREEN (indicator is within acceptable limits), YELLOW (indicator may need to be addressed), or RED (indicator needs immediate attention). Each row of this report represents a separate key indicator with a description of the indicator and additional insightful comments about the usage of the indicator, how that indicator was computed, etc. The report's overall summary is calculated by summing all green indicators with a value of 1, yellow with a value of ½, and red with a value of 0. The sum of all green and yellow are divided by the total number of indicators (green, yellow, and red) to arrive at the value of "Summary of Database Health" at the top of the report.

The data in this report are from the last collected data before midnight of the date specified.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- Oracle DDL4 / Oracle V$BackupDatafile
- Oracle DDL5 / Oracle V$LogFile
- Oracle DDL5 / Oracle V$DataFile
- Oracle DDL5 / Oracle V$ControlFile
- Oracle DDL5 / Oracle DBADataFiles
- Oracle DDL5 / Oracle V$Parameter
- Oracle Monitors / Oracle V$SessionWait
- Oracle Monitors / Oracle V$Session
- Oracle Monitors / Oracle V$OpenCursor
- Oracle Monitors / Oracle V$Lock
- Oracle Monitors / Oracle V$Latch
- Oracle Monitors / Oracle V$Process
- Oracle Statistics / Oracle V$WaitStat
- Oracle Statistics / Oracle V$RowCache
- Oracle Statistics / Oracle V$RollStat
- Oracle Statistics / Oracle V$SysStat
- Oracle Statistics / Oracle V$RollStat
- Oracle Statistics / Oracle V$License
- Oracle Statistics / Oracle V$LibraryCache
- Oracle Statistics / Oracle V$BufferPoolStatistics
- Oracle Statistics / Oracle V$BufferPool
- Oracle Statistics / Oracle V$Backup

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

- Database Block Buffer Effectiveness

  Measure of the Database Buffer Cache Hit Ratio from:

  ```
  v$stasstat (1-(PhysicalReads/(DBBlockGets+ConsistentGets)))*100
  ```

  If the ratio is less than 90%, then red; if the ratio is 90-95%, then yellow; otherwise green.

- Datafiles backup mode status

  Ensure all “AVAILABLE” datafiles have not been left in backup mode as reported by `v$backup`.

- Shared Pool Effectiveness

  Measure of the Shared Pool Hit Ratio from:

  ```
  v$library_cache (SUM(pins)/(SUM(reloads)+SUM(pins))
  ```

  If the ratio is less than 90%, then red; if the ratio is 90-95%, then yellow; otherwise, green.

- Redo Log Buffer Effectiveness

  Measure the ratio of redo log sync times vs. user commits from `v$sysstat`. If the Ratio is less than .5, then green; between .5 and .75, yellow; above .75, red. This is generally an indicator that the redo log buffer is too large.

- Rollback Segment Contention

  Ratio of waits to gets from `v$rollstat` as a percent. If any rollback segment has a ratio greater than 6, it is red; between 3 and 6, yellow; and less than 3, green. red or yellow is generally an indication that an additional rollback segment is needed.

- Buffer Pool Partition Check

  Ensure more than one buffer pool is in use as reported by `v$buffer_pool`. If only one is configured, then yellow. More than one pool is green.

- Buffer Pool Usage

  Ensure more than one buffer pool is in use as reported by `v$buffer_pool_statistics`. If any buffer pools have no usage, then yellow. If all buffer pool used, then green.

- Process and Session Limits

  Check if the process or session limits as specified in `v$parameter` are being approached. If within 90%, then red; within 80%, then yellow; otherwise green.

- Blocked Locks Check

  Check for any locks on objects that are blocking another process for more than 200 seconds as reported by `v$lock` and `v$locked_object`.

- Open Cursor Check

  Check the open cursor limit. Compare the `open_cursor init.ora` parameter against the maximum number of cursors specified in `v$sesstat`. If within 80%, then red; within 70% then yellow; otherwise green.

- Datafile Limit Check

  Check the datafile limit. Compare the `db_files init.ora` parameter against the number of datafiles currently used by the database as specified in `v$datafile`. If within 90%, then red, within 80%, then yellow; otherwise green.

- Controlfile Multiplexing Check

  Ensure at least one controlfile and ensure the controlfiles are in different directories as reported by `v$controlfile`. Green if multiplexed; otherwise red.

- Redolog Multiplexing Check

  Ensure each redo log group has at least two members as reported by `v$logfile`. Green if multiplexed; otherwise red.

- Database Blocksize Check

  Ensure the database blocksize is at least 8KB. Green if at least 8KB; otherwise red.

- Latch Contention Check (Required HitRatio >= 99%)

  Ensure the latch contention ration is greater than or equal to 99% as reported by `v$latch` for cache buffers lru chains. Green if >= 99; otherwise red.

- Online Redo Logfile Status Check

  Ensure no online log files are either STALE or INVALID. If Stale or invalid, then red; otherwise green.

- Disk Sorts to Memory Sorts Ratio

  Disk to memory sort ratio as reported by:

  ```
  v$sysstat ("sorts (disk)", "sorts (memory)")
  ```

  If within 5%, then green; within 8%, then yellow; otherwise red.

- Rollback Segment Header Contention

  Ratio of waits to gets as reported by `v$rollstat`. If any rollback segments have less than 95% waits, then yellow; otherwise green.

- Database Backup Schedule

  Verify RMAN is used for database backups. If true then green; otherwise ignored.

- Archive Log Backup Schedule

  Verify RMAN is used for archivelog backups. If true then green; otherwise ignored.

- Multiple Users with Same DB Login

  Verifies that no more than one record in `v$session` with the same user name excluding NULL OS Users and UserNames.

- Rollback Segment Extent Contention

  Ensure the ratio of Undo Block waits to Consistent Gets as reported by `v$waitstat` and `v$sysstat` is less than 1%. If less than 1%, then green; otherwise yellow.

- Data Dictionary Cache Hit Ratio

  The ratio of gets to gets + misses as reported by `v$rowcache`. If greater than 95%, then green; between 95% and 90%, then yellow; otherwise red.

- ORA- Errors Logged to Alert Log in Past Seven Days

  The number of ORA- errors reported in the alert log in the last 7 days as fetched from the `alert.log` file. If none, green; otherwise yellow.

- Sessions Waiting for More than Five Minutes in Past Seven Days

  The number of sessions that waited more than 5 minutes as reported by `v$sessionwait`.

➤ The collection interval for this table is generally 30 minutes, so waits less than the collection interval may not be reported. The following waits are ignored: “pmon timer,” “rdbms ipc message,” “smon timer,” and “SQL*Net message from client.”

Oracle: Top SQL Report

The Top SQL report displays the most resource consumptive SQL over the specified time period. This allows you to know which SQL statements consume the most resources and likely which are the worst performing. If the queries are vendor supplied, then tuning may be possible using outlines.

The value for the “SQL Weight” column can be modified by the “Oracle.Top.SQL” parameter in both the “Oracle V$SQL” and “Oracle V$SQLText” Agent routines. By default, this parameter is set to “(disk_reads + buffer_gets) / executions” but can be any expression of columns in the V$SQL table. Changes to these parameters should only be done at the request of Senior DBAs.

Only the first 80 characters of the SQL statement are displayed in the initial report screen. Click the hyperlinked SQL to display the entire statement in the area below the initial report page.

Accurate if DB Rebooted between Start/s: No

If the database was rebooted at any time during the time period, the statistics used to determine the worst performing SQL are reset, so the report may not be entirely accurate.

Required agent modules and routines:

- Oracle Monitors / Oracle V$SQL
- Oracle Monitors / Oracle V$SQLText
- Oracle Monitors / Oracle V$Session

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters:

- Sample Size:
  - The number of SQL statements to display. If set to 10, then the top 10 SQL statements are reported. If set to 100, then the top 100 SQL statements, and so forth.
  - The default value is 10.

Primary report information:

- SQL Weight

  The value calculated by the formula specified by the “Oracle.Top.SQL” Agent property.

- Executions

  Number of executions that took place on this object since it was brought into the library cache. This is reset when the instance is started.

- Disk Reads

  The number of disk reads for this child cursor since previous instance startup.

- Buffer Gets

  The number of buffer gets for this child cursor since the previous instance startup.

- Rows Processed

  Total number of rows the parsed SQL statement returns since the previous instance startup.

- Sorts

  Number of sorts that were done for this child cursor since the previous instance startup.

- OS User

  The name of the OS User identified executing this SQL statement during the specified time period.

Many reports will display NULL for the OS User and/or DB User. It is difficult to catch a SQL statement in the process of being executed because the default Agent frequency for the “Oracle Monitors” Agent module is 30 minutes. For example, a query that takes 20 minutes to run will show in the report (since it is still in the SQL area) but the executing user may have either logged out or is on to another SQL statement that does not show in the report.

It is possible that the same user (as specified by OS User and DB User) will be associated with multiple SQL statements. This could be the result of shared logins, parallel query and/or parallel execution (original query spawning child queries concurrently) or the SQL statements being executed at different times during the specified time interval.

- DB User

  The name of the Database user identified executing this SQL statement during the specified time period.

Many reports will display NULL for the OS User and/or DB User. It is difficult to catch a SQL statement in the process of being executed because the default Agent frequency for the “Oracle Monitors” Agent module is 30 minutes. For example, a query that takes 20 minutes to run will show in the report (since it is still in the SQL area) but the executing user may have either logged out or is on to another SQL statement that does not show in the report.

It is possible that the same user (as specified by OS User and DB User) will be associated with multiple SQL statements. This could be the result of shared logins, parallel query and/or parallel execution (original query spawning child queries concurrently) or the SQL statements being executed at different times during the specified time interval.

- SQL Statement

  The first 80 characters of the SQL statement hyperlinked to display the entire text of the SQL statement (in the area below the report).

Oracle: Latch Contention Report

The Latch Contention report displays Oracle Database latch statistics by hour, minute, and second. High latch wait times or a high number of latch misses can indicate potential performance problems. Some latches indicate problems with the application code, others with database initialization parameters, etc. Contention with the “library cache” latch, indicates that bind variables need to be used instead of hardcoded parameters in the application SQL. Contention with the “redo allocation latch” may be corrected by adjusting the `log_small_entry_max_size init.ora` parameter. Contention with the “row cache” latch may be corrected by adjusting the `shared_pool_size` larger, and so forth. You should refer to this report if latch waits are long on examination of the System or Session level wait events.

This report is sorted by the total number of misses during the report's time period as reported by `v$latch.misses + v$latch.immediatemisses`. Note that if the “Time Period” column is 0, then only one collection was made during the specified time period resulting in zero values for the calculated columns (Total Miss Change / Hour, Total Miss Change / Minute, and Total Miss Change / Second).

Accurate if DB Rebooted between Start/s: No.

➤ If the instance was restarted during the period between a report’s start time and end time, the statistics are reset, resulting in incorrect data. This may also result in negative numbers.

Required agent modules and routines:

- Oracle Monitors / Oracle V$Latch

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

- Latch Name

  Name of the latch.

➤ Each type of latch is dealt with in a unique way. Due to the volume of latches, only a few are documented. See the Oracle Latch Documentation for additional information on any particular latches.

- Gets

  Number of times the latch was requested in willing-to-wait mode.

- Misses

  Number of times the latch was requested in willing-to-wait mode and the requestor had to wait.

- Immediate Gets

  Number of times a latch was requested in no-wait mode.

- Immediate Misses

  Number of times a no-wait latch request did not succeed (that is, missed).

- Waits Holding Latch

  Number of waits for the latch while the waiter was holding a different latch.

- Spin Gets

  Willing-to-wait latch requests which missed the first try but succeeded while spinning.

- Total Misses

  The sum of all misses (misses + immediate misses).

- Time Period (Days)

  The number of days being reported over.

- Average Total Miss Change/Hour

  The average of total misses per hour over the reporting interval.

- Average Total Miss Change/Minute

  The average of total misses per minute over the reporting interval.

- Average Total Miss Change/Second

  The average of total misses per second over the reporting interval.

Oracle: System-Wide Statistics Report

System-Wide Statistics displays system statistics sorted by the amount the statistic changes by time. This report provides a look into the most active statistics for this database. These most active statistics are targets for possible performance tuning.

➤ The statistics do not have consistent units, so be careful when comparing. The details of each statistic are documented by Oracle. Note that if the "Time Period" column is 0, then only one collection was made during the specified time period resulting in zero values for the calculated columns (Total Miss Change / Hour, Total Miss Change / Minute, and Total Miss Change / Second).

Accurate if DB Rebooted between Start/s: No

➤ If the instance was restarted during the period between a report's start time and end time, the statistics are reset, resulting in incorrect data. This may also result in negative numbers.

Required agent modules and routines:

- Oracle Statistics / Oracle V$SysStat

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters:

- Sample Size:
  - The number of SQL statements to display. If set to 10, then the top 10 SQL statements are reported. If set to 100, then the top 100 SQL statements, and so forth.
  - The default value is 10.

Primary report information:

- Statistic

  The name of the system statistic.
- Last Value

  The value of this statistic since the database was started.
- Time Period

  The number of days this report encompasses.
- Average Change/Hour

  The average change per hour of this statistic over the time period of the report. Since the units are not the same by statistic, take care when comparing lines.
- Average Change/Minute

  The average change per minute of this statistic over the time period of the report. Since the units are not the same by statistic, take care when comparing lines.

- Average Change/Second

  The average change per second of this statistic over the time period of the report. Since the units are not the same by statistic, take care when comparing lines.

DB2 Reports

HP DMA provides the following DB2 reports:

- DB2: Database Health Check Report on page 115

  Shows several Key Performance Indicators (KPI) grouped in Red (Alert), Yellow (Warning) and Green (Normal) zones. The report also shows the overall Database Health which is a percentage value computed using KPI values and their associated weight.

- DB2: Buffer Pool Report on page 119

  Shows various performance-related stats for buffer pools.

- DB2: Top SQL Report on page 121

  Shows the most expensive SQL statements and associated statistics running against the database.

- DB2: Lock Wait Report on page 123

  This is a line graph for “Total Lock Waits,” “Locks Currently Waiting,” and “Locks Currently Held” plotted against time (every hour).

- DB2: Tablespace Performance Report on page 124

  Shows various performance-related statistics for tablespaces.

- DB2: Table Reorg Check Report on page 125

  Shows the reorg check output for tables and indices. Reorg check determines if tables/ indices need reOrganization.

- DB2: Index ReOrganization Check Report on page 127

  Shows indices reorg check output. This report is similar to the “Table Reorg Check Report” but reports on all the indices in the database.

- DB2: Agent Analysis Report on page 128

  Shows various DB2 Agent-related stats (Agents from pool, Agents created due to empty pool, Agents stolen, Agents registered, Agents waiting on token, idle Agents, and max Agent overflows) plotted against time (every hour).

- DB2: Database Usage Report on page 130

  Shows “Applications connected,” “Applications Executing,” and “Total Agents” plotted against time (every hour).

- DB2: Lock Escalations Report on page 131

  Shows “Lock Escalations” and “Exclusive Lock Escalations” plotted against time (every hour).

- DB2: Sort Activity Report on page 133

  Shows “Active Sorts,” “Total Sorts,” and “Sort Overflows” plotted against time (every hour).

- DB2: Log Usage Report on page 134

  Shows “Log space Used” and “Total Log” plotted against time (every hour).

- DB2: Database Size Report on page 135

  Shows “Database Size” and “Database Capacity” (max size that database can grow to) plotted against time (every hour).

- DB2: Backup and Archive Log Report on page 136

  Shows the details of Database Backups and Archive Log operations.

- DB2: SQL Statement Distribution Report on page 138

  Shows the horizontal bar chart of different types of SQL statements (Static, Dynamic, Failed, Commit, Rollback, Select, Update/Insert/Delete, DDL) executed in database.

- DB2: Tablespace Usage Report on page 140

  Shows tablespace space usage and remaining free space in the tablespaces.

- DB2: Container File System Usage Report on page 142

  Shows the space used on different file systems associated with tablespace containers.

- DB2: Invalid Objects Report on page 144

  Shows the invalid objects and associated dependencies.

DB2: Database Health Check Report

Report on key indicators of database health. Each indicator is marked as GREEN (indicator is within acceptable limits), YELLOW (Indicator may need to be addressed), or RED (indicator is in need of immediate attention). Each row of this report represents a separate key indicator with a description of the indicator and additional insightful comments about the usage of the indicator, how that indicator was computed, etc. The overall summary is computed as follows: Each indicator has associated weight between 0 to 10 (items with a weight of 0 are not included in the overall summary calculation).The overall summary is calculated by summing all green indicators with a value of 1 multiplied by associated weight, yellow with a value of 0.5 multiplied by associated weight, and red with a value of 0.2 multiplied by associated weight. The sum of all green, yellow and red are divided by the sum of associated weights of all the indicators (green, yellow, and red) to arrive at the value of “Summary of Database Health” at the top of the report.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 DBM Snapshot
- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

Some of the key indicators described here may be displayed as “Undefined.” This is not an error but rather an indication that the result could not be calculated for that indicator (e.g. denominator becoming zero in the expression).

- Overall Buffer Pool Hit Ratio (Weight 10)

  Measure of the buffer pool effectiveness. It is calculated as (1-((Data & Index Physical Reads)/(Data & Index Logical Reads)))*100. If the ratio is less than 70%, then red; if the ratio is 70-80%, then yellow; otherwise green.

- Buffer Pool Data Page Hit Ratio (Weight 8)

  Measure of the buffer pool effectiveness for data pages. It is calculated as (1-((Data Physical Reads)/(Data Logical Reads)))*100. If the ratio is less than 60%, then red; if the ratio is 60-70%, then yellow; otherwise green.

- Buffer Pool Index Page Hit Ratio (Weight 8)

  Measure of the buffer pool effectiveness for index pages. It is calculated as (1-((Index Physical Reads)/(Index Logical Reads)))*100. If the ratio is less than 70%, then red; if the ratio is 70-80%, then yellow; otherwise green.

- Average Buffer Pool Page Write Time (Weight 8)

  Measure of the write speed of pages from buffer pools. It is calculated as (Data & Index write time in ms)/ (Total number of data & index pages written). If the average write time is more than 100 ms, then red; if it is 20-100 ms, then yellow; otherwise green.

- Average Buffer Pool Page Read Time (Weight 8)

  Measure of the read speed of pages into buffer pools. It is calculated as (Data & Index physical read time in ms)/ (Total number of data & index pages read). If the average read time is more than 50 ms, then red; if it is 10-50 ms, then yellow; otherwise green.

- Dirty Page Threshold Cleaner Triggers (Weight 7)

  Percentage of dirty page threshold page cleaners. It is calculated as ((Dirty page threshold cleaner triggers)/ ((Dirty page threshold cleaner triggers)+(LSN gap cleaner triggers)+(Dirty page steal cleaner triggers)))*100. If dirty page threshold cleaner trigger is more than 50%, then green; otherwise yellow.

- Lsn Gap Cleaner Triggers (Weight 7)

  Percentage of LSN gap page cleaners. It is calculated as ((LSN gap cleaner triggers)/ ((Dirty page threshold cleaner triggers)+(LSN gap cleaner triggers)+(Dirty page steal cleaner triggers)))*100. If LSN gap cleaner trigger is 10- 50%, then green; otherwise yellow.

- Dirty Page Steal Cleaner Triggers (Weight 7)

  Percentage of dirty page steal cleaners. It is calculated as ((Dirty page steal cleaner triggers)/ ((Dirty page threshold cleaner triggers)+(LSN gap cleaner triggers)+(Dirty page steal cleaner triggers)))*100. If dirty page steal cleaner trigger is less than 20%, then green; if it is 20-40%, then yellow; otherwise red.

- Maximum Total Log Space Used (Weight 8)

  Measure of transaction log space used. It is calculated as ((Max log space used)/((Log space available)+(Log space used)))*100. If max total log space used is greater than 95%, then red; if it is 80-95%, then yellow; otherwise green.

- Sorts Overflow (Weight 10)

  Measure of sort overflows. It is calculated as ((Sort overflows)/(Total sorts))*100. If sort overflow is greater than 20%, then red; if it is 5-20%, then yellow; otherwise green.

- Average Sort Time (Weight 0)

  Measure of average time spent for sort operation. It is calculated as (Total sort time)/ (Total sorts). This is just an informative indicator and it's not included in calculating the summary database health. This indicator is always shown as green.

- Piped Sorts Accepted vs Piped Sorts Requested (Weight 8)

  Measures the acceptance of piped sorts requested. It is calculated as ((Piped sorts accepted)/(Piped sorts requested))*100. If this ratio is greater than 80%, then green; if it is 60-80%, then yellow; otherwise red.

- Asynchronous Data Reads (Weight 8)

  Measure of asynchronous data reads. It is calculated as ((Asynchronous data reads)/(Data physical reads))*100. If this ratio is greater than 70%, then green; otherwise yellow.

- Asynchronous Index Reads (Weight 8)

  Measure of asynchronous index reads. It is calculated as ((Asynchronous index reads)/ (Index physical reads))*100. If this ratio is greater than 70%, then green; otherwise yellow.

- Asynchronous Data Writes (Weight 8)

  Measure of asynchronous data writes. It is calculated as ((Asynchronous data writes)/ (Total data writes))*100. If this ratio is greater than 70%, then green; otherwise yellow.

- Asynchronous Index Reads (Weight 8)

  Measure of asynchronous index writes. It is calculated as ((Asynchronous index writes)/(Total index writes))*100. If this ratio is greater than 80%, then green; if it is 60-80%, then yellow; otherwise red.

- Database Files Closed (Weight 5)

  Number of database files closed in last 24 hours because the applications reached the limit specified by MAXFILOP database configuration parameter. It is calculated as (Value of Files Closed counter)-(Value of Files Closed counter 24 hours ago). Files Closed counter is taken from database snapshot. If no files were closed in last 24 hours, then green; if 1-100 files were closed, then yellow; otherwise red.

- Average Lock Wait Time (Weight 6)

  Measure of average time spent on lock wait. It is calculated as (Total lock wait time in ms)/(Number of lock waits). If the average time is greater than 5000 ms, then red; if it is 2000-5000 ms; then yellow, otherwise green.

- Lock Timeouts (Weight 5)

  Number of lock timeouts in last 24 hours. It is calculated as (Value of Lock Timeouts counter)-(Value of Lock Timeouts counter 24 hours ago). If this number is less than 100, then green; otherwise yellow.

- Lock Waits (Weight 5)

  Number of lock waits in last 24 hours. It is calculated as (Value of Lock Waits counter)-(Value of Lock Waits counter 24 hours ago). If this number is less than 200, then green; otherwise yellow.

- Internal Rollbacks Due to Deadlock (Weight 7)

  Measure of internal rollbacks due to deadlocks. It is calculated as ((Internal rollbacks due to deadlocks)/(Total internal rollbacks))*100. If the ratio is more than 90%, then red; if it is 50-90%, then yellow; otherwise green.

- Deadlock Detected (Weight 5)

  Number of deadlocks detected in last 24 hours. It is calculated as (Value of Deadlocks counter)-(Value of Deadlocks counter 24 hours ago). If no deadlocks were detected in last 24 hours, then green; if 1-5 deadlocks detected, then yellow; otherwise red.

- Exclusive Lock Escalations (Weight 2)

  Measure of exclusive lock escalations. It is calculated as ((Exclusive lock escalations)/(Total lock escalations))*100. If the ratio is less than 80%, then green; otherwise yellow.

- Lock Escalations (Weight 5)

  Number of lock escalations in last 24 hours. It is calculated as (Value of Lock Escalations counter)-(Value of Lock Escalations counter 24 hours ago). If more than 200 locks were escalated in last 24 hours, then yellow; otherwise green.

- Rows Selected/Rows Read (Weight 7)

  Rows selected versus rows read. This indicates whether rows are read wisely in order to serve the selected rows. It is calculated as ((Rows selected)/(Rows read))*100. If the ratio is greater than 10%, then green; if it is 1-10%, then yellow; otherwise red.

- Package Cache Hit Ratio (Weight 7)

  Measure of package cache performance. It is calculated as (1- ((Package cache inserts)/(Package cache lookups)))*100. If the ratio is greater than 90%, then green; if it is 70-90%, then yellow; otherwise red.

- Package Cache Overflow Ratio (Weight 7)

  Measure of absence of package cache overflows. It is calculated as (1- ((Package cache overflows)/(Package cache lookups)))*100. If the ratio is greater than 90%, then green; if it is 60-90%, then yellow; otherwise red.

- Catalog Cache Hit Ratio (Weight 7)

  Measure of catalog cache performance. It is calculated as (1- ((Catalog cache inserts)/(Catalog cache lookups)))*100. If the ratio is greater than 90%, then green; if it is 70-90%, then yellow; otherwise red.

- Catalog Cache Overflow Ratio (Weight 7)

  Measure of absence of catalog cache overflows. It is calculated as (1- ((Catalog cache overflows)/(Catalog cache lookups)))*100. If the ratio is greater than 90%, then green; if it is 60-90%, then yellow; otherwise red.

- Hash Join Overflow Ratio (Weight 7)

  Measure of absence of hash join overflows. It is calculated as (1- ((Hash join overflows)/(Total hash joins)))*100. If the ratio is greater than 90%, then green; if it is 60-90%, then yellow; otherwise red.

- Hash Join Small Overflow vs. Hash Join Overflow (Weight 2)

  Measure of hash join small overflows. It is calculated as ((Hash join small overflows)/(Total hash join overflows))*100. If the ratio is less than 10%, then green; otherwise yellow.

- Total Prefetch Wait Time (Weight 5)

  Total time spent on prefetch wait in last 24 hours. It is calculated as (Value of Prefetch Wait Time counter)-(Value of Prefetch Wait Time counter 24 hours ago). If prefetch wait time is less than 3,600,000 ms (1 hour), then green; otherwise yellow.

- Agents Created Due to Empty Agent Pool/Assigned from Pool (Weight 8)

  Measure of Agents created due to empty Agent pool. It is calculated as ((Agents created due to empty Agent pool)/((Agents created due to empty Agent pool)+(Agents from pool)))*100. If the ratio is greater than 70%, then red; if it is 30-70%, then yellow; otherwise green.

- Post Threshold Sorts (Weight 8)

  Number of post threshold sorts in last 24 hours. It is calculated as (Value of Post threshold sorts counter)-(Value of Post Threshold sorts counter 24 hours ago). If this number is 0, then green; if it is 1-20, then yellow; otherwise red.

- Maximum Number of Agents Waiting (Weight 8)

  Maximum number of Agents waiting for token (simultaneously) since the instance startup. This value is available in database manager snapshot. If this number is less than 10, then green; otherwise yellow.

- Agents Stolen (Weight 6)

  Number of Agents stolen in last 24 hours. It is calculated as (Value of Agents Stolen counter)-(Value of Agents Stolen counter 24 hours ago). If this number is 0, then green; otherwise yellow.

DB2: Buffer Pool Report

Buffer pool plays a critical role in overall database performance. Tablespaces are associated with in-memory buffer pools for caching data in order to provide faster access. Hence, the buffer pools should be regularly monitored to verify that buffer pools are being effectively utilized to reduce physical reads from disk, thereby increasing performance. Buffer Pool Performance is a tabular report, one row per buffer pool, providing several performance related statistics for each buffer pool.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-2 / DB2 Buffer Pool Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

A value of -1 in any of these values indicates that its data is not available in the repository for the specified report period.

- Buffer Pool

  Name of the buffer pool.

- Buffer Pool Hit Ratio (BPHR)

  BPHR is an indicator of buffer pool effectiveness. It is calculated as (1-((Physical Reads)/(Logical Reads)))*100. Logical & Physical reads include both data and index pages in this formula and therefore is an indication of the percentage of total reads that were served from the buffer pool.

- Data BPHR

  Data BPHR is the BPHR computed for data pages only (excluding the index data). It is calculated as (1-((Data Physical Reads)/(Data Logical Reads)))*100.

- Index BPHR

  Index BPHR is the BPHR computed for index pages only. It is calculated as (1-((Index Physical Reads)/(Index Logical Reads)))*100.

- Avg Read Time (ms)

  Average time in milliseconds to read a page from disk.

- Avg Write Time (ms)

  Average time in milliseconds to write a page to disk.

- % Async Reads

  Percentage of asynchronous reads out of total (synchronous & asynchronous) physical reads. Asynchronous reads are performed by database manager prefetchers such that the data is read from disk, in advance, and kept in buffer pool before the user process needs that data.

- % Async Writes

  Percentage of asynchronous writes out of total (synchronous & asynchronous) writes to disk. Asynchronous writes are performed by asynchronous page cleaners such that the dirty pages (changed data) in the buffer pool are written to disk in advance to make room for new pages in the buffer pool. A prefetcher may also write dirty pages to disk to make space for the pages being prefetched.

- Avg Async Read Time (ms)

  Average time in milliseconds to read a page from disk for all the asynchronous reads. Do not get confused that the asynchronous writes are faster, they are not; they are simply performed in advance.

- Avg Async Write Time (ms)

  Average time in milliseconds to write a page to disk for all the asynchronous writes. Do not get confused that the asynchronous writes are faster, they are not; they are simply performed in advance.

- Avg Data Pages Read / Async Req

  Average number of data pages read per asynchronous read request. Reading a greater number of pages per request is more efficient.

- Unread Prefetch Pages

  Number of pages which were prefetched by prefetchers, but not required or read by any user requests. If there are too many unread prefetch pages, it indicates that prefetchers are too aggressive and are unnecessarily bringing pages into memory which are not required.

- Physical Reads

  Total number of pages read from disk. Since unread prefetch pages are cumulative and will keep on increasing as time passes, it is difficult to determine whether a particular value of unread prefetch pages are acceptable or too high. Hence unread prefetch pages should be analyzed by comparing with physical reads (which is also cumulative).

DB2: Top SQL Report

This report shows the SQL statements with the highest resource consumption and/or the worst performance. After reviewing the information in this report, DBA's or others can determine if SQL needs to be tuned to improve performance. This report can also help in diagnosing overall database performance issues. For example, in some cases, you may decide not to run a particular SQL statement on the database during certain critical time periods.

SQL statements are ranked based on associated weights which are calculated using relevant performance statistics of the SQL. The default formula to calculate the SQL weight is:

```
(Total System CPU Time + Total User CPU Time) + 0.01 * (Total Execution
Time) + 0.001 * (Number of sorts in the SQL statement)
```

This formula may be changed using the HP DMA Web Server. However, note that any such changes will only affect SQL statement information collected after the formula was changed. All SQL statement data collected prior to the formula change will have their calculations based upon the old formula.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-6 / DB2 Dynamic SQL Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters:

- Sample Size:
  - The number of SQL statements to display. If set to 10, then the top 10 SQL statements are reported. If set to 100, then the top 100 SQL statements, and so forth.
  - The default value is 10.

Primary report information:

This report is tabular with its columns described here. The Statement Text column can be selected in order to see the full SQL statement.

If a SQL statement is costly but not executed frequently, it should not be considered the Top (worst performing) SQL. Conversely, if a SQL statement has average cost but it is executed very frequently, it has more of an effect on overall performance and is therefore a candidate for SQL tuning. Based upon that fact, the statistics shown in this report are the cumulative value since the DB2 instance startup and allows the number of executions to be considered in the statistics. The same statistics are used to calculate the SQL weight and hence the SQL weight also considers the number of executions of each SQL statement.

- SQL Weight

  Weight of the SQL (described above) as calculated during data collection from Dynamic SQL Snapshot.

- Total System CPU Time (ms)

  Total system CPU time in milliseconds of all the executions of the SQL since DB2 instance startup.

- Total User CPU Time (ms)

  Total user CPU time in milliseconds of all the executions of the SQL since DB2 instance startup.

- Total Execution Time (ms)

  Total execution time in milliseconds of all the executions of the SQL since DB2 instance startup.

- Total Sorts

  Total number of sorts performed during all the executions of the SQL since instance startup.

- Rows Read

  Total number of rows read during all the executions of the SQL since DB2 instance startup.

- Rows Written

  Total number of rows written during all the executions of the SQL since DB2 instance startup.

- Executions

  Total number of executions of the SQL since DB2 instance startup.

- Statement Text

  Shows first 50 characters of the SQL statement. This column is hyperlinked and, when selected, you can see the complete text of the SQL below the tabular report.

DB2: Lock Wait Report

Lock-wait is a situation when one application is waiting for another application to release the lock on a database object. Lock-wait results in long response time (poor performance) and may also result in errors (if the lock wait is timed out). For these reasons, lock-wait is an undesirable scenario and should be avoided wherever possible. This report plots Total Lock Waits (since instance startup) over time (maximum granularity of one hour). It also shows an instantaneous value of Total Locks Waiting and Total Locks Held over time in the same graph. This helps to identify lock-wait situations and the interval of time when there are more lock waits. If the report indicates high lock waits, more detailed data can then be researched using the HP DMA Database monitoring modules (Lock-Wait Snapshot, Lock Snapshot, Application Snapshot) to identify the applications involved in lock waits. These applications can be analyzed and corrected to reduce the lock waits (by choosing proper isolation level or refining the application logic so to avoid contention).

Accurate if DB Rebooted between Start/s: No.

Lock-waits shown in this report is a cumulative value since instance start-up. Hence this value is supposed to be either constant or increasing. However, the value will be reset to 0, if the instance is rebooted between the start time and end time of this report.

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

This report displays a line graph of the following three statistics over time. These statistics are collected from Database snapshot.

- Lock Waits

  Total number of lock-wait events in the database since instance startup.

- Locks Waiting

  Instantaneous value of total number of locks waiting in the database.

- Locks Held

  Instantaneous value of total number of locks held in the database.

DB2: Tablespace Performance Report

All the tablespaces associated with the same buffer pool may perform differently. Hence it becomes important to analyze the performance statistics on the tablespace level also. Whether to have a single tablespace associated with each buffer pool or to have multiple tablespaces grouped together with the same buffer pool are some of the initial design considerations. For example, associating a tablespace with mostly read-only operations and static data and a tablespace with read-write operations and dynamic data to the same buffer pool might not result in optimal performance. The Tablespace Performance Report is a tabular report with one row per tablespace and provides performance related statistics against each tablespace (and can reveal any inefficient groupings of tablespaces and buffer pools).

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-2 / DB2 Tablespace Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

This report has similar information as the Buffer Pool Performance report but the values are per tablespace instead of per Buffer Pool (recognizing that multiple tablespaces can be associated with the same buffer pool). See DB2: Buffer Pool Report on page 119 for a detailed description of the information in this report.

DB2: Table Reorg Check Report

DB2 provides a utility to reorganize the tables and indexes to eliminate fragmentation, thereby compacting the data. ReOrganization reduces the storage used by tables & indexes and improves the performance because I/O is reduced when data is compacted. The Table ReOrganization Check report is a tabular report with one row per table. It shows statistics required to determine whether the table needs to be reorganized. Click on any table name's hyperlink to show its indexes, along with information about the need to reorganize indexes.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Reorg Check / DB2 Reorg Check Table Stats
- DB2 Reorg Check / DB2 Reorg Check Index Stats

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

Table reOrganization is determined based on F1, F2 and F3 formula values. F1, F2 and F3 formulae and their recommended value ranges are as follows:

- **F1**: 100 * OVERFLOW / CARD < 5
- **F2**: 100 * (Effective Space Utilization of Data Pages) > 70
- **F3**: 100 * (Required Pages / Total Pages) > 80

  If their values fall outside the recommended range, the table needs reOrganization. Generally, if two out of three of the values indicate reOrganization, that table should be reorganized.

- Table Schema

  Schema name of the table.

- Table Name

  Name of the Table.

- Reorg Check

  This is a 3-character field, each character mapping to one of the three formulas: F1, F2 and F3 respectively. The character displayed is either a dash "-" or an asterisk "*". A dash means that the formula value is in the recommended range, and an asterisk means that the formula value is out of the recommended range, indicating a possible need for reOrganization.

- Cardinality

  Number of rows in the table.

- Overflow

  Number of rows that overflowed.

- Total Pages

  Total number of pages used by the table.

- Free Pages

  Number of free pages in the table.

- Active Blocks

  Number of active blocks for a multi-dimensional clustered (MDC) table.

- Table Size

  Size of the table in bytes.

Additional report detail:

Index reOrganization is determined based on F4, F5, F6, F7 and F8 formula values. F4, F5, F6, F7, and F8 formulae and their recommended value ranges are as follows:

- **F4**: CLUSTERRATIO or normalized CLUSTERFACTOR > 80
- **F5**: 100 * (KEYS * (ISIZE + 9) + (CARD - KEYS) * 5) / ((NLEAF - NUM EMPTY LEAFS) * INDEXPAGESIZE) > 50
- **F6**: (100 - PCTFREE) * ((INDEXPAGESIZE - 96) / (ISIZE + 12)) ** (NLEVELS - 2) * (INDEXPAGESIZE - 96) / (KEYS * (ISIZE + 9) + (CARD - KEYS) * 5) < 100
- **F7**: 100 * (NUMRIDS DELETED / (NUMRIDS DELETED + CARD)) < 20
- **F8**: 100 * (NUM EMPTY LEAFS / NLEAF) < 20

  If the index values fall outside the recommended range, the index needs reOrganization. Normally, if three out of five values indicate reOrganization, that index should be reorganized.

➤ Some of the columns may have a -1 value, which indicates that the corresponding statistic is not available for that table/index. Statistics can be generated by using the DB2 "runstats" utility.

- Index Schema

  Schema name of the index.

- Index Name

  Name of the index.

- Reorg Check

  This is a 5-character field, each character mapping to one of the five formulas: F4, F5, F6, F7 and F8. The character is either a dash "-" or an asterisk "*". A dash means that the formula value is in the recommended range, and an asterisk means that the formula value is out of the recommended range, indicating a possible need for reOrganization.

DB2: Index ReOrganization Check Report

DB2 provides a utility to reorganize the indexes. The reOrganization operation eliminates fragmentation, reducing the storage used by indexes and improving performance (because I/O is reduced). The Index ReOrganization Check Report helps in identifying the indexes that need to be reorganized. It is a tabular report with one row per index that shows the statistics required to determine whether the table needs to be reorganized.

Index reOrganization check statistics are also available from the Table ReOrganization Check Report. However, that report shows the index statistics as a drill-down on each table, one table at a time. This report lists all the indexes in the database.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Reorg Check / DB2 Reorg Check Index Stats

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

This report has all of the columns shown in the Addition Report Detail section of the Table ReOrganization Check Report.

➤ Similar to the Table ReOrganization Check Report, some of the columns in this report can have a “-1” value which indicates that the corresponding statistic is not available for that index. Statistics can be generated by using DB2 “runstats” utility.

Also, this report has the following two additional columns: Table Schema and Table Name.

- Table Schema

  Schema name of the table for which the index has been created.

- Table Name

  Name of the Table for which the index has been created.

DB2: Agent Analysis Report

For each database that an application uses, various database processes or threads start to perform the various application tasks. These tasks include logging, communication and prefetching. Database Agents are engine dispatchable unit (EDU) processes or threads that perform these tasks. In UNIX® environments, these Agents run as processes. In Windows environments, the Agents run as threads. This report plots line graphs of several Agent related statistics (Agents from pool, Agents created due to empty pool, Agents stolen, Agents registered, Agents waiting on token, Idle Agents, Maximum Agent overflows) against time between StartDate and EndDate.

➤ This report runs on a DB2 instance level. Hence the "Database" report criteria actually shows a drop-down of DB2 instances instead of databases.

Accurate if DB Rebooted between Start/s: No

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 DBM Snapshot

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

This report plots a line graph of the following statistics against time between start and end. These statistics are collected from Database Manager (DBM) Snapshot.

- From Pool

  The number of Agents assigned from the pool since instance startup. This is a cumulative value and does not decrease unless the instance is restarted.

- Created Empty Pool

  The number of Agents created (because the Agent pool was empty) since instance startup. It includes the number of Agents started at DB2 startup (as defined in `num_initAgents` configuration parameter). This is a cumulative value and does not decrease unless the instance is restarted.

- Stolen

  The number of times since instance startup that Agents are stolen from an application. Agents are stolen when an idle Agent associated with an application is reassigned to work on a different application. This is a cumulative value and does not decrease unless the instance is restarted.

- Registered

  The number of Agents registered in the database manager instance that is being monitored.

- Waiting on Token

  The number of Agents waiting for token so thy can execute a transaction in the database manager.

- Idle

  The number of Agents in the Agent pool that are currently unassigned to an application and are therefore "idle".

- Max Overflows

  The number of times a request to create a new Agent was received when the maxAgents configuration parameter had already been reached.

DB2: Database Usage Report

This report provides a high level view of how heavily the database is being used during a specified interval. It plots the number of "Agents," "Connected Applications" and "Executing Applications" against time. By looking at this graph for 24 hour interval, one can easily identify the peak time of database usage during that 24 hour interval. A similar analysis can be performed for a week or a month interval to identify the database usage pattern for different days of the week or different days of the month.

Addition useful conclusions can also be derived by looking at this report. For example, if the number of Agents is much higher than the connected applications, it indicates that many applications have more than one Agent working for it - indicative of intra or inter partition parallelism. If the executing applications are much less than connected applications, it's indicative of many idle connections.

This report can also help in tuning the Agent related configuration parameters (MAXAGENTS, NUM_POOLAGENTS, NUM_INITAGENTS, MAX_COORDAGENTS, MAXCAGENTS) and connections and applications related configuration parameter (MAX_CONNECTIONS, MAXAPPLS, AVG_APPLS).

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

Primary Report Information

This report plots a line graph of following three statistics against time between start and end. These statistics are collected from Database snapshot.

- Agents

  Number of subAgents for all applications connected to the database at any point in time.

- Connected Apps

  Number of applications that are currently connected to the database.

- Executing Apps

  Number of applications that are currently connected to the database, and for which the database manager is currently processing a request.

DB2: Lock Escalations Report

Lock escalation is an event in which the database escalates several row level locks to a single table level lock. A lock is escalated when the total number of locks held by an application reaches the maximum amount of lock list space available to the application or the lock list space consumed by all applications is approaching the total lock list space. Exclusive lock escalation is an event when locks are escalated from several row locks to one exclusive table lock or when an exclusive lock on a row causes the table lock to become an exclusive lock. This report plots both "Lock Escalations" and "Exclusive Lock Escalations" against time and provides a quick way to look at the trend of lock escalations and exclusive lock escalations over a period of time. Both of these statistics are cumulative over time and get reset when the instance is restarted. Hence if their value is constant over a period of time, it indicates no lock escalation (or exclusive lock escalation) has occurred during that interval. Also, their values are not supposed to drop; if the value drops, the instance has been restarted around that time.

Accurate if DB Rebooted between Start/s: No

Lock Escalation Guidelines

Lock escalations reduce the concurrency of data and exclusive lock escalations are even worse for concurrency. Hence they should be minimized as much as possible. The following guidelines may help in resolving excessive lock escalation issues:

- The most important thing to consider while analyzing any locking related issue is application programs. Application programs should be revised to reduce the excessive locking by using proper isolation levels. If the number of locks is reduced, most likely it will reduce the lock escalations as well.
- If a cursor is read only, it should be declared so by using "FOR READ ONLY" clause.
- If a table is known to be harmless for table level locking, the LOCKSIZE of the table should be specified as TABLE.
- If the size of LOCKLIST database configuration parameter is too low, it should be increased to allow more memory for locking purpose.
- MAXLOCKS database configuration parameter controls the max percentage of LOCKLIST memory that a single application can use. Hence this parameter should also be tuned for optimal locking performance.

Report Information

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

This report plots a line graph of following two statistics against time start and end. These statistics are collected from Database snapshot.

- Lock Escalations

  Total number of lock escalations in the database since instance startup.

- Exclusive Lock Escalations

  Total number of exclusive lock escalations in the database since instance startup.

DB2: Sort Activity Report

Sorting means ordering the rows in a table into the order specified by one or more of its columns, optionally eliminating duplicate entries. Sorting is required when no index exists that satisfies the requested ordering or when sorting will be less expensive than an index scan.

➤ In general, it is better to avoid sorting because index scan is less expansive than sorting in most of the cases. Sorting becomes even more expensive when sorting can not be performed in sort memory area (sort heap) and it is spilled to temporary tables (on disk), which is known as sort overflow.

The statistics shown in this report provide one aspect of overall SQL performance in the database. Excessive number of sorts is a general indicator of poor SQL performance. Hence most expensive queries should be checked for sorts and necessary indexes should be created to avoid the sorts wherever possible. This report shows "Active Sorts," "Total Sorts" and "Sort Overflows" plotted against time.

➤ Reducing the number of sorts should help to reduce the sort overflows as well. Additionally SORTHEAP, SHEAPTHRES and SHEAPTHRES_SHR configuration parameters can be tuned to avoid sort overflow.

Accurate if DB Rebooted between Start/s: No

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

This report plots a line graph of the following three statistics against time between Start and End. These statistics are collected from Database Snapshot.

- Active Sorts

  The number of sorts in the database that currently have a sort heap allocated.

- Sort Overflows

  The total number of sorts since instance startup that ran out of sort heap and may have required disk space for temporary storage. This is a cumulative value and does not decrease unless the instance is restarted.

- Total Sorts

  The total number of sorts since instance startup. This is also a cumulative value and does not decrease unless the instance is restarted.

DB2: Log Usage Report

All databases have logs associated with them that keep a transactional history of database changes. If a database needs to be restored to a point beyond the last full offline backup, logs are required to roll the data forward to the point of failure. If the transaction logs become full, all subsequent transactions fail with the transaction log full error. This report is a line graph for "Total Log Used" and "Total Log" against time. It provides a trend analysis for log usage and helps in quickly identifying the intervals of high log usage as well as the gap between Log Used and Total Log. If that gap is very low, either increase the log size or identify the transactions which are using abnormal amount of log space. If some of the transactions are using huge amounts of transaction log space, those transactions should issue commit or rollback statements more frequently in order to free up the active log space.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

This report plots a line graph of the following two statistics against time between Start and End. The statistics in this report are collected from Database snapshot.

- Total Log Used

  The amount of log space actually used, in bytes.

- Total Log

  The total log space of the database, in bytes.

DB2: Database Size Report

This report shows a trend of database size and database capacity. Database capacity is the maximum size that a database can grow to and depends upon the space available on the mount points or drives used by the database. If the mount points or drives used by a database are dedicated for that database, the database capacity will most likely be constant. However, if those mount points or drives are used for other purposes, the database capacity will change. The report is a line graph of "Database Size" and "Database Capacity" plotted against time between report parameters.

➤ The primary difference between this report and the Storage Space Report is that the size and overall capacity of the DB2 Database, regardless of its physical layout into mount points or disk drives, is summarized in Database Size (Storage Space shows each individual mount point and drive sizes and capacities but not the aggregate total).

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Alerts (once a day) / DB2 Alert - Database Size Info

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

This report plots a line graph of following two statistics against time between start and end. These statistics are collected using the GET_DBSIZE_INFO procedure.

- Database Size

  The size of database in bytes.

- Database Capacity

  The capacity of database in bytes. This statistic is not available for partitioned database.

DB2: Backup and Archive Log Report

The Backup and Archive Log report provides a quick summary of all the backup operations on a database and associated information such as Backup Type, Time Spent, Backup Target, any errors encountered, etc. between two specified dates. In addition the report also shows the details of any archive log operation between the same two dates.

When the database is in archive logging mode, the oldest log file that is full that does not contain any active transactions is marked as archived by the database. Those log files can be moved to different location for archival. Such archived log files are not required for regular functioning of the database. However they are required when a roll forward recovery is needed on the database.

This report is useful as an auditing report to quickly find out if the database is backed up on a regular interval and whether the backup policy in place meets the data recovery requirements. In addition it is also useful when planning a recovery operation as it helps to quickly identify a suitable backup image for the recovery and the log files required to roll forward the database. Both the backup and archive log details presented in the report are collected from the database history file.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 DB History / DB2 DB History

Required report criteria:

- Organization
- Server
- Database
- Time span

Parameters: None

Primary report information:

- Partition#

  Partition number. In a partitioned database, each database partition needs to be backed up separately. The catalog partition has to be backed up first, and the rest of the partitions can be backed up simultaneously, serially or a combination of the two depending upon the backup infrastructure available for the database backup.

- Backup Type

  The type of backup operation (Offline/Online, Full/Incremental/Delta etc.)

  Backup in "year-month-date hour:minute:second" format. This time collected is the time zone of Agent machine which monitors the target database. If the report is being run in a different time zone, the time should be offset for the time zone difference (local vs. collected) to get the time in the time zone where report is being run.

- Duration (Mins)

  The time spent in minutes during the backup operation.

- First Active Logfile

  The first active log file when the backup started. For a roll forward recovery only, this log file and the subsequent log files are required after restoring the corresponding backup image.

- Backup Target

  The device type and location where the backup image is created.

- Error

  The SQLCODE and SQLSTATE of any error encountered during the backup operation. When there is no error, the backup operation is successful.

Additional report detail:

- Partition#

  Partition number. In a partitioned database each database partition has a separate set of log files and the log archival in one partition is independent from other partitions.

- Archive Log Type

  Type of log archival: Primary log path, Secondary log path, Failover archive path, Primary log archive method, Secondary log archive method.

  Log archival in "year-month-date hour:minute:second" format. This time collected is the time zone of Agent machine which monitors the target database. If the report is being run in a different time zone, the time should be offset for the time zone difference (local vs. collected) to get the time in the time zone where report is being run.

- Duration (Mins)

  The time spent in minutes between log archival start and end.

- First Active Logfile

  First active log file at the time of this log archival.

- Error

  The SQLCODE and SQLSTATE of any error encountered during the log archival. When there is no error, the log archival is successful.

DB2: SQL Statement Distribution Report

The SQL Statement Distribution report provides a high level overview of the workload on the database by showing a horizontal bar chart of different types of SQL statements executed against the database since instance startup. By looking at this report, one can quickly find out the most common SQL statements running on that database.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-3 / DB2 Database Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

The report shows one horizontal bar for each of following types of SQL statements:

- Commit SQL Statements

  The number of Commit SQL statements since instance startup.

- Rollback SQL Statements

  The number of Rollback SQL statements since instance startup.

- Dynamic SQL Statements

  The number of Dynamic SQL statements since instance startup.

Dynamic SQL statements are not pre-compiled. The dynamic SQL must be explicitly compiled at run time. DB2 caches dynamic SQL statements, so that the statements do not need to be compiled often by DB2, but they must be compiled at least once when the application is executed.

- Static SQL Statements

  The number of Static SQL statements since instance startup.

Static SQL statements are precompiled and are not compiled at run time.

- Failed SQL Statements

  The number of failed SQL statements since instance startup. The number of successful SQL statements can be computed as ((Static SQL Statement + Dynamic SQL statements) - Failed SQL statements).

- Select SQL Statements

  The number of Select SQL statements since instance startup.

- DDL SQL Statements

  The number of DDL (Data Definition Language) SQL Statements since instance startup.

- UID SQL Statements

  The number of UID (Update/ Insert/ Delete) SQL statements since instance startup.

DB2: Tablespace Usage Report

The storage within a database is managed by logical grouping into multiple tablespaces. As the database size grows, more and more storage needs to be allocated to the database's tablespaces or purge activities need to be initiated. The Tablespace Usage report shows the space (number of pages) used in each tablespace of a database. For a DMS (Database Managed Space) tablespace, it also shows percent usage. However, for a SMS (System Managed Space) tablespace, space is acquired from the operating system and hence percent usage is not applicable. The report is sorted by percent usage in descending order. In addition, each tablespace has one or more containers. This report provides a drill down on tablespace name to show all the containers of that tablespace and size of each of those containers.

Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-2 / DB2 Tablespace Config Snapshot
- DB2 Snapshot-2 / DB2 Container Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

- Tablespace

  Name of the tablespace. Clicking on this name will open a drill-down report that shows you the containers that are allocated to this tablespace.

- Type

  Type of tablespace (DMS or SMS).

- State

  Current state of the tablespace, coded as a decimal number. State 0 indicates normal state. The description of any other state can be found using the "db2tbst" utility.

- Contents Type

  Type of contents in the tablespace: Any, Long, Sys Temp, User Temp.

- Page Size

  Page size of the tablespace in bytes.

- Extent Size

  The extent size of a tablespace represents the number of pages of table data that will be written to a container before data will be written to the next container.

- Prefetch Size

  The maximum number of pages the prefetcher gets from the disk at a time.

- Total Pages

  Total number of pages allocated to the DMS tablespace. This column has no value for a SMS tablespace.

- Usable Pages

  The total number of pages in a tablespace minus overhead pages. This column only applies to a DMS tablespace.

- Used Pages

  The total number of pages that are currently used in a tablespace.

- Free Pages

  The total number of pages that are currently free in a tablespace. This column only applies to a DMS tablespace.

- Percent Usage

  Percentage of pages used out of total usable pages.

Additional report detail:

- Container Name

  Name of the container. Usually the fully qualified path name or file name.

- Type

  Type of container: Directory path, File, Raw device, Striped raw device, Striped file.

- Total Pages

  Total number of pages in the container.

- Usable Pages

  Total number of pages in the container minus overhead pages. For SMS tablespace, the value of usable pages is same as total pages.

- Accessible

  Indicates whether the container is currently accessible or not.

- Stripe Set

  The stripe set to which the container belongs.

DB2: Container File System Usage Report

Each tablespace in DB2 has one or more containers that define the storage of the tablespace. The containers can be a directory, a file, or a raw device. Containers of a SMS (System Managed Space) tablespace are directories and the disk space is not pre-allocated to the tablespace (they are used as needed). This means that the disk space available to a SMS tablespace is determined by the free space on the file systems where the containers of the tablespace are placed. In the case of DMS (Database Managed Space) tablespace, the space is pre-allocated to tablespace. However, the space available to increase the size of tablespace depends on free space on the file systems where the containers of the tablespace are placed. This report provides a quick way to look at the free space on all the file systems where the containers of the tablespace are placed and shows the space free on the file system (or mount point) of each container for all the tablespaces.

- Accurate if DB Rebooted between Start/s: Yes

Required agent modules and routines:

- DB2 Snapshot-2 / DB2 Container File System Snapshot
- DB2 Snapshot-2 / DB2 Tablespace Config Snapshot

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

- Tablespace

  Name of the tablespace.

- Partition#

  Database partition number.

- Container

  Fully-qualified path/file of the container.

- Page Size

  Page size of the tablespace in bytes.

- Container Size (pages)

  Size of the container in pages.

- FS Mount Point

  File system mount point.

- FS Type

  File system type (JFS, NTFS etc.)

- FS Total Size (bytes)

  Total size of file system in bytes.

- FS Free Size (bytes)

  Free space on file system in bytes.

- FS % Free

  Percent free space on file system.

DB2: Invalid Objects Report

This report lists any invalid database object (table, view, alias, nickname, function, method, procedure, trigger and package). The most common reason for an object to become invalid is that the objects on which it is dependent have either changed or have been dropped. Hence the report also provides a drill-down feature on invalid objects to list all the base objects on which this invalid object is dependent.

Required agent modules and routines:

- DB2 System Catalog-3 / DB2 Package Dependency
- DB2 System Catalog-3 / DB2 Packages
- DB2 System Catalog-4 / DB2 Routine Dependency
- DB2 System Catalog-4 / DB2 Routines
- DB2 System Catalog-5 / DB2 Table Dependency
- DB2 System Catalog-5 / DB2 Tables
- DB2 System Catalog-5 / DB2 Trigger Dependency
- DB2 System Catalog-5 / DB2 Triggers

Required report criteria:

- Organization
- Server
- Database

Parameters: None

Primary report information:

- Object Schema

  Schema name of the invalid object.
- Object Name

  Name of the invalid object.
- Object Type

  Type of invalid object: table, view, alias, nickname, Function, method, procedure, trigger or package.
- Date Created

  Date and time when the object was created.

Additional report detail:

- Base Type

  The object type of base object on which the invalid object is dependent.
- Base Schema

  The schema name of the base object on which the invalid object is dependent.
- Base Name

  The name of the base object on which the invalid object is dependent.

8 Environment

This chapter contains the following topics:

Dashboard

The environment dashboard presents a listing of all the objects HP DMA can currently monitor or use in automation. Select an Organization to view the associated servers, or create a new one. Servers, Instances, and Databases are not available within HP DMA until you add them to the Environment. Use Discovery to add new servers, instances, and databases to the Environment.

Navigating through the Organization Browser

Objects within an Organization can be found by “selecting” objects. Select an Organization to view Servers that belong to it, select a Server to view Instances, etc.

An object's editor is opened by double-clicking on an object in the Organization Browser. This opens the editor directly below the Organization browser. Here you will be able to view specific properties for that object. If you have the necessary permissions, you will also be able to modify values and add additional objects to the object hierarchy.

The hierarchy of objects that compose an Organization are as follows:

- Organization
- Server
- Instance
- Database

Creating an Organization

An Organization is a logical grouping of servers, whether that be for separating dev/stage/prod or separating logical business units. Since user security for running Workflows is defined at the Organization level, Organizations should be broken into units with user security in mind. You can define against which Organizations an HP DMA user can run Workflows.

1. Navigate to Environment > Dashboard.
2. Click **New Organization**.

   The Properties tab displays.
3. Type the Organization's Name.

➤ Organization names must be unique.

Deleting an Organization

Deleting an object deletes all objects "under" that object.

1. Navigate to Environment > Dashboard.
2. Double-click the object from the Organization browser that you want to delete. This opens the editor for that object.
3. Click and confirm delete.

Viewing Server Information

1. Navigate to Environment > Dashboard.
2. Select the Organization to which the Server belongs. This opens an editor below the Organization browser. The editor contains several tabs, which separate the server properties into logical groupings.

Creating a New Server

1. Navigate to Environment > Dashboard.
2. Select the Organization to which the Server belongs.
3. Click **New server**.
4. Name the new server.
5. To add a server, click **New Server**.
6. Click **Save**.

Adding a Discovered Server

The **Add server** button is available from within an instance.

1. Click **Add server**.
2. Select the desired server you want to add.

3. Click **Add**.
4. Click **Save**.

Deleting a Server

1. Navigate to Environment > Dashboard.
2. Select the Organization to which the Server belongs.
3. Click ⛔.

Creating a New Instance

The **New instance** button is available from within the Server editor.

1. Navigate to Environment > Dashboard.
2. Select the Organization to which the Server belongs.
3. Click **New instance**.
4. Edit the properties.
   - General
     - — Name: This is a required property.
     - — Type
     - — ASM
   - Connection
     - — User
     - — Password
     - — Host
     - — Port
   - Servers
   - Databases
5. Click **Save**.

Agents and Oracle Automatic Storage Manager

HP DMA Agents can communicate with Oracle Automatic Storage Manager (ASM). The ASM Agent Routines are:

- Oracle ASM Disk Space: Total space in disk groups and how much space you have available.
- Oracle ASM Disk to Mountpoint: Maps the ASM disk group to the mountpoint.

Related procedures are described here:

- Searching for a Collection Routine

  Use the **Filter** box to perform a search.

- Testing Instance Connection Information

  a Navigate to Environment > Dashboard.

  b Under the Properties tab, there is a Connection section where you need to specify the correct information for connecting to the instance.

> If the target Instance is Oracle ASM, be sure to check the **ASM** box under **General**.

  c In the Connections area, type the necessary information, including the Instance to which you want to connect:

  — User

  — Password

  — Host

  — Port

> Select the **ASM** box if you want the Agent to connect using ASM.

  d Click the **Test Connection** link under any of the databases listed.

> If the instance has no databases you must add one before attempting to test the connection.

Adding a Discovered Instance

See Discovery on page 149.

Deleting an Instance

1 Navigate to Environment > Dashboard.

2 Select the Organization to which the Instance belongs.

3 Click ⛔.

Viewing Database Information

1 Navigate to Environment > Dashboard.

2 Select the Organization to which the database belongs. This opens an editor below the Organization browser. The editor contains several tabs, which separate the database properties into logical groupings.

Creating a New Database

1 Navigate to Environment > Dashboard.

2 Double-click the desired Instance.

3 To add a database, click **New Database**.

Adding an Existing Database

See Discovery on page 149.

Deleting an Existing Database

1. Navigate to Environment > Dashboard.
2. Select the Organization to which the database belongs.
3. Click and confirm delete.

Smart Groups

Smart groups are dynamic groups of servers, instances, or databases defined by some criteria. As information about the object changes its membership in the groups is re-evaluated. For example, if a server has a custom field called `sshd_running` set to true, it may belong to an SSH Group of servers. When `sshd_running` for this server turns false, it is no longer in the SSH Group.

Grouping servers by dynamic criteria is useful for automation. A workflow deployment can contain a list of static servers that the workflow can run on as well as a list of smart groups. Deploying a workflow to the SSH Group allows the flow to run only on servers with ssh enabled. You don't need to update the deployment each time sshd is started and stopped. The smart group takes care of that for you.

Each smart group is assigned to a role for Roles Based Access Control. A user can only create smart groups for roles they are assigned to. The role acts as the master server list that the smart group will filter on. The role must give the user both read and deploy permission on an organization for that organization's servers to be used in the smart group.

Discovery

The discovery process is activated when an Agent is started on a target machine. Discovery finds information about the server, network, and database instances on the machine and presents that information here. Adding discovered objects from this view into the HP DMA managed environment allows them to be monitored and automated.

The discovery tool saves you time and effort by finding and setting up the databases and instances for all Agents that are pointed at the Expert Engine.

Initially Adding Discovery Objects to Organizations

1. Navigate to Environment > Discovery.
2. Select an Organization to which you want to add servers.
3. The Discovery results display showing a list of all found servers, instances, and databases. Selecting a server populates the list of instances discovered on that server; selecting an instance does the same for databases discovered on that instance.

4 Click ⊕ to add an object to the Organization which you specified. When you add a single machine/database at the database level, discovery automatically creates and adds not only the specified database but also the necessary server and instance. When you add objects, Discovery creates the Agent configuration for the server.

> With Oracle, running Discovery returns only instances and never databases.

5 To view discovered objects within an environment, navigate to Environment > Dashboard. Double-click an object to view the objects associated by Discovery. You can return to Discovery by clicking the Discovery link on the Dashboard, or by clicking Discovery in the Menu bar.

Agents

Agents are installed on target servers for automation and monitoring. When automation Workflows are run, the Agent runs the individual Workflow steps, and returns results to the Expert Engine. Agents are also used to monitor the server and its instances and databases. The data collected by these Agents can result in the creation of alerts, or by automation rules to run Workflows. Each Agent needs a configuration to know what to collect and when to run the collections.

Searching for an Agent

You can perform a real-time filter on any Agent. Type what you are searching for in the filter field and see the filter results display as you type.

Editing Agent Properties

There are four main sections in the Agent configuration editor:

- Properties
- OS
- Database
- Pauses

To edit Agent properties, follow these steps:

1 Navigate to Environment > Agent.

2 Click the desired Host.

The Properties tab displays. From here, you can access any of the Host's properties tabs.

Properties tab

Agents can encounter several common problems which prevent proper data collection and automation execution. These properties allow you to specify to what level an alert is set when an error of the types listed here occur. HP DMA detects these errors, and sends alerts when these problems are encountered. The severity of the alert send can be configured here. See Escalation on page 69 for more information.

1. In the Alert levels area, you can edit the following properties:
   - Module overlap: Level to which a module overlap alert is set.
   - Module overlap count: Sets the number of overlaps that can happen before an alert is sent.
   - OS error: Level to which an OS error alert is set.
   - Connect error: Level to which an Connect error alert is set.
   - SQL error: Level to which an SQL error alert is set.
2. Select or clear “Enable auto update.” Enabling auto update allows Agents to be automatically updated with new versions by the Expert Engine.
3. Click **Save**.

OS tab

1. On the OS tab, you can edit the following properties:
   - Active: Select or clear this check box to activate or deactivate the desired Module. Selecting the **Active** check box at the header level activates all modules in the list.
   - Module: A list of modules that can be scheduled to run on the server.
   - Schedule: Select default schedule or create a custom schedule.
2. Click **Save**.

Database tab

You can set values to module parameters at an instance level or at a database level.

The value assigned at instance level is assigned across all databases, unless an override value is assigned at a database level.

On the Database tab, you can edit the following properties:

- Active: Select or clear this check box to activate or deactivate the desired Module. Selecting the **Active** check box at the header level activates all modules in the list.

You must select the **Active** check box to enable the **Edit** link.

- Module: Click the desired Module to view its attributes.
- Schedule: Select default schedule or create a custom schedule.
- Params: You can set module parameters values at an instance or database level.

Pauses tab

On the Pauses tab, you can edit the following properties:

- Global pauses
  - Schedule
  - Minutes
  - Description
- Instance pauses
  - Instance
  - Schedule
  - Minutes
  - Description

To add a New Pause, click **New Pause**.

Modules

Modules are groups of data collection routines that are scheduled together and run by the Agent. Each routine in the module runs in order, returning statistics to the Expert Engine. Modules have a default schedule that you can customize using the Agent editor.

In this area, you can access the HP DMA Agent collection modules. The default modules cannot be changed, but can be copied, and new modules can be created. When you create a module, you can set its default schedule, and select which routines are part of the module, as well as the order in which the modules are run.

Viewing/Opening a Collection Module

From the Environment > Modules screen, you can view all existing modules as well as some of the modules' more important properties. You can also filter modules by type (database platform, OS, etc.) or name.

It is possible to copy read-only modules, and customize your copy. All modules shipped with HP DMA are Read Only. See Copying a Module on page 153 for more information.

Creating a New Collection Module

1. On the Environment > Modules screen, click **New module**.

   The General tab opens.

2. Type the following in the General information area:

   - Name
   - Schedule: Select default schedule or create a custom schedule for the new module.
   - Module Type

3. Add one or more routines.

> Modules must have unique names. If you attempt to save a module by a previously-used name, “Module name already used.” displays in a red bar at the top of the Modules screen. Choose a different name for your module.

Copying a Module

Copy is available from all the tabs in the Environment > Modules area. Creating a copy of a module saves time by allowing you to reuse information. Since modules are read-only, you must copy the module before you can modify it.

1. Click **Copy**.

   The General tab displays and the module name changes to “Copy of <module name>.”

2. Make any changes to the copy.
3. Click **Save**.

Scheduling Modules for Agents

Modules are only effective when scheduled by an Agent. See Agents on page 150 for more information on how to add/modify/remove module schedules.

Viewing Agents that Use a Module

1. Navigate to Environment > Modules.
2. Click the module you want to view.
3. Click the **Agent** tab.

Routines

Routines are the smallest unit of data collection in the HP DMA system. Routines are grouped into modules and scheduled by the Agent to run on target servers. The collected data is available to rules, automation, monitoring, alerting, and reporting systems.

Viewing a Collection Routine

From the Environment > Routines screen, you can view all existing routines. You can also view routine-specific information about a single collection routine. This information includes the routine's code, statistics generated by the routine, any properties available to the routine, and all of the modules in which the routine is included.

➤ All routines are Read Only.

Viewing Modules Associated with a Specific Routine

1. Navigate to Environment > Routines.
2. Select the desired routine.
3. Click the **Modules** tab.

Viewing Routine Code

1. Navigate to Environment > Routines.
2. Select the desired routine.
3. Click the **Code** tab.

Viewing Routine Statistics

1. Navigate to Environment > Routines.
2. Select the desired routine.
3. Click the **Statistics** tab.

Viewing Routine Properties

1. Navigate to Environment > Routines.
2. Select the desired routine.
3. Click the **Properties** tab.

Adding a Routine to a Collection Module

1. Navigate to Environment > Modules.
2. Open the editor for the desired Module.
3. Find the routine you want to add in the Available Routine box on the lower-left area of the screen.
4. Click the **Add** link.

Solutions tab

The Solutions tab tells you in which Solution Packs the current routine is being used. For more information on Solutions, see Chapter 9, Solutions, on page 159 of this guide.

Removing a Routine from a Collection Module

1. Navigate to Environment > Modules.
2. Open the editor for the desired Module.
3. Find the routine you want to remove in the Selected Routine box on the lower-right area of the screen.
4. Click the **Remove** link.

Custom Fields

Custom fields are primarily used to customize Workflows and rules. Custom fields can be used in Workflow steps to automatically apply values that are specific to an Organization, server, instance, database. As an example, you can have a custom field that identifies a database as "Production" or "Test," then use this field in Workflows to choose between different behavior for the different type of database.

Searching for Custom Fields

You can perform a real-time filter on any custom field. Type what you are searching for in the **Custom Fields** field and see the filter results display as you type.

Viewing Custom Fields

From the Environment > Custom Fields screen, you can view all existing custom fields associated with an Organization, a Server, an Instance, or a Database. You can also see and define the Custom Fields, if any, that are associated with a specific object. See Defining Custom Fields on page 156.

Creating Custom Fields

1. Navigate to Environment > Custom Fields.
2. Click **New field**.
3. Type the following information on the Attributes tab:
   - Name
   - Object
   - Type
   - Options: (Only available if the Custom field is a "Type = List.") Use the ⊕ to add options or use the ⊖ to delete options.
4. Navigate to the Usage tab to see which Workflows or deployments use a particular custom field.
5. To define the custom field you just created, see Defining Custom Fields on page 156.

Editing Custom Fields

1. Navigate to Environment > Custom Fields.
2. Click the custom field you want to edit.
3. Change the following information on the Attributes tab:
   - Name
   - Object: Cannot be changed for existing custom fields as changing this state could result in broken automation.
   - Type: Cannot be changed for existing custom fields as changing this state could result in broken automation.
   - Options: (Only available if the Custom field is a "Type = List.") Use the ⊕ to add options or use the ⊖ to delete options. Options can be reordered with drag-and-drop.
4. Navigate to the Usage tab to see which Workflows or deployments use a particular custom field.

Defining Custom Fields

1. Navigate to Environment > Dashboard.
2. Double-click on the object for which you want to view the Custom Fields.

   The Properties tab displays below the main Dashboard.
3. Click the Custom Fields tab to define the custom field you created in Creating Custom Fields on page 156.

Removing Custom Fields

1. Navigate to Environment > Custom Fields.
2. Open the editor for the custom field you want to delete.
3. Click and confirm delete.

Viewing Policies

1. Navigate to Environment > Dashboard.
2. Double-click on the object for which you want to view the Policies.

   The Properties tab displays below the main Dashboard.
3. Click the **Policies** tab.

Viewing Roles

1. Navigate to Environment > Dashboard.
2. Double-click on the object for which you want to view the Custom Fields.

   The Properties tab displays below the main Dashboard.
3. Click the **Roles** tab.

   From the Roles tab, you can view the permissions for the selected object. If you are logged in as an administrator, you may click on the Role to assign the permissions for each user. See Assigning/Removing Users to Roles on page 172.

9 Solutions

This chapter contains the following topics:

- Understanding Solution Packs on page 159
- Working With Solution Packs on page 161

Understanding Solution Packs

What is a Solution Pack?

A Solution Pack is a set of HP DMA components grouped together to address client-specific needs. Each Solution Pack is a purchased component that runs on the HP DMA platform. Solution Packs are deployed as a patch file, which means that no downtime is required, and they can be deployed in five to ten minutes. Each Solution Pack contains the following:

- Workflow Templates for commonly-recurring IT administration pain points
- Workflow Steps to provide an automation library
- Policies that define desired automation behavior
- Agent Modules to monitor configuration changes, events, and statistics
- Rules to detect critical events and execute Workflow Templates
- Reports for a comprehensive view of the processes and systems
- Documentation to define best practices followed by Workflow Templates

For information about available Solution Packs, contact your HP sales representative.

Where can I view available Solution Packs?

Available Screen

The Solutions > Available screen shows the available Solution Packs that are not yet installed. It also shows any available updates to installed Solution Packs.

Installed Screen

You can view all purchased/installed Solution Packs from the Solutions > Installed screen.

History Screen

You can view a recursive history of Solution Pack activity on the Solutions > History screen, shown in Figure 7.

Figure 7 Solutions > History screen

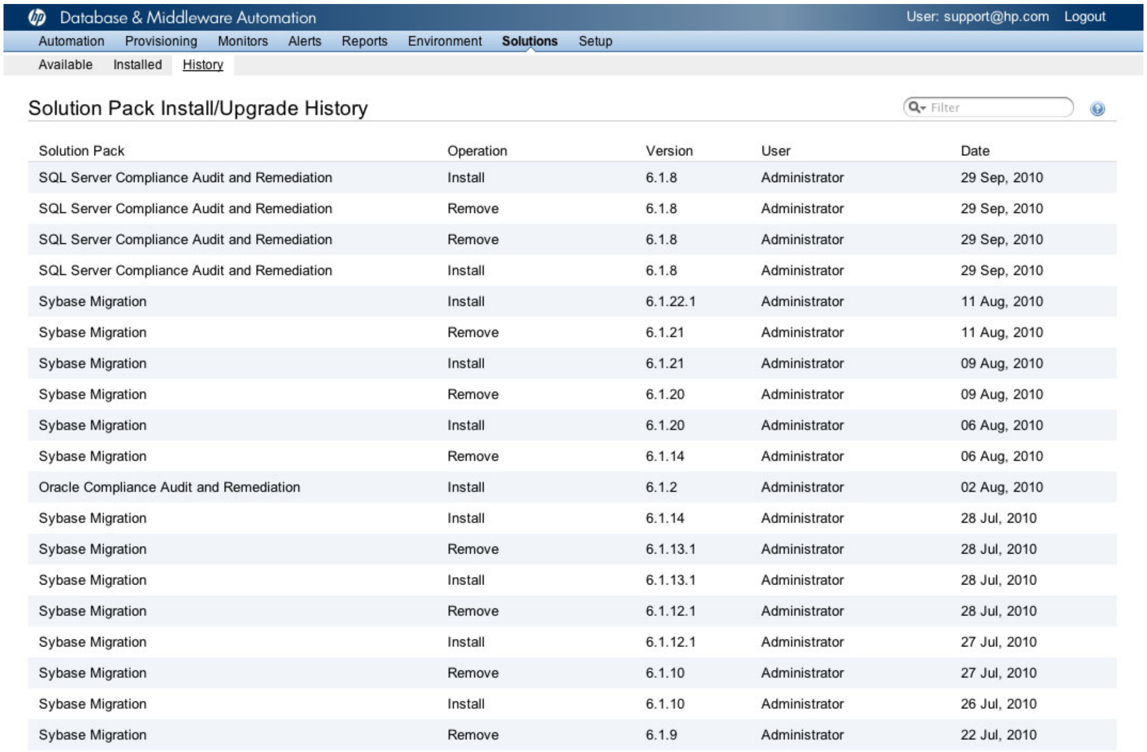

Solution Pack Install/Upgrade History

Solution Pack	Operation	Version	User	Date
SQL Server Compliance Audit and Remediation	Install	6.1.8	Administrator	29 Sep, 2010
SQL Server Compliance Audit and Remediation	Remove	6.1.8	Administrator	29 Sep, 2010
SQL Server Compliance Audit and Remediation	Remove	6.1.8	Administrator	29 Sep, 2010
SQL Server Compliance Audit and Remediation	Install	6.1.8	Administrator	29 Sep, 2010
Sybase Migration	Install	6.1.22.1	Administrator	11 Aug, 2010
Sybase Migration	Remove	6.1.21	Administrator	11 Aug, 2010
Sybase Migration	Install	6.1.21	Administrator	09 Aug, 2010
Sybase Migration	Remove	6.1.20	Administrator	09 Aug, 2010
Sybase Migration	Install	6.1.20	Administrator	06 Aug, 2010
Sybase Migration	Remove	6.1.14	Administrator	06 Aug, 2010
Oracle Compliance Audit and Remediation	Install	6.1.2	Administrator	02 Aug, 2010
Sybase Migration	Install	6.1.14	Administrator	28 Jul, 2010
Sybase Migration	Remove	6.1.13.1	Administrator	28 Jul, 2010
Sybase Migration	Install	6.1.13.1	Administrator	28 Jul, 2010
Sybase Migration	Remove	6.1.12.1	Administrator	28 Jul, 2010
Sybase Migration	Install	6.1.12.1	Administrator	27 Jul, 2010
Sybase Migration	Remove	6.1.10	Administrator	27 Jul, 2010
Sybase Migration	Install	6.1.10	Administrator	26 Jul, 2010
Sybase Migration	Remove	6.1.9	Administrator	22 Jul, 2010

Searching for a Solution

You can perform a real-time filter on any Solution Pack name or version. Type what you are searching for in the Solution Packs field and see the filter results display as you type.

Viewing a Solution

The HP DMA Solutions portal allows you to view conveniently from one area all the components that compose a Solution Pack. From the Available Solutions screen, you can view all existing Solutions that are available to you. From the Installed Solutions screen, you can view all the Solution Packs that your company already owns, as well as very detailed information about each component contained and used within each Solution Pack.

1 In the Solutions pane, point to the Solutions Name. As you point to the Solutions, you can view specific information about the Solution in the Details pane.

   — Name

   — Version

   — Released

   — Description

   — Usage

2. Click the Solution you want to view.

   Click through the following tabs to view the selected Solution's components, associations, and details:

   - General tab: Provides Solution Pack description, active version, and associated usage notes.
   - Policies tab: Displays the policies associated with the Solution. See Policies on page 47.
   - Workflows tab: Displays the Workflows associated with the Solution. See Workflows on page 21.
   - Rules tab: Displays the Rules associated with the Solution. See Rules on page 36.
   - Modules tab: Displays the Modules associated with the Solution. See Modules on page 152.
   - Steps tab: Displays the Steps associated with the Solution. See Steps on page 26.
   - Reports tab: Displays the Reports associated with the Solution. See Introduction on page 75.

Working With Solution Packs

Once you find a Solution Pack that fits your needs, you can import the Solution Pack into your environment.

Installing/Upgrading a Solution

If you need to purchase a new Solution Pack or if a more recent version of your Solution Pack exists, you will want to import the new or updated Solution Pack.

1. On the Solutions > Available tab, click **Browse** to find the Solution you want to import.

   The File Upload screen opens.

2. Locate the Solution you want to import and click **Open**.
3. Click **Import solution pack**.
4. View the installed solution:

   - Navigate to Solutions > Installed tab to view the installed Solution.
   - Navigate to Solutions > History tab to view a recursive history of Solution Pack activity.

Versioning and Importing Solution Packs

You may not import a Solution Pack with a lower version than your currently existing Solution Pack. To return to a previous Solution Pack, you must use the Rollback feature. See Rolling Back a Solution on page 162.

Also, if you import two Solution Packs with shared components the shared component is only imported once, and the higher-versioned component takes precedence over the lower-versioned component. For example, if you import Solution Pack 1 with step version 1 and Solution Pack 2 with step version 2, and they share the step, the shared step is only imported once and the higher-versioned step takes precedence and is shared between the two Solution Packs.

Steps and routines are the only components that can be shared across Solution Packs. This fact is of particular importance when you are removing Solution Packs. See Removing a Solution on page 163.

Modifying a Solution

You may need to modify an installed Solution Pack to fit your company's needs. Solution Packs are fully-supported by HP, but modifications to Solution Pack contents are supported by the customer to whom the modifications belong.

1. Navigate to the Solutions > Installed screen.
2. Select the Solution Pack you wish to modify.
3. Select the tab of the component you wish to modify.
4. Click on the desired component you wish to modify.
5. Click **Copy** to copy the component.
6. Modify the component.
7. Click **Save**.

Rolling Back a Solution

You can roll back a Solution Pack to its previous state after an import or an upgrade.

Roll back a Solution Pack import if you discover that you accidentally overwrote a necessary version of a Solution Pack or if you encounter any issues with a newly-imported Solution Pack. The most recently-installed Solution Pack is removed when you perform a rollback.

Some Solution Pack components can be modified and used without copying the entire Solution Pack:

- Rules can be modified to associate a Workflow and certain deployments.
- Polices can be modified to change values of existing attributes, and to assign targets.

For example, if you import version 1, then you import version 2, and then perform a rollback, all Solution Pack components are reset to version 1, regardless of any modifications made. You can only have one version of a specific Solution Pack on your system at once, and if you want to modify an installed Solution Pack, you have to copy it and give it an original name. See Modifying a Solution on page 162.

Another scenario is if you roll back a Solution Pack that has only been imported once, the end result is the same as if you deleted a Solution Pack. For example, if you initially import version 3, and then performed a rollback, HP DMA removes version 3 because there is not another previously-existing version to which you can roll back.

If an upgrade was performed on a Solution Pack after another Solution Pack was deleted, the rollback ignores the removed Solution Pack in the rollback sequence. Similarly, if the last action was to delete a Solution Pack, the rollback ignores the removed Solution Pack in the rollback sequence.

To roll back a solution, follow these steps:

1 Navigate to Solutions > History.

2 Click **Rollback**.

Removing a Solution

You may remove a previously-installed Solution Pack from the Installed Solution Pack list.

If you roll back a Solution Pack whose version is the only version installed on your system, the History list will display a "Remove" as the Operation.

➤ Remember that steps and routines are the only components that can be shared across Solution Packs. If a step is shared with another Solution Pack that you are removing, once you remove the Solution Pack, that shared step remains in the system.

1 Navigate to Solutions > Installed.

2 Select the Solution Pack you want to remove.

3 Click and confirm delete.

➤ Deleting a Solution Pack or selecting Rollback both display as a Remove operation on the History screen.

10 HP DMA Administration (Setup)

From the Setup screen, you can perform the following tasks:

- Add users to access the HP DMA Web Server.
- Configure Expert Engine parameters and change settings that alter the behavior of certain parts of the product. For example, the Alert Severity Level to use for various problems encountered by the product, email configuration, or setup of LDAP authentication.
- Import custom reports.

This chapter contains the following topics:

Users

A user can be an administrator or a regular user, and can be set to active or inactive. Active users are allowed login access to the Web Server. Inactive users cannot log in. Administrator users have full control over every product feature and are not restricted by security permission settings.

Viewing Existing Users

From the Setup > Users screen, you can view all existing users as well as some of the user's properties

1. To view a user, hover over a user's name.
2. Click the user you want to view.

   The Account tab displays, which includes all information about that particular user's account.

Searching for a User

You can perform a real-time filter on any user. Type what you are searching for in the **Users** field and see the filter results display as you type.

Creating a New User

1 On the Setup > Users screen, click **New user**.

2 Type the following Contact information:

- Name: Required field.
- Email: Required field.
- Phone
- Mobile
- Pager
- Pager 2

3 Type the following Login account information:

- Username: Required field.
- Password: Required field.
- Confirm password: Required field.
- Administrator: Select this check box to grant a user administrative privileges. If you select the Administrator box, the Roles tab disappears.
- Active: Default setting is "active." Clearing the **Active** check box prevents the user from logging into the Web Server.

Contact points are used by Alerts > Escalation. If you want a user to receive notification of Alerts, they will first need to have one or more contact points enabled. Next, add these contact points into the appropriate escalation. See Defining Escalation Paths on page 69.

4 Specify the following Contact points by selecting or clearing the appropriate check boxes:

- Email
- Pager
- Pager 2

5 Assign a user to at least one role. See Assigning Users to Roles on page 166.

Assigning Users to Roles

1 Create a new user. See Creating a New User on page 166.

2 Click the **Roles** tab.

3 Assign users to a role by clicking **Add** next to each role.

4 Click **Save**.

Modifying an Existing User

1. Navigate to Setup > User.
2. Click the user you want to modify.
3. Make any desired changes to the user.
    - Update user information on the Account tab.
    - Assign a user to a role on the Roles tab.
4. Click **Save**.

Logging on to the HP DMA Web Server

1. Open a web browser.
2. Type in the URL for accessing the Web Server.

    For example: `http://server_name:8080`
3. Log in to the Web Server using your user name and password.

Revoking Access to an Organization

You must be an HP DMA administrator in order to revoke a user's access to an Organization.

To revoke access to an Organization:

1. Navigate to Setup > Users.
2. Click the user you want to edit.
3. Click the **Access** tab.
4. Locate the desired Organization and clear the following options:
    - Read
    - Write
    - Execute
5. Click **Save**.

Roles

Roles define access (read-write) permissions for objects such as Organizations, Workflows, Steps, Policies, Rules, and Deployments. Deployments have an extra permission: execute. Users are assigned to roles and gain access to objects according to the permissions defined for their roles.

Roles can be defined in one of two ways: native or LDAP groups. Native roles define groups of HP DMA users in the repository. LDAP groups are retrieved from the LDAP server configured on the Setup > Expert Engine screen. No user information is stored in the repository for LDAP groups. This allows you to use your corporate directory for defining users and their permissions making security audits easier.

Understanding Roles

Here is a breakdown of how roles present themselves throughout HP DMA to administrative and to non-administrative users.

Workflows, Steps, Rules, Policies, and Deployments

- Roles tab

  Show all roles

  - Administrators:
    - All roles are links that navigate user to its details.
    - All check boxes for all roles are enabled.
  - Non Administrators:
    - Only read and write check boxes for roles in which the user is a member of will be enabled.
    - See Assigning Deployments to Roles on page 53 for details on assigning execute permissions.

Workflows

- Index
  - List all readable workflows.
  - Preview panel: All steps display for the highlighted workflow.
- New/View/Update
  - Disable all inputs when not writable.
  - Button bar

    Only display "Run" link if workflow has at least one executable deployment.
  - Deployments tab

    Show all deployments using this workflow.

    - Readable deployments are links that navigate user to its details.
    - Unreadable deployments just show the deployment name (no link).
- Saving

  Only display "Would you like to run the workflow now" message if saved workflow has at least one executable deployment.

Steps

- Index

  List all steps

  Preview panel: All workflows display for the highlighted step.
- View/Update
  - Disable all inputs when not writable.
  - Workflow tab

    Show all workflows using this step.
    - Readable workflows are links that navigate user to its details.
    - Unreadable workflows just show the workflow name (no link).

Rules

- Index
  - List all readable rules
  - Preview panel: All deployments display for the highlighted rule.
- New/View/Update
  - Disable all inputs when not writable.
  - Workflow tab

    Only list workflows that have at least one executable deployment.
    - Hide “view workflow” action when the selected workflow isn't readable.
    - On an update, include and retain already selected workflow in list (no matter what user permissions are).
  - Display all deployments for the selected workflow.
    - Disable active check box when deployment isn't executable.
    - Remove link when the deployment isn't readable.

Policies

- Index

  List readable policies
- New/View/Update
  - Disable all inputs when not writable.
  - Usage Tab

    User can remove any selected target (no matter what user permissions are). Removed targets that are not readable by the user, will be added to the Available box but only until the policy has been saved.
    - Available Usage: List all unselected, readable targets.
    - Selected Usage: List all selected targets (no matter what user permissions are).

  - Deployments Tab

    Show all deployments using this policy.

    - Readable deployments are links that navigate user to its details.
    - Unreadable deployments just show the deployment name (no link).

Deployments

- Index
  - Workflows list:
    - Display all workflows with readable deployments.
    - Preview panel: All readable deployments display for the highlighted workflow.
  - Targets list:
    - Display all targets with readable deployments.
    - Preview panel: All readable deployments (and associated workflow) display for the highlighted target.
- New/View/Update
  - Disable all inputs when not writable.
  - Button bar

    Only display "RUN" link if deployment is executable.

  - Attributes tab

    Workflow Field:

    - On create, only list readable workflows.
    - On update, always display selected workflow (no matter what user permissions are).
    - Remove 'view workflow' action if workflow isn't readable.

    Available Targets: List all unselected, deployable targets.

    Selected Targets: List all selected targets (no matter what user permissions are).

    - User can remove any selected target (no matter what user permissions are).
    - Removed targets that are not deployable by the user, will be added to the Available box but only until the deployment has been saved.

  - Parameters tab
    - Only list readable policies in "." drop-down.
    - Do NOT allow users to save parameter text values that match unreadable policies (throw error on save).
  - Rules tab

    Show all rules using this deploymen.t

    - Readable rules are links that navigate user to its details.
    - Unreadable rules just show the rule name (no link).

  - Roles tab
    - If ALL selected targets are deployable:

      Execute check boxes for roles in which the user is a member of will be enabled.
    - If ANY selected target is NOT deployable:

      ALL execute check boxes will be disabled.
- Save

  Only display "Would you like to run the workflow now" message if saved Deployment is executable.

Run

- Left column: List workflows having at least one executable deployment.
- Middle column: List executable deployments for selected workflow.
- Right column: Display all targets for selected deployment (no matter what user permissions are).

Console

Only enable the cancel workflow button for Administrators.

Monitors, Alerts, and Reports

- Monitors
  - Only display readable organizations in left column.
- Alerts
  - Only display tickets associated to readable organizations.
  - Only display readable organizations/entities in the resource drop-down.
- Reports
  - Only display readable organizations in the organization drop-down.

Viewing Existing Roles

From the Setup > Role screen, you can view all existing roles as well as some of the role's properties

1. To view a role, hover over a role's name.
2. Click the role you want to view.

Searching for a Role

You can perform a real-time filter on any role. Type what you are searching for in the **Roles** field and see the filter results display as you type.

Creating a New Role

New installations will have a Default role created with no members and no privileges.

1. Navigate to Setup > Roles.
2. Click **New role**.

   The Role tab displays.
3. Type a role name.
4. Click **Save**.

Assigning/Removing Users to Roles

You must be an administrator to modify a role.

1. Navigate to Setup > Roles.
2. Select the role to which you want to assign or remove users.

   The Role tab displays.
3. Perform one of the following:
   - To assign a user to a role, in the Available area, click **Add**.
   - To remove a user from a role, in the Selected area, click **Remove**.
4. Click **Save**.

Modifying a Role

You must be an administrator to modify a role.

1. Navigate to Setup > Roles.
2. Select the role that you want to modify.

   The Role tab displays.
3. Click the tab that corresponds to the component you wish to modify:
   - Deployments
   - Workflows
   - Steps
   - Policies
   - Rules
   - Organizations

Use the **Read All**, **Write All**, and **Execute All** links, as available, to select or clear permissions simultaneously for the selected component.

4. Click **Save**.

Filtering a Role

1. Navigate to Setup > Roles.
2. Select the desired Role.
3. Click the **Deployments** tab.
4. Perform one of the following tasks:
   - Read All
   - Write All
   - Execute All
5. Click **Save**.

Deleting a Role

Only administrators can delete roles. You cannot delete a role if it is the only role for a user. You may not delete default roles, such as those created after a fresh installation.

1. Navigate to Setup > Roles.
2. Select the role you wish to delete.
3. Click **Delete**.

Expert Engine

Expert Engine Configuration information is located on two tabs: General tab and Mail tab.

- General tab
  - — Agent properties
  - — Packet cache settings
  - — Prune settings
  - — LDAP settings
- Mail tab
  - — Outgoing Mail settings
  - — Incoming Mail settings
  - — Alert Forwarding

Expert Engine Configuration

In order to view or modify Expert Engine Configuration properties, you must be an Administrator.

Agent Settings

- Update Directory: Where you store the zip files on the Expert Engine. It is from this location that your Agent checks for updates.
- Alert Level: Alerting when an Agent unexpectedly goes offline. If the Expert Engine does not hear from an Agent over a certain time span, the Expert Engine sends an alert at this specified level.
- Minutes: The number of minutes an Agent has to communicate with the Expert Engine before a “Agent down alert is sent.”

Agents communicate with the Expert Engine every 30 seconds.

Packet Cache Settings

Agents gather data and send it to the Expert Engine. The Expert Engine temporarily stores these data packets in a directory. An alert is triggered when the number of packets in the Temporary directory exceeds the designated number of packets.

These settings control the packet cache overflow alert. When you receive this alert, it usually means that the Repository database is down or its tablespaces are full. This alert warns you of problems before they happen. If the overflow limit is reached, an alert is triggered.

- Cache directory: Where to store the data files.
- Overflow limit: The maximum number of packets before an alert is triggered.
- Overflow alert level: If the overflow limit is reached, an alert of the specified severity level is triggered.

Prune Settings

Data Prune information specifies how you want to handle excess or old information in your Repository. The Repository stores collected data, metrics, and Workflows. As you monitor your databases, your Repository becomes populated with excessive amounts of historical data. It becomes necessary for you specify parameters around how much historical data you want to maintain.

- Days: How many days to wait before old data is deleted. For example, Days=7, data older than 7 days is deleted.
- Stats Days: Stats are used to run reports. For example, if Stats Days=500, anything older than 500 days is deleted.
- Alert level: If the prune process fails, HP DMA sends out an alert at this level.

LDAP Settings

HP DMA has two ways to authenticate users: native and LDAP. The native option is enabled by default and stores user names and passwords in the HP DMA repository. Enabling the LDAP option disables the native user storage and all authentication requests are routed to your LDAP server instead. This allows you to maintain user names and passwords in your corporate directory server. HP DMA supports Microsoft Active Directory.

- Enabled: If you select “Enabled,” then users will log on to an LDAP server. Native authentication is disabled in this mode.
- SSL: Connections to the LDAP server will use SSL encryption.
- User DN: The LDAP management user that has query privileges for users and groups. This user logs into LDAP for each user authentication request.
- Password: The management user’s password.
- Base DN: Tells LDAP where to locate someone’s User ID, group, or user domain. User searches start at this node in the LDAP database.
- Server: The server where the LDAP directory resides.
- Port: The port the LDAP server is listening on. Typically 389.
- User Attribute: Where the user name is stored in LDAP. Typically uid.

Outgoing Mail Settings

- Server: SMTP Server that sends outgoing alert e-mails.
- Sender: The “From” address, which is customizable to avoid possible issues with spam blockers.

Forwarding Settings

- Enabled: Select this check box if you want an alert email to be forwarded.
- Forward to: To whom you would like to send an alert email.

Incoming Mail Settings

- SSL: Select if IMAP server uses encryption.
- Server: Host name or IMAP IP Address
- Port: Default IMAP port is 143. If you altered the default IMAP port, ensure that you type the correct IMAP port in this field.
- User: This is the email address to which you are sending alerts.
- Password: Choose your own password to log in to the IMAP server. This password is only for logging into the IMAP server.

Reports

Standard, built-in reports can be downloaded as an XML file and modified to suit a particular environment. The modified report file can then be uploaded as a custom report which is immediately available in the Reports area of HP DMA.

Searching for a Report

You can perform a real-time filter on any Report. Type what you are searching for in the **Filter** field and see the filter results display as you type.

Viewing/Opening a Report

From the Setup > Reports screen, you can view and open all existing reports.

1. Navigate to Setup > Reports.
2. Click the report you want to view.
3. When you click on a step, the file begins downloading to your machine. Open the file on your machine with the text editor of your choice.

Importing Reports

From the Setup > Reports screen, you can import reports.

1. Navigate to Setup > Reports.
2. Click **Browse** to select a report to import.
3. Click **Import report**.

Application Types

Application types allow users to extend HP DMA functionality by defining local applications which can be discovered and have their information collected via custom routines and application-specific custom fields.

Security

Granting Access to an Organization

Security and user permissions are assigned at the Organization level. A user's permission to an Organization determines if they can read, modify, or execute (automation) for objects that belong to that Organization. You must be an HP DMA administrator in order to grant a user access to an Organization. You can associate users with one or more Organizations by granting them unique read, write, and/or execute access to each Organization.

Understanding Administrative Access

Once logged into HP DMA, users with Administrator access can:

- View the Setup menu.
- View the Setup link in the Reports module and import and export custom reports.
- View all Organizations, their associated objects, and any users in the system.
- Delete any object from the system
- Edit user profiles to include deactivating users and granting them Administrator access.

Understanding Non-Administrative Access

Once logged into HP DMA, users without Administrator access can:

- View only Organizations for which they have been given read access.
- Edit Organizations and associated objects for which they have write access.
- Run Workflows against Organizations and associated objects for which they have execute access.
- Non-Administrative users may not view the Setup menu.

➤ A user without administrative privilege can still update their own account information. The only way to access their account information is by clicking on their email address in the upper right corner of the HP DMA Web Server, just to the left of the “logout” link.

To grant access to an Organization:

1. Navigate to Setup > Users.
2. Click the user you want to edit.
3. Click the **Access** tab.
4. Locate the desired Organization and select the desired access level for the Organization:
    - Read
    - Write
    - Execute
5. Click **Save**.

Understanding Security Details

Administrative-Only Abilities

- Add Users
- Configure Expert Engine
- Upload Licence Files
- Import/Export Custom Reports

- Delete Workflows and Steps

HP DMA Navigational Security

Environment Tab

- Dashboard
  - Can view all objects to which you have at least R access.
  - Can update all objects to which you have at least R/W access.
  - Cannot view objects to which you do not have at least R access.
- Discover
  - Users that do not have R/W access to any existing Organization can still use discovery. Adding an object causes a new Organization “Default” to be created, and the user will have R/W/E permission to the newly-created Organization.
  - Can only Discover objects to an Organization to which you have at least R/W access.
- Agents
  - Are only able to view the list of Agents (no hyperlink to details) or Agents belonging to Organizations to which you do not have any access.
  - Are able to select and view the details of Agents (Hyperlink) or Agents belonging to Organizations to which you have at least R access.
  - Are able to update/save details of Agents of Organizations to which you have at least R/W access.
- Modules/Routines/Custom Fields

  Unaffected by permissions

Monitors

- Graphs/Database
  - Can only view Graphs and Monitors for Organizations to which you have at least R access.
  - W/E permissions are ignored.

Alerts

- Tickets/Escalations/Instance Log Expressions/Thresholds
  - Can only view Tickets, Escalations, Instance Log Expressions, and Thresholds for an Organization to which you have at least R access.
  - Can only update Tickets, Escalations, Instance Log Expressions, and Thresholds for an Organization to which you have at least R/W access.
  - E permissions are ignored.

Reports

- Can only execute reports for an Organization to which you have at least R access.
- W/E permissions are ignored.

Provisioning

- Machines

  You must have at least R/W/E permissions to an Organization in order to provision a machine under it. In the event that you do not have permissions to *any* Organization, the machine will automatically be provisioned under a Default Organization.

- Template/History

  Unaffected by permissions.

- Software

  Only administrators are able to upload files and add/delete folders and files. However, all users should be able to view and navigate the file structure.

Automation

- Workflows/Steps
  - R/W/E permissions are ignored.
  - Non-Administrators are unable to delete Workflows and Steps.
- Rules
  - Add/Delete/Updating Rule Tab: R/W/E permissions are ignored.
  - Workflow Tab: In order to select/deselect a Deployment, you must have at least E permissions to ALL selected targets within that Deployment.
- Functions

  Unaffected by permissions.

- Policies
  - Add, Delete, Updating, Attributes Tab: R/W/E permissions are ignored.
  - Usage Tab:
    - View Usage: You must have at least R/W access to an Organizational object in order for it to show up as Available Usage.
    - View Usage: No access is required in order to view Selected Usage.
    - Remove Usage: You will only be able to remove Organizational objects to which you have R/W access.
- Deployments
  - Add, Delete, Updating, Attributes Tab
    - View Targets: You must have R/W/E access to an Organizational object in order for it to show up as an Available Target.
    - View Targets: No access is required in order to view Selected Targets.
    - Remove Targets: You will only be able to remove Selected Targets to which you have R/W/E access.

  - Parameters Tab

    Unaffected by permissions.

- Run

  You must have at least R/E permissions to an Organizational object (server, instance or database) in order to execute a Deployment.

- Console

  You must have at least R permissions to an Organizational object (server, instance or database) in order to execute a Deployment.

- History

  You must have at least R permissions to an Organizational object (server, instance or database) in order to view its execution history.

11 Web Services Interface

All web services in HP DMA are written in a REST style. This means all services are implemented using standard HTTP features and are easily accessible from every programming language. Special care is taken with the HTTP status codes the services return to the caller as well as the HTTP methods (GET, POST, PUT, DELETE).

Each API feature is exposed as a URL representing some resource. For example, the list of running workflows is accessed via `http://server_name/api/sop/running`.

Access to the web service URLs is controlled with HTTP Basic authentication. The user name and password provided in the Authentication header is a standard HP DMA user (the same credentials you use to login via a web browser).

This chapter contains the following topics:

Running Workflows

Workflows can be started using a web service interface. Running workflows' status can also be queried using this interface.

URL http://server_name/api/sop/running

POST

Starts the execution of the workflow defined in the request body.

- Request Body

  The body of the request must be URL form encoded. This is the same format that a web browser uses to submit forms on a web page. For example:

  ```
  workflow=Backup%20Database&deployment=Production&server=dev
  &instance=db1
  ```

- Parameters

  Italicized parameters are required.

  — *workflow* – the name of the workflow to execute

  — *deployment* – the name of the deployment to execute

  — *server* – the target server to run the workflow on

  — instance – the database instance to run against

- database – the database to run against
- params – any runtime parameters the workflow requires to run

- Example usage:

  The UNIX curl command is an easy way to test the web service.

```
curl -u user1:pass1 -d workflow="Backup Database"
-d deployment=Production -d server=prod1
http://server_name/api/sop/running
```

- Status Codes
  - 201 Created – returned when the workflow has started successfully. The Location HTTP header is set in the response to the URL of the running workflow.
  - 400 Bad Request – returned when the workflow could not start because parameters are missing from the request.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

GET

Returns an Atom XML feed of the running workflows.

- Example usage:

```
curl -u user1:pass1 http://server_name/api/sop/running
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:sop='http://
www.hp.com/datapal/api/sop'>
  <id>http://server_name/api/sop/running</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:56:24.249Z</updated>
  <title>Active Workflows</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/sop/running' />
  <entry>
    <id>http://server_name/api/sop/running/workflow/1230763715789</id>
    <published>2008-12-31T23:39:13.374Z</published>
    <updated>2008-12-31T23:40:24.180Z</updated>
    <title>Backup Database</title>
    <link rel='alternate' type='application/atom+xml' href='http://
server_name/api/sop/running/workflow/1230763715789' />
    <link rel='alternate' type='text/html' href='http://server_name/
sop/workflow/view/19787' />
    <sop:target server='loki' />
    <sop:status state='Waiting' />
  </entry>
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed is sent successfully.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

URL http://server_name/api/sop/running/workflow/ID

GET

Returns an Atom XML feed of a workflow's steps. Each entry contains a link to the step's resource and it's current state.

- Example usage:

```
curl -u user1:pass1 http://server_name/api/sop/running/workflow/
1230763715789
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:sop='http://
www.hp.com/datapal/api/sop'>
  <id>http://server_name/api/sop/running/workflow/1230763715789</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:40:24.180Z</updated>
  <title>Backup Database</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/sop/running/workflow/1230763715789' />
  <link rel='alternate' type='text/html' href='http://server_name/sop/
workflow/view/19787' />
  <entry>
    <id>http://server_name/api/sop/running/workflow/1230763715789/step/
1230763715789</id>
    <published>2008-12-31T23:39:13.386Z</published>
    <updated>2008-12-31T23:39:44.197Z</updated>
    <title>A test Step</title>
    <link rel='alternate' type='application/atom+xml'
    href='http://server_name/api/sop/running/workflow/1230763715789/
step/1230763715789' />
    <link rel='alternate' type='text/html' href='http://server_name/
sop/step/view/19224' />
    <sop:status state='Finished' rc='0' />
  </entry>
  <entry>
    <id>http://server_name/api/sop/running/workflow/1230763715789/step/
1230763715788</id>
    <published>2008-12-31T23:39:44.197Z</published>
    <updated>2008-12-31T23:40:04.185Z</updated>
    <title>A test Step</title>
    <link rel='alternate' type='application/atom+xml'
    href='http://server_name/api/sop/running/workflow/1230763715789/
step/1230763715788' />
    <link rel='alternate' type='text/html' href='http://server_name/
sop/step/view/19224' />
    <sop:status state='Finished' rc='0' />
  </entry>
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed is sent successfully.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 404 Not Found - returned when the workflow has completed and is no longer available in the API.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

URL http://server_name/api/sop/running/workflow/ID/step/ID

GET

Returns an Atom XML feed of a step's output. Each entry contains the step's stdout and stderr output messages.

- Example usage:

```
curl -u user1:pass1 http://server_name/api/sop/running/workflow/ID/
step/1230763715789
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:sop='http://
www.hp.com/datapal/api/sop'>
  <id>http://server_name/api/sop/running/workflow/1230763715789/step/
1230763715789</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:39:44.197Z</updated>
  <title>A test Step</title>
  <link rel='self' type='application/atom+xml'
  href='http://server_name/api/sop/running/workflow/1230763715789/
step/1230763715789' />
  <link rel='alternate' type='text/html' href='http://server_name/sop/
step/view/19224' />
  <entry>
    <id>http://server_name/api/sop/running/workflow/1230763715789/step/
1230763715789/stdout</id>
    <published>2008-12-31T23:39:13.386Z</published>
    <updated>2008-12-31T23:39:44.042Z</updated>
    <title>Standard Output</title>
    <content>Test output on stdout</content>
  </entry>
  <entry>
    <id>http://server_name/api/sop/running/workflow/1230763715789/step/
1230763715789/stderr</id>
    <published>2008-12-31T23:39:13.386Z</published>
    <updated>2008-12-31T23:39:13.085Z</updated>
    <title>Standard Error</title>
    <content></content>
  </entry>
  <sop:status state='Finished' rc='0' />
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed is sent successfully.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 404 Not Found - returned when the workflow has completed and is no longer available in the API.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

Environment Discovery

The environment web service API can be used to query environmental information from HP DMA. The API provides access to organizations, servers, instances, and databases as well as their custom fields. This can be used to populate a third-party CMDB. New objects can be created and updated through this API as well which enables third-parties to populate HP DMA with discovered objects.

URL http://server_name/api/env/organization

GET

Returns an Atom XML feed of the organizations in HP DMA.

- Example Usage:

```
curl -u user1:pass1 http://server_name/api/env/organization
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:env='http://
www.hp.com/datapal/api/env'>
  <id>http://server_name/api/env/organization</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:40:32Z</updated>
  <title>HP DMA Organizations</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/organization' />
  <entry>
    <id>http://server_name/api/env/organization/10676</id>
    <published>2008-12-31T23:40:32Z</published>
    <updated>2008-12-31T22:39:38Z</updated>
    <title>HP Software</title>
    <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/organization/10676' />
    <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/organization/10676' />
    <link rel='alternate' type='text/html' href='http://server_name/
env' />
  </entry>
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed was returned successfully.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

POST

Creates a new organization in the HP DMA repository.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 201 Created – returned when the workflow has started successfully. The Location HTTP header is set in the response to the URL of the running workflow.
  - 400 Bad Request – returned when the workflow could not start because parameters are missing from the request.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

URL http://server_name/api/env/organization/ID

GET

Returns an Atom XML feed of the organization data and its servers.

- Example Usage:

```
curl -u user1:pass1 http://server_name/api/env/organization/10676
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:env='http://
www.hp.com/datapal/api/env'>
  <id>http://server_name/api/env/organization/10676</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:42:29Z</updated>
  <title>HP Software</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/organization/10676' />
  <env:organization name='HP_Software'>
    <env:custom-field name='test 1' value='SOME VALUE 1' />
    <env:custom-field name='test 2' value='SOME VALUE 2' />
  </env:organization>
  <entry>
    <id>http://server_name/api/env/server/10681</id>
    <published>2008-12-31T23:42:29Z</published>
    <updated>2008-12-31T22:39:40Z</updated>
    <title>server_name</title>
    <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/server/10681' />
    <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/server/10681' />
    <link rel='alternate' type='text/html' href='http://server_name/
env' />
  </entry>
</feed>
```

- Status Codes
  - — 200 OK – returned when the Atom feed was returned successfully.
  - — 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - — 405 Method Not Allowed – returned for accessing the URL with any method other than POST or GET.

PUT

Updates an organization with new information.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - — 200 OK – returned when the organization was updated successfully
  - — 400 Bad Request – returned when the organization entry is missing information.
  - — 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - — 405 Method Not Allowed – returned for accessing the URL with any method other than PUT or GET

URL http://server_name/api/env/server/ID

GET

Returns an Atom XML feed of the server data and its instances.

- Example Usage:

```
curl -u user1:pass1 http://server_name/api/env/server/10681
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:env='http://
www.hp.com/datapal/api/env'>
  <id>http://server_name/api/env/server/10681</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:45:48Z</updated>
  <title>server_name</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/server/10681' />
  <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/server/10681' />
```

```
  <link rel='parent' type='application/atom+xml' href='http://
server_name/api/env/organization/10676' />
  <env:server name='server_name' os='linux' os-version=''>
    <env:network-interface dns-name='server_name.example.com'
ip-address='192.168.6.159' />
    <env:custom-field name='Server Data' value='Test 1----' />
    <env:custom-field name='Server Data 2' value='Test 2-' />
  </env:server>
  <entry>
    <id>http://server_name/api/env/instance/22692</id>
    <published>2008-12-31T23:45:48Z</published>
    <updated>2008-12-31T23:45:36Z</updated>
    <title>db2inst1</title>
    <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/instance/22692' />
    <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/instance/22692' />
    <link rel='alternate' type='text/html' href='http://server_name/
env' />
  </entry>
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed was returned successfully
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

POST

Creates a new server in the HP DMA repository.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 201 Created – returned when the server was created successfully
  - 400 Bad Request – returned when the server entry is missing information.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

PUT

Updates a server with new information.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 200 OK – returned when the server was updated successfully
  - 400 Bad Request – returned when the server entry is missing information.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

URL http://server_name/api/env/instance/ID

GET

Returns an Atom XML feed of the instance data and its databases.

- Example Usage:

```
curl -u user1:pass1 http://server_name/api/env/instance/22692
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:env='http://
www.hp.com/datapal/api/env'>
  <id>http://server_name/api/env/instance/22692</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:47:28Z</updated>
  <title>db2inst1</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/instance/22692' />
  <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/instance/22692' />
  <link rel='parent' type='application/atom+xml' href='http://
server_name/api/env/server/10681' />
  <env:instance name='db2inst1' type='DB2' asm='false'>
    <env:connection host='server_name' port='50000' windows-domain=''
username='rdc' />
  </env:instance>
  <entry>
    <id>http://server_name/api/env/database/22708</id>
    <published>2008-12-31T23:47:28Z</published>
    <updated>2008-12-31T23:45:44Z</updated>
    <title>TEST</title>
    <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/database/22708' />
    <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/database/22708' />
    <link rel='alternate' type='text/html' href='http://server_name/
env' />
  </entry>
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed was returned successfully
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

POST

Creates a new instance in the HP DMA repository.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 201 Created – returned when the instance was created successfully
  - 400 Bad Request – returned when the instance entry is missing information.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

PUT

Updates an instance with new information.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 200 OK – returned when the instance was updated successfully
  - 400 Bad Request – returned when the instance entry is missing information.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

URL http://server_name/api/env/database/ID

GET

Returns an Atom XML feed of the database data.

- Example Usage:

```
curl -u user1:pass1 http://server_name/api/env/database/22708
<feed xmlns='http://www.w3.org/2005/Atom' xmlns:env='http://
www.hp.com/datapal/api/env'>
```

```
  <id>http://server_name/api/env/database/22708</id>
  <author><name>HP DMA</name></author>
  <updated>2008-12-31T23:48:56Z</updated>
  <title>TEST</title>
  <link rel='self' type='application/atom+xml' href='http://
server_name/api/env/database/22708' />
  <link rel='edit' type='application/atom+xml' href='http://
server_name/api/env/database/22708' />
  <link rel='parent' type='application/atom+xml' href='http://
server_name/api/env/instance/22692' />
  <env:database name='TEST'></env:database>
</feed>
```

- Status Codes
  - 200 OK – returned when the Atom feed was returned successfully.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET.

POST

Creates a new organization in the HP DMA repository.

- Request Body

  The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 201 Created – returned when the database was created successfully
  - 400 Bad Request – returned when the database entry is missing information.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET

PUT

Updates a database with new information.

- Request Body

The body of the request must be an Atom XML `<entry>` object. The entry should be formatted like the GET example above.

- Status Codes
  - 200 OK – returned when the database was updated successfully.
  - 400 Bad Request – returned when the organization entry is missing information.
  - 401 Unauthorized – returned when the user name and password credentials sent in the Authorization header are incorrect.
  - 405 Method Not Allowed – returned for accessing the URL with any method other than POST, PUT or GET.

12 Datapal Tools

HP DMA provides several built-in tools. All tools are called by passing a `-t` flag and the tool name to the datapal script. For example, `datapal -t Hostname`. The datapal script is automatically installed on both Agent and Expert Engine machines.

This chapter contains the following topics:

Running the datapal Script

For UNIX, log in as the datapal user and use:

```
datapal
```

For Windows, log in as the datapal user and use:

```
cd "C:\Program Files\Data Palette"
jython\jython.bat bin\datapal
```

CertificateGen

The CertificateGen tool is used to create a self-signed certificate. This is used by the HP DMA Web Server to allow encrypted https connections. There are two parameters that are passed to this command.

- password – The password the web server will use to manage the keystore.
- hostname – The server name the HP DMA Web Server is running on. This must match the URL users will use to access the HP DMA Web Server in their web browsers.

Upon completion of this command, a file named .keystore will be generated and placed in the HP DMA user’s home directory (/opt/datapalette).

CertificateGen example:

```
datapal -t CertificateGen my_password server_name.example.com
Saving the keystore: /opt/datapalette/.keystore
Add the following XML fragment to /opt/datapalette/web/tomcat/conf/
server.xml:
    <Connector port="443" protocol="HTTP/1.1" SSLEnabled="true"
               scheme="https" secure="true" sslProtocol="TLS"
               keystoreFile="/opt/datapalette/.keystore"
               keyAlias="dma" keystorePass="my_password"/>
```

CreateDB

The CreateDB tool is used to install or upgrade the HP DMA repository. The input to this tool is a repository zip file included with HP DMA. This zip file will be named `repository-install-oracle-version.zip` or `repository-upgrade-oracle-version.zip`. Upon running the CreateDB tool, an interactive session will be created where database specific information will be required. You can also run CreateDB without the prompts by providing this information on the command line.

CreateDB example:

```
datapal -t CreateDB
System user (system): system
System password (manager): manager
Schema user (rdr): rdr
Schema password (rdeq123): rdeq123
Application user (rd): rd
Application password (eq123): eq123
Database hostname (localhost): localhost
Database port (1521): 1521
Database name (RD1): RD1
Index datafile directory (optional, /u01/oradata/SID): /u01/oradata/RD1
Data datafile directory (optional, /u02/oradata/SID): /u02/oradata/RD1
Repository zip file location: /opt/datapalette/
repository-install-oracle-5.0.0.zip
```

Hostname

The Hostname tool displays the primary, or canonical, name and IP address that the current system reports. This combination of canonical name and IP is what is used by the Expert Engine to match Agents with their Agent Configurations and will disambiguate when a system has more than one available network interface.

Hostname example:

```
datapal -t Hostname
server_name.example.com
192.168.6.207
```

SOPExecute

The SOPExecute tool is used to execute a workflow for a specific deployment target. SOPExecute can accept the following parameters:

- -u userid: The HP DMA user running the workflow.
- -p password: The HP DMA user's password.
- -w workflow: The workflow to be executed.
- -d deployment: The deployment name to be executed.
- -S server: The server on which the workflow runs.
- -I instance: Optional. The database instance on which the workflow runs. If not specified, it is assumed that the deployment is targeted at a server level.
- -D database: Optional. The database on which the workflow runs. If not specified, it is assumed that the deployment is targeted at an instance or server level.
- -E web_server: Optional. The IP or DNS address of the HP DMA Web Server. If not specified, it will default to localhost.
- -o port: Optional. The port to communicate to the HP DMA Web Server on. If not specified, it will default to 80.
- -F file_location: Optional. Specify the location of the parameter file. This parameter file is used only for deployments that have specified runtime parameters.
- -b: Optional. Enable blocking mode. This will cause SOPExecute to wait until the workflow is finished before returning control to the console.
- -v: Optional. Enable verbose mode. If used in combination with blocking mode, will cause the standard output stream of the workflow to be printed to the console.

SOPExecute parameter file:

- For deployments that have a parameter that is specified as "runtime entered", a file name must be passed using the -F option. The format of the file is:

  ```
  step.parameter1=value1
  step.parameter2=value2
  ```

- Steps or parameters with spaces in their names must be escaped using a backslash character. For example:

  ```
  Step1.parameter1=value1
  Second\ step.second\ parameter=second value
  ```

SOPExecute example:

In this example, we have a workflow named "Backup", a deployment named "Production" which is running on a server named "thistles". There are two steps, "Step1" and "Second step" which only print out the value of their input parameters. These input parameters are both "runtime specified" and use the file defined above.

```
@datapal -t SOPExecute -u user -p password -w Backup -d Production -S
server_name -F /tmp/my_parameters -b -v@
[Starting Step1]
parameter1 is 'value1'
[Finished Step1 : RC=0]
[Starting Second step]
second parameter is 'second value'
```

```
    [Finished Second step : RC=0]
```

SOPExecute reads the values from the `/tmp/my_parameters` file and assigns these values to the input parameters of the workflow before running it.

A DB2 Default Modules

DB2 Alerts (10-15 min)

The following collection routines, except DB2 Alert - Database Down, collect the DB2 Statistics for which they are named. The collected statistics are used for threshold-based alerting. The DB2 - Database Down routine alerts when a database is down. This module should be scheduled to run every 10 to 15 minutes.

This module contains the following collection routines:

- DB2 Alert – Agents Registered
- DB2 Alert – Lock Waits
- DB2 Alert – Transaction Log Used
- DB2 Alert – Database Down

DB2 Alerts (30-60 min)

The following collection routines, except DB2 Alert – Admin Notification Log, collect the DB2 Statistics for which they are named. The collected statistics are used for threshold-based alerting. DB2 Alert – Admin Notification Log collects messages from the DB2 Admin Notification Log. Alerts can be configured for these messages using alert log expressions. This module should be scheduled to run every 30 to 60 minutes.

This module contains the following collection routines:

- DB2 Alert – DMS Tablespace Used
- DB2 Alert – Tablespace State
- DB2 Alert – Lock Escalations and Timeouts
- DB2 Alert – Catalog Cache, Package Cache and Sort Overflows
- DB2 Alert – Admin Notification Log

DB2 Alerts (6 hours)

The following collection routines check for invalid database objects. If a routine finds an invalid object, rule-based alerting generates an alert. This module should be scheduled to run every 6 hours.

This module contains the following collection routines:

- DB2 Alert – Invalid Views/Nicknames
- DB2 Alert – Invalid Routines
- DB2 Alert – Invalid Triggers
- DB2 Alert – Invalid Packages

DB2 Alerts (once a day)

This collection routine collects Database Size and Database Capacity (maximum database size). The collected stats can send an alert when database size is close to reaching database capacity. Currently, the collected statistics are used in the Database Size report. This module should be scheduled to run once a day.

This module contains the following collection routine:

- DB2 Alert – Database Size Information

DB2 Analytics (once a day)

The Storage Space Projections report uses the statistics collected by the following routines. This module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 DMS Tablespace Used/Allocated
- DB2 SMS Tablespace Used
- DB2 Container Mount Point

DB2 Snapshot-1

The following collection routines collect Agent, application, statement, and lock-related snapshots. You can use the data collected by the routines to check the details of connected applications. The collected data may help in performance or lock-related troubleshooting. This module should be scheduled to run every 15 to 20 minutes.

This module contains the following collection routines:

- DB2 Agent Snapshot
- DB2 Agent Memory Pool Snapshot
- DB2 Application Snapshot Information
- DB2 Application Snapshot
- DB2 Lock Snapshot
- DB2 Lock-Wait Snapshot
- DB2 Statement Snapshot
- DB2 Sub-Section Snapshot

DB2 Snapshot-2

This module should be scheduled to run every 15 to 20 minutes.

This module contains the following collection routines:

- DB2 Buffer Pool Snapshot
- DB2 Buffer Pool Config Snapshot
- DB2 Storage Path Snapshot
- DB2 Tablespace Snapshot
- DB2 Tablespace Config Snapshot
- DB2 Container File System Snapshot
- DB2 Container Snapshot
- DB2 Ranges Snapshot

The above routines collect buffer pool and tablespace related snapshots. The data collected by the above routines is used in following reports.

- Buffer Pool Performance
- Tablespace Performance
- Tablespace Usage
- Container File System Usage

DB2 Snapshot-3

This module should be scheduled to run every 15 to 20 minutes.

This module contains the following collection routines:

- DB2 Database Snapshot
- DB2 Database Memory Pool Snapshot
- DB2 Transaction Log Snapshot
- DB2 HADR Snapshot
- DB2 DBM Snapshot
- DB2 DB Manager Memory Pool Snapshot
- DB2 FCM Snapshot
- DB2 FCM Node Snapshot
- DB2 Switches Snapshot

The above collection routines collect Database, Transaction Log, Database Manager (Instance), FCM (Fast Communication Manager), and Monitor Switches-related snapshots. The data collected by the above routines is used in following reports:

- Database Health Check
- Database Usage
- Log Usage
- Lock Wait
- Lock Escalations
- SQL Statement Distribution
- Agent Analysis
- Sort Activity

DB2 Snapshot-4

These collection routines collect snapshots related to Quiescers and to the DB2 Utility (backup, restore, load etc.) run. If required, this module should be scheduled to run every 15 to 20 minutes.

This module contains the following collection routines:

- DB2 Quiescers Snapshot
- DB2 Utility Snapshot
- DB2 Utility Progress Snapshot

DB2 Snapshot-5

These collection routines collect the table and table reorg-related snapshots. If required, this module should be scheduled to run every 15 to 20 minutes.

This module contains the following collection routines:

- DB2 Table Snapshot
- DB2 Table Reorg Snapshot

DB2 Snapshot-6

This collection routine collects dynamic SQL statements details in the Database package cache. The data collected by above routine is used in Top SQL report. This module should be scheduled to run every 15 to 20 minutes.

This module contains the following Agent routine.

- DB2 Dynamic SQL Snapshot

DB2 Reorg Check

This module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Reorg Check Table Stats
- DB2 Reorg Check Index Stats
- DB2 Table Admin Information

The collection routines above collect table and index reorg check statistics as well as administrative information about tables. The data collected by above collection routines is used in following reports:

- Table ReOrganization Check
- Index ReOrganization Check

DB2 Config

These collection routines collect database manager configuration parameters, registry variables, and database partition information. The first two routines help track changes in database manager configuration parameters and in registry variables. The partition information collected by the DB2 Partition routine is used by the snapshot collection routines. This module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 DB Manager Config
- DB2 Registry Variable
- DB2 Partition

DB2 DB Config

This collection routine collects database configuration parameters and helps track hangs in database configuration parameters. This module should be scheduled to run once a day.

This module contains the following Agent routine:

- DB2 Database Config

DB2 DB History

This collection routine collects entries from the database recovery history file. The data collected is used in Backup and Archive Log report. This module should be scheduled to run once a day.

This module contains the following collection routine:

- DB2 DB History

DB2 Environment Information

These collection routines collect environment information about the system, the installed DB2 product, and the DB2 instance. This module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Instance Environment Information
- DB2 Product Environment Information
- DB2 System Environment Information

DB2 Database DDL

This module should be scheduled to run once a day.

This module contains the following collection routine:

- DB2 DDL

The above collection routine collects the DDL of following database objects:

- Tables
- Constraints
- Indexes
- Views
- Triggers
- Functions
- Stored Procedures
- Tablespaces

DB2 System Catalog – Objects

These collection routines collect system catalog information about different database objects. The data collected is used in the Invalid Objects report. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Tables
- DB2 Routines
- DB2 Triggers
- DB2 Packages
- DB2 Views
- DB2 Tablespaces
- DB2 Schema
- DB2 Buffer Pools
- DB2 Sequences
- DB2 Indexes
- DB2 Routine Parameters
- DB2 Statements
- DB2 Index Options

DB2 System Catalog – Authority

These collection routines collect information about the permissions and authorities of different database objects. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Sequence Authority
- DB2 Schema Authority
- DB2 DB Authority
- DB2 Column Authority
- DB2 Index Authority
- DB2 Tablespace Authority
- DB2 Table Authority
- DB2 Package Authority
- DB2 Routine Authority

DB2 System Catalog – Checks & Dependency

These collection routines collect information about checks, constraints, and dependencies. The data collected by these routines is used in the Invalid Objects report. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Table Dependency
- DB2 Trigger Dependency
- DB2 Package Dependency
- DB2 Routine Dependency
- DB2 Index Dependency
- DB2 Constraint Dependency
- DB2 Checks
- DB2 Column Checks
- DB2 Table Constraints
- DB2 References
- DB2 Table Detached Dependency

DB2 System Catalog – Columns

These collection routines collect information about columns and data types. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Columns
- DB2 Data Types
- DB2 Column Use
- DB2 Column Identity Attributes
- DB2 Key Column Use
- DB2 Index Column Use

DB2 System Catalog – Partitions

These collection routines collect information about DB2 table partitions and database partitions. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Data Partitions
- DB2 Data Partition Expression
- DB2 DB Partition Group Definition
- DB2 DB Partition Groups
- DB2 Buffer Pool DB Partitions
- DB2 Partition Maps

DB2 System Catalog – Event Monitors

These collection routines collect information about DB2 event monitors. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Events
- DB2 Event Tables
- DB2 Event Monitors

DB2 System Catalog – Security

These collection routines collect information about security policy, label based access control, and surrogate authorization IDs. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Security Label Access
- DB2 Security Label Component Elements
- DB2 Security Label Components
- DB2 Security Labels
- DB2 Security Policies
- DB2 Security Policy Component Rules
- DB2 Security Policy Exemptions
- DB2 Surrogate Authorization IDs

DB2 System Catalog – Federated Server

These collection routines collect information about federated servers and related components. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Wrappers
- DB2 Wrapper Options
- DB2 User Options
- DB2 Nicknames
- DB2 Servers
- DB2 Server Options
- DB2 Pass-Through Authority
- DB2 Column Options
- DB2 Table Options
- DB2 Routine Parameter Options
- DB2 Routine Options
- DB2 Federated Routines
- DB2 Function Mapping Options
- DB2 Function Mapping Parameter Options
- DB2 Function Mappings
- DB2 Type Mappings

DB2 System Catalog – Index Extension

These collection routines collect information about index extension. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Index Extensions
- DB2 Index Extension Parameters
- DB2 Index Extension Dependency
- DB2 Index Extension Methods
- DB2 Index Exploit Rules

DB2 System Catalog – Libraries

These collection routines collect information about libraries. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Libraries
- DB2 Library Authority
- DB2 Library Bind Files
- DB2 Library Versions

DB2 System Catalog – XML

These collection routines collect information about XML storage. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 XDB Map Graphs
- DB2 XDB Map Shred Trees
- DB2 XSR Object Authority
- DB2 XSR Object Components
- DB2 XSR Object Dependency
- DB2 XSR Object Hierarchies
- DB2 XSR Objects

DB2 System Catalog – Misc

These collection routines collect miscellaneous information from system catalogs. If required, this module should be scheduled to run once a day.

This module contains the following collection routines:

- DB2 Attributes
- DB2 Cast Functions
- DB2 Name Mappings
- DB2 Transforms
- DB2 Predicate Specifications
- DB2 Column Group Columns
- DB2 Column Groups
- DB2 Full Hierarchies
- DB2 Hierarchies

We appreciate your feedback!

If an email client is configured on this system, click Send Email

If no email client is available, copy the following information to a new message in a web mail client and send the message to sa-docs@hp.com.

Product name and version: HP Database and Middleware Automation version 1.00

Document title: *User Guide: Database and Middleware Automation*

Feedback: