EMC Unisphere for VMAX Database Storage Analyzer

Version 8.1.0

Online Help (PDF version)

EMC²

Copyright © 2014-2015 EMC Corporation. All rights reserved. Published in USA.

Published September, 2015

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (https://support.emc.com).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Chapter 1 Database Storage Analyzer 5

Welcome 6
Login to Database Storage Analyzer 7
Tour of the interface 7
Viewing the dashboard 9
Viewing monitored databases and instances 11
Adding monitored Oracle databases 12
Adding monitored MS SQL server instances 14
Editing monitored Oracle databases 15
Editing monitored MS SQL server instances 17
Removing monitored databases or instances 17
Starting statistics collection 18
Stopping statistics collection 18
Running device mapping 18
Setting configuration parameters 18
Viewing the Performance Page 19
Viewing the Analytics Page 21
Adding hints from Analytics screen 27
Adding hints from Hinting screen 28
Adding database objects to an existing hint 29
Editing hints 29
Viewing hints 30
Enabling hints 31
Disabling hints 31
Removing hints 32
Viewing hint logs 32

CHAPTER 1

Database Storage Analyzer

Topics include:

- Welcome 6
- Login to Database Storage Analyzer 7
- Tour of the interface 7
- Viewing the dashboard 9
- Viewing monitored databases and instances 11
- Adding monitored Oracle databases 12
- Adding monitored MS SQL server instances 14
- Editing monitored Oracle databases 15
- Editing monitored MS SQL server instances 17
- Removing monitored databases or instances 17
- Starting statistics collection 18
- Stopping statistics collection 18
- Running device mapping 18
- Setting configuration parameters 18
- Viewing the Performance Page 19
- Viewing the Analytics Page 21
- Adding hints from Analytics screen 27
- Adding hints from Hinting screen 28
- Adding database objects to an existing hint 29
- Editing hints 29
- Viewing hints 30
- Enabling hints 31
- Disabling hints 31
- Removing hints 32
- Viewing hint logs 32

Welcome

Unisphere for VMAX includes a Database Storage Analyzer (DSA) application that provides a database to storage performance troubleshooting solution for Oracle and SQL server databases running on the storage systems.

DSA is a feature in the Foundation Suite, Advanced Suite, and the Unisphere Suite (for single tier arrays). It supports database to storage correlation by providing a shared view on how performance issues correlate to database level activity and storage level activity. This view is accessible by a database administrator (DBA) and a storage administrator (SA). The view presents: I/O metrics such as Input/Output Operations per second (IOPS), throughput and response time from both the database and the storage system which helps to immediately identify any gap between the database I/O performance and the storage I/O performance.

For example, the screenshot below shows a variable database response time in the same time window as a more or less constant storage response time, indicating that the database level variance is not caused by the storage layer.



In addition to databases running on physical servers, DSA supports the mapping of database files located on VMware virtual disks to their storage system volumes.

Note

An Oracle diagnostic pack license is required for the monitored database.

In addition, DSA supports hinting capabilities for Oracle databases on storage systems running HYPERMAX OS 5977 or higher that allows users to prioritize mission-critical database processes in order to achieve improved response time.

DSA offers the following benefits:

- Provides a unified view across database and storage.
- Quickly identifies when a database is suffering from high I/O response times.
- Reduces troubleshooting time for database or storage performance issues—DBAs and SAs can look at a unified database and storage I/O metrics view and quickly identify performance gaps or issues on both layers.
- Identifies database bottlenecks that are not related to the storage.
- Allows better coordination between the SA and DBA.

- Reduces repetitive manual drill downs for troubleshooting.

Email any feedback on this help content to mailto:VMAXContentFeedback@emc.com.

Login to Database Storage Analyzer

You can create local users using **Unisphere for VMAX**. See the Unisphere for VMAX help topic **Create local users** for details. A DSA user does not have any access permission to **Unisphere for VMAX** automatically. If you create a DSA user in **Unisphere for VMAX** without setting access permission to a storage system, the created DSA user is not allowed to log in to **Unisphere for VMAX**. However, if the user was created with permissions to at least one storage system, the user name is the same for both DSA and Unisphere for VMAX. To perform hinting operations, you must be a DSA Admin with hinting permission.

If Unisphere was installed with enabled X.509 client authentication, the **Use Client's X.509 Certificate** checkbox, checked by default, appears on the **Login** dialog.

When the **Use Client's X.509 Certificate** checkbox is checked, the user name and password fields are not required and disabled. If authentication fails, the **Failed to authenticate with the smart card** error message is presented to the user.

When the **Use Client's X.509 Certificate** checkbox is not checked, the user name and password fields are enabled.

For successful authorization, the DSA user should be created in **Security › Authorized Users** and **Groups** rather than in **Local Users**.

The **Login** dialog box contains the following elements:

- **`User`**—Your **Unisphere for VMAX** Database Storage Analyzer (DSA) user name.
- **`Password`**—Your DSA password.
- **`Login`**—Opens the DSA console.
- **Use Client's X.509 Certificate** checkbox.

Tour of the interface

The Unisphere for VMAX Database Storage Analyzer (DSA) interface consists of the following components.

Title bar:

Exits the DSA console.

Opens the entire help system.

Clicking help in a wizard page or view opens a help topic specifically for that page or view.

DASHBOARD tab

- View a full list or filtered list of Oracle databases/SQL server instances running on the storage system and their attributes.
- View performance chart for Top 5 Database/Instance IOPS for databases filtered by database name, host and ID.
- View the response time of the databases/instances displayed in the Top 5 Database/Instance IOPS filtered by database name, host and ID.

- View a bubble chart that provides a more visual comparison between DB and Storage response time for all monitored databases/instances.
- Quickly identify if the database/instance response time is meeting its threshold.

ADMINISTRATION tab

- View and filter the list of Oracle databases/SQL server instances being monitored.
- View properties associated with each monitored database/instance.
- Add database/instance that you wish to monitor.
- Start statistics collection for the monitored database/instance.
- Stop statistics collection for the monitored database/instance.
- Remove database/instance that you no longer wish to monitor.
- Edit monitored database/instance attributes.
- Run device mapping.
- Set database/instance configuration parameters (including the set of the response time thresholds).

PERFORMANCE tab

- View performance chart for the **I/O Wait** for the following Oracle database events: **I/O Wait**, **Background**, and **Non-I/O Wait** over a specified time.
- View performance chart for the **Active Session Wait** for the following Oracle database events: **Application**, **Commit**, **Concurrency**, **Network**, **ORA CPU**, **ORA Wait CPU**, **OS CPU**, **Other**, **System I/O**, and **User I/O** over a specified time.
- View performance chart for the **I/O Wait** vs **Non-I/O Wait** for the following MS SQL server instance events: **I/O Wait**, **Background**, and **Non-I/O Wait** over a specified time.
- View the database response time, log write response time, IOPS and MB/s for each database within the instance.
- View performance chart for the **Wait Classes Over Time** for the following MS SQL server instance events: **Buffer I/O**, **CPU**, **Latch**, **Locks**, **Logging**, **Network I/O**, **Waiting for CPU** , and **Other** over a specified time.
- View performance chart for Database versus Storage response time, IOPS, throughput, and I/O size over time for Reads, Writes and Redo Writes over a specified time.
- View storage back-end performance charts for IOPS, Throughput and Tier Capacity on the following disk types: **EFD**, **FC**, and **SATA** in terms of Reads or Writes over a specified time. Density can also be viewed.

ANALYTICS tab

- View Oracle database performance chart for **Wait Events/Classes over Time** for the following operations: **Single Block Read**, **System I/O**, and **Other I/O**—filtered by a specific context (**IO Wait Classes**, **IO Wait Events**, **All Wait Classes**, or **All Wait Events**) or by specified database properties over a specified time.
- View SQL server instance performance chart for **Wait Events/Classes over Time** for the following operations: **Buffer I/O**, **CPU**, **Latch**, **Locks**, **Logging**, **Network I/O**, **Waiting for CPU** , and **Other**) or by specified insance properties over a specified time.
- View the list of database/instance objects and their properties.
- Group database objects by object name. When viewing database objects, you can group all object partitions or sub-partitions belonging to the same object in one row, or you can un-group them and have one row for each partition or sub-partition of the object.

- View the list of storage volumes and their properties.
- View the list of the database sessions (filtered by session type—hosts, program names, NT domains, NT users, Login names, original Login names, Login databases) and how much time they are waiting for I/O.
- Select one or more or all Oracle database objects and create a hint.
- Select one or more or all Oracle database objects and add them to a hint.

HINTING tab

- View and filter the list of database object hints.
- View properties associated with each hint.
- Edit hint.
- Enable hint.
- Disable hint.
- Remove hint.
- View hint logs.

Viewing the dashboard

The dashboard displays a view of all of the monitored databases.

Procedure

1. Click the **DASHBOARD** tab and then **ORACLE** to view the full list of monitored Oracle databases and their associated properties.

   Note that all I/O statistics in the dashboard are database statistics only.

   The following properties display for each monitored database:

   - **DB Name**—Name of the monitored database. Hovering over a database name instance allows you to view the database size, database version, identity of the storage system that the database is running on, and the storage groups. For storage systems running HYPERMAX OS 5977, you can view the storage group service level, for example, Diamond, Platinum, and so on.
   - **DB Host**—Name or IP address of the monitored database host.
   - **DB Read RT ms**—Average database read response time in milliseconds. If you hover over the column, you will see a tool tip displaying the DB read RT values for small reads (single block) and large reads.
   - **Storage Read RT ms**—Average storage read response time in milliseconds.
   - **DB RT status**—Percentage of occurrences where the response time was as expected, over, or poor. By default, the thresholds are taken automatically from the associated Source group (SG) thresholds. These can be changed by selecting a database from the **ADMIN** tab and clicking **Configure**.
   - **DB Redo Write RT ms**—Average Database Redo Write response time for the selected period of time based on Oracle statistics. If you hover over the column, you will see a tool tip displaying the DB redo write RT values for small and large writes.

     Note

     For Oracle12 PDB, the redo statistics can be seen at the CDB level only.

- **Storage Redo Write RT ms**—Average storage write response time for storage volumes where the database redo files are located.
- **DB IOPS**—Average Database Input/Output (I/O) per second (IOPS) is the number of reads from data file and temporary files and the number of writes to data files, temporary files, redo logs and archive files per second. If you hover over the column, you will see a tool tip displaying the DB IOPS values for small reads (single block), large reads, small writes, large writes, small redo writes and large redo writes.
- **Storage IOPS**— Average Storage IOPS is the number of reads and writes per second to the storage volumes that are associated with the database files.
- **DB MB/s**—Average data transfer rate which comprises of reads from data files and temporary files and writes to data files, temporary files and redo logs. If you hover over the column, you will see a tool tip displaying the DB MB/s values for small reads (single block), large reads, small writes, large writes, small redo writes and large redo writes.
- **Storage MB/s**— Average reads and writes transfer rate for the storage volumes that are associated with the database files.
- **DB R/W Ratio**—Average ratio between reads from data files and temporary files and writes to data files, temporary files and redo logs.
- **Collection Status**—Represents the current state of the collector regardless of the date filter. The status is one of the following:
  - OK (collection is up and running properly)
  - Failed
  - Pending
  - Installing
  - Stopped

2. Click the **DASHBOARD** tab and then **MS SQL Server** to view the full list of monitored MS SQL Server instances and their associated properties.

   Note that all I/O statistics in the dashboard are database statistics only.

   The following properties display for each monitored instance:

   - **Instance Name**—Name of the monitored instance. Hovering over an instance name allows you to view the instance size, instance version, identity of the storage system that the instance is running on, and the storage groups. For storage systems running HYPERMAX OS 5977, you can view the storage group service level, for example, Diamond, Platinum, and so on.
   - **Instance Host**—Name or IP address of the monitored instance host.
   - **Instance Read RT ms**—Average instance read response time in milliseconds.
   - **Storage Read RT ms**—Average storage read response time in milliseconds.
   - **Instance RT status**—Percentage of occurrences where the response time was as expected, over, or poor. By default, the thresholds are taken automatically from the associated Source group (SG) thresholds. These can be changed by selecting an instance from the **ADMIN** tab and clicking **Configure**.
   - **Instance Log Write RT ms**—Average instance Redo Write response time for the selected period of time based on SQL server statistics.
   - **Storage Log Write RT ms**—Average storage write response time for storage volumes where the database redo files are located.

- **Instance IOPS**—Average instance Input/Output (I/O) per second (IOPS) is the number of reads from data file and temporary files and the number of writes to data files, temporary files, and log writes per second. If you hover over the column, you will see a tool tip displaying the DB IOPS values for reads, writes, and log writes.
- **Storage IOPS**— Average Storage IOPS is the number of reads and writes per second to the storage volumes that are associated with the database files.
- **Instance MB/s**—Average data transfer rate which comprises of reads from data files and temporary files and writes to data files, temporary files and log writes. If you hover over the column, you will see a tool tip displaying the DB MB/s values for reads, writes, and log writes.
- **Storage MB/s**— Average reads and writes transfer rate for the storage volumes that are associated with the database files.
- **Instance R/W Ratio**—Average ratio between reads from data files and temporary files and writes to data files, temporary files and log writes.
- **Collection Status**—Represents the current state of the collector regardless of the date filter. The status is one of the following:
  - OK (collection is up and running properly)
  - Failed

3. Set the time range for the monitoring activity.
4. View the bubble chart for the database or instance Storage Read Response Time Correlation.

   The x-axis represents database response time in milliseconds and the y-axis represents storage response time in milliseconds. Hovering over a bubble instance allows you to view its Database Name (database on the host), Database Read Response Time, Storage Read Response Time, IOPS, and Symmetrix ID.
5. View the performance chart for Top 5 Database or Instance IOPS for the monitored databases. The x-axis represents time and the y-axis represents the IOPS in terms of bytes.
6. View the performance chart for Top 5 Database or Instance Read Response Time ms for the monitored databases. The x-axis represents time and the y-axis represents the response time (greater of maximum response time and 10ms) of the top database IOPS.
7. Optional: Create a filter for the list by typing any combination of DB or Instance Name, Host and Symmetrix ID.
8. Optional: Click **Apply** to apply the filter.
9. Optional: Click **Clear** to clear the applied filter.
10. Optional: Select a chart and click to export the chart as one of the following:
    - JPEG image
    - PNG image
    - PDF document
    - SVG Vector image

Viewing monitored databases and instances

This task allows you to view the databases and instances being monitored by DSA.

Procedure

1. Click the **ADMINISTRATION** tab to view the full list of monitored databases or instances and their associated properties

   The following properties display for each database or instance:

   - **DB Name**—Name of the monitored database or instance.
   - **DB Host**—Name or IP address of the monitored database or instance host.
   - **Symmetrix ID**—Identity of the storage system that the database or instance is running on.
   - **DB Version**—Monitored database or instance version.
   - **System status**—Status of the system.
   - **Mapping status**—Status of the device mapping process.
   - **Status Description**—Monitored database or instance status description.
   - **Oldest Date**—Oldest date where there is data in the system.
   - **Latest Available Data**—Date of latest available data.
   - **Last Mapping Date**— Date of the last time the mapping process ran successfully. You can run the mapping process manually from the **ADMINISTRATION** page. You can also define the frequency of the device mapping process in the **Set Configuration Parameters** dialog.

2. Optional: Create a filter for the list by typing any combination of Environment Name, Host and Symmetrix ID.
3. Optional: Click **Apply** to apply the filter.
4. Optional: Click **Clear** to clear the applied filter.
5. Optional: Click to export the screen contents to Excel.
6. Optional: Click **Refresh** to refresh the screen display. You will be able to see the progress of installation and device mapping operations.

Adding monitored Oracle databases

This task allows you to add Oracle databases to be monitored by DSA.

Procedure

1. Click the **ADMINISTRATION** tab to view the full list of monitored databases and their associated attributes.
2. Click **Add**.
3. Click **Oracle**, then click **Create DSA User** button or **Use Existing DSA User** and click **Next**.

   Creating a DSA user is used when a new DSA database user is added by the installation process. The sys user must be provided to DSA during the next step of the installation.

   Using an existing DSA user requires creating the user manually prior to the installation using the script provided (click **DSA Script** to access the script), and then manually entering the new user during the installation.

4. Modify any number of the following values:

- **DB User Name**—Name of the database user.
- **DB User Password**—User password for the database.
- **DB SYS Password**—SYS password required by DSA in order to create the DSA database user.
- **DB Host Name or IP**—Name or IP address of the database host.
- **Port**—Database listener port number.
- **Service**—Service name that may contain up to 16 alpha numeric characters.
- **SID**—System identifier (SID) is a unique name for an Oracle database instance on a specific host.
- **Enable Device Mapping**—Enables mapping of the database files to the devices.
  Note : In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by clicking **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then clicking **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by clicking **Edit VM settings** on the **vSphere** client.
- **Mapper Host Name or IP**—Host name or IP address where the mapping process is executed.
- **Root User**:
  - **Root Password**—Password.
  - If login to remote server as root is disabled, specify SSH login user and password.
    - **Login User**—SSH user name.
    - **Login Password**—SSH user password.
  - **Use login user home directory to store mapping files**— Select this option to use the login user's home directory to store the mapping files.
  - **Run Device Mapping immediately after installation**— Select this option to run device mapping immediately after installation. Device mapping is a process running in the monitored environment It maps all of the database files to devices.
    Note : In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by clicking **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then clicking **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by clicking **Edit VM settings** on the **vSphere** client.
- **Sudo User**:
  - **Login User**—Sudo user name.
  - **Login Password**—Sudo user password.
  - **Run sudo interactively**— Select this option to run substitute user do (sudo) interactively. This allows users to run **Use login user home directory to store mapping files** programs with the security privileges of another user (normally the superuser, or root).
  - **Use login user home directory to store mapping files**— Select this option to use the login user's home directory to store the mapping files. The temporary mapping files are stored by default in temp directories.
  - **Alternative path to sudo**—Use an alternative path to sudo if sudo is not available by default or you would like to specify a different sudo path.

- **Use sudo with root password**—If sudo is configured in a way where the root password needs to be provided before running privileged actions, provide the root password here.
- **Run Device Mapping immediately after installation**— Select this option to run device mapping immediately after installation. Device mapping is a process running in the monitored environment It maps all of the database files to devices.
  Note : In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by clicking **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then clicking **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by clicking **Edit VM settings** on the **vSphere** client.
- **Mapper command line information**—Select this option to display the mapper command line details The file path displayed correlates to the directories created when you selected **Use login user home directory to store mapping files**. You can confirm that the appropriate permissions are set for the sudo login users in the sudoers file.

5. Click **Finish**.

Adding monitored MS SQL server instances

This task allows you to add MS SQL server instances to be monitored by DSA.

Procedure

1. Click the **ADMINISTRATION** tab to view the full list of monitored databases and their associated attributes.
2. Click **Add**.
3. Click **MS SQL Server**, then click **Create DSA Login** or **Use Existing DSA Login** button and click **Next**.

   If you select the **Create DSA Login** option to create a DSA user, the user can only to connect to the database via the SQL server authentication mode. If you select **Use Existing DSA Login**, the user also has the option to connect using Windows authentication mode.

   Creating a DSA user is used when a new DSA login is added by the installation process. A SYSADMIN login must be provided during the next step of the installation.

   Using an existing DSA login requires creating the user manually prior to the installation by running DSA script (click **DSA Script** to access the script) provided using the SYSADMIN login, and then manually entering the new user during the installation. Note that SYSADMIN role is required in order to collect object data. However, DSA login can be created with read only privileges, yet any functionality associated with object collection is disabled. In order to map database files to Symmetrix devices, DSA needs to install a DSA Listener on the monitored environment. The listener can be installed by the installation process or manually. In order to install the DSA listener manually on the monitored environment, click **DSA Listener** to download it, copy it to the monitored environment and then follow the instructions in the readme.txt file.
4. Modify any number of the following values:
   - **Authentication**—Windows Authentication Mode or SQL Server Authentication Mode.

- **Login (DOMAIN\Username)** or **Login** —Login name.
- **Password**—Password
- **SYSADMIN Login**—SYSDAMIN login name. A SYSADMIN login is temporarily used to set up the DSA login.
- **SYSADMIN Password**—SYSADMIN password.
- **DB Host Name or IP**—Name or IP address of the database host.
- **DB Port**—Database listener port number.
- **Instance**—Instance name or ID.
- **Default Database Name**—Default database name.
- **Enable Device Mapping**— Select this option to enable mapping of the database files to the devices.
  Note : In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by clicking **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then clicking **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by clicking **Edit VM settings** on the **vSphere** client.
- **Mapper Port**—Port where the mapping process is executed.
- **Windows User**—Windows username. This is the user running the mapping executable. This user needs to be part of the Administrator group on the database server.
- **Windows Password**—Windows password.
- **Run Device Mapping immediately after installation**— Select this option to run device mapping immediately after installation. Device mapping is a process running in the monitored environment It maps all of the database files to devices.

5. Click **Finish**.

Editing monitored Oracle databases

This task allows you to edit Oracle database records already being monitored by the Database Storage Analyzer (DSA) application.

Procedure

1. Click the **ADMINISTRATION** tab to view the full list of monitored databases and their associated properties
2. Select a Oracle database and click **Edit** to open a dialog.
3. Modify any number of the following values:
   - **DB User Name**—Name of the database user.
   - **DB User Password**—User password for the database.
   - **DB SYS Password**—SYS password required by Database Storage Analyzer in order to create the Database Storage Analyzer database user.
   - **DB Host Name or IP**—Name or IP address of the database host.
   - **Port**—Database listener port number.
   - **Service**—Service name that may contain up to 16 alpha numeric characters.
   - **SID**—System identifier (SID) is a unique name for an Oracle database instance on a specific host.

- **Enable Device Mapping**— Select this option to enable mapping of the database files to the devices.
  Note: In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by clicking **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then clicking **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by selecting **Edit VM settings** on the **vSphere** client.
- **Mapper Host Name or IP**—Host name or IP address where the mapping process is executed.
- **Root User:**
  - **Root Password**—Password.
  - If login to remote server as root is disabled, specify SSH login user and password.
    - **Login User**—SSH user name.
    - **Login Password**—SSH user password.
  - **Use login user home directory to store mapping files**— Select this option to use the login user's home directory to store the mapping files.
  - **Run Device Mapping immediately after installation**— Select this option to run device mapping immediately after installation. Device mapping is a process running in the monitored environment It maps all of the database files to devices.
- **Sudo User**:
  - **Login User**—Sudo user name.
  - **Login Password**—Sudo user password.
  - **Run sudo interactively**—Select this option to run substitute user do (sudo) interactively. This allows users to run programs with the security privileges of another user (normally the superuser, or root).
  - **Use login user home directory to store mapping files**— Select this option to use the login user's home directory to store the mapping files. The temporary mapping files are stored by default in temp directories.
  - **Alternative path to sudo**—Use an alternative path to sudo if sudo is not available by default or you would like to specify a different sudo path.
  - **Use sudo with root password**—If sudo is configured in a way where the root password needs to be provided before running privileged actions, provide the root password here.
  - **Run Device Mapping immediately after installation**— Select this option to run device mapping immediately after installation. Device mapping is a process running in the monitored environment It maps all of the database files to devices.
    Note : In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by clicking **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then clicking **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by clicking **Edit VM settings** on the **vSphere** client.
  - **Mapper command line information**—Select this option to display the mapper command line details The file path displayed correlates to the directories created when you selected **Use login user home directory to store mapping**

files. You can confirm that the appropriate permissions are set for the sudo login users in the sudoers file.

4. Click **Finish**.

Editing monitored MS SQL server instances

This task allows you to edit MS SQL server instances to be monitored by DSA.

Procedure

1. Click the **ADMINISTRATION** tab to view the full list of monitored databases and their associated properties
2. Select an instance of MS SQL Server and click **Edit** to open a dialog.
3. Modify any number of the following values:
   - **Authentication**—Windows Authentication Mode or SQL Server Authentication Mode.
   - **Login** —Login name.
   - **Password**—Password .
   - **DB Host Name or IP**—Name or IP address of the database host.
   - **DB Port**—Database listener port number.
   - **Instance**—Instance name or ID.
   - **Default Database Name**—Default database name.
   - **Enable Device Mapping**—Select this option to enable mapping of the database files to the devices.
     Note : In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. You do this by selecting **UI Hosts => Virtual Servers** on the **Unisphere for VMAX** GUI and then selecting **Add VM Server**. In addition, you need to set the parameter disk `EnableUUID = "TRUE"` in the ESX server. You do this by selecting **Edit VM settings** on the **vSphere** client.
   - **Mapper Port**—Port where the mapping process is executed.
   - **Windows User**—Windows username.
   - **Windows Password**—Windows password.
   - **Run Device Mapping immediately after installation**—Select this option to run device mapping immediately after installation. Device mapping is a process running in the monitored environment It maps all of the database files to devices.
4. Click **Finish**.

Removing monitored databases or instances

This task allows you to remove databases or instances so that they are no longer monitored by DSA.

Procedure

1. Click the **ADMINISTRATION** tab to view the full list of monitored databases and their associated attributes.

2. Select a database or instance and click **Remove** to remove the selected database or instance.
3. Click **Yes** to confirm.

Starting statistics collection

This task allows you to start statistics collection for the monitored database or instance.

Procedure

1. Click the **ADMINISTRATION** tab to view the list of monitored databases or instances and their associated attributes.
2. Select a database and click **Start** to start statistics collection for the database or instance.
3. Click **OK** to confirm.

Stopping statistics collection

Procedure

1. Click the **ADMINISTRATION** tab to view the list of monitored databases or instances and their associated attributes.
2. Select a database and click **Stop** to stop statistics collection for the database or instance.
3. Click **OK** to confirm.

Running device mapping

This task allows you to manually run device mapping for a monitored database or instance. During device mapping, the list of database or instance files is copied using ssh to the monitored database or instance host. A process executing on the monitored database or instance host identifies the list of host physical devices associated with those files. The list is sent back and loaded into the DSA repository.

Procedure

1. Click the **ADMINISTRATION** tab to view the list of monitored databases and their associated attributes.
2. Select a database and click **Run Mapping** to run device mapping for the database record.
3. Click **OK** to confirm.

Setting configuration parameters

This task allows you to set database configuration parameters for data retention, thresholds, and device mapping.

Procedure

1. Click the **ADMINISTRATION** tab to view the list of monitored databases and their assoicated attributes.
2. Select a database or instance and click **Configure** to open the **Set Configuration Parameters** wizard.

3. Modify any number of the following values:
   - **Data retention for 5 min data (days)**—numbers of days to retain 5 minute data.
   - **Data retention for hourly data (months)**—numbers of months to retain hourly data.
   - **Data retention for daily data (months)**—numbers of months to retain daily data.
   - **First threshold for DB response time (ms)**—first threshold for database response time.
   - **Second threshold for DB response time (ms)**—second threshold for database response time.
   - **Mapping Schedule**—mapping schedule.
   - **Mapping Start Time**—mapping start time.
4. Click **Finish**.

Viewing the Performance Page

The Performance page provides database and storage information in a single view allowing DBAs and SAs to troubleshoot performance issues in a more timely manner.

Procedure

1. Click the **DASHBOARD** tab to view the full list of monitored databases and their associated attributes.
2. Click on a database name instance to view the **PERFORMANCE** tab.

   By default, the performance chart for **I/O Wait vs.Non-I/O Wait** is displayed.
3. Set the time range for the monitoring activity.
4. For Oracle databases, view the performance chart for **I/O Wait**. The x-axis represents time and the y-axis represents the following::
   - **I/O Wait**—Time spent waiting for I/O operations to complete.
   - **Non-I/O Wait**—Time spent waiting for non-I/O operations to complete.
   - **Redo**—Time spent waiting for redo operations to complete.
   - **Background**—Time spent by background process waiting for a database resource.
5. For Oracle databases, select **Active Session Wait** and view the performance chart for **Average Active Session Wait**. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class. The Oracle wait classes are listed below:
   - **Administrative**—Waits resulting from DBA commands (for example, an index rebuild).
   - **Application**—Waits resulting from user application code (for example, lock waits caused by row level locking or explicit lock commands).
   - **Cluster**—Waits related to Real Application Cluster resources (for example, global cache resources such as 'gc cr block busy').
   - **Commit**—This wait class only comprises one wait event - wait for redo log write confirmation after a commit (that is, 'log file sync').
   - **Concurrency**—Waits for internal database resources (for example, latches).
   - **Configuration**—Waits caused by inadequate configuration of database or instance resources (for example, undersized log file sizes, shared pool size).

- **Network**—Waits related to network messaging (for example, 'SQL*Net more data to dblink').
- **ORA CPU**—Sessions that are consuming CPU.
- **ORA Wait CPU**—Waits for CPU – Sessions waiting in the CPU queue.
- **OS CPU**—Displays the non-database process waiting/consuming CPU on the host.
- **Other**—Waits which should not typically occur on a system (for example, 'wait for EMON to spawn').
- **Scheduler**—Resource Manager related waits (for example, 'resmgr: become active').
- **System I/O**—Waits for background process I/O (for example, DBWR wait for 'db file parallel write').
- **User I/O**—Waits for user I/O (for example 'db file sequential read').

6. For MS SQL server instances, view the performance data for each database.
   - **DB Name**—Name of the database.
   - **RT Status**—Response time status.
   - **DB Read RT ms**—Database read response time in milliseconds.
   - **DB Log Write RT ms**—Database log write response time in milliseconds.
   - **DB IOPS**—Database IOPS.
   - **DB MB/s**—Database throughput.
   - **DB R/W Ratio**—Database Read/Write ratio.
7. For MS SQL server instances, view the performance chart for **I/O Wait vs Non-I/O Wait**. The x-axis represents time and the y-axis represents the following:
   - **I/O Wait**—Time spent waiting for I/O operations to complete.
   - **Non-I/O Wait**—Time spent waiting for non-I/O operations to complete.
   - **Redo**—Time spent waiting for redo operations to complete.
   - **Background**—Time spent by background process waiting for a database resource.
8. For MS SQL server instances, select **Wait Classes over Time** and view the performance chart for **Wait Classes over Time**. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class. The SQL server wait classes are listed below:
   - **Buffer I/O**
   - **Buffer Latch**
   - **CPU**
   - **Latch**
   - **Locks**
   - **Logging**
   - **Memory**
   - **Network I/O**
   - **Waiting for CPU**
   - **Other**

9. View the performance chart for **Response Time**. The x-axis represents time and the y-axis represents the response time for Reads, Writes and Redo Writes (Oracle only) or Log Writes (SQL server only) for both Database and Storage. All of the Storage statistics represents the storage devices that are associated with the database rather than the entire array. Clicking DB Read Response Time Thresholds enables the display of threshold lines on the chart. The thresholds are based on the database thresholds as defined in the system. If you hover over the line, you will see the threshold information. The first threshold for database response time is 20ms. The second threshold for database response time is 30ms.
10. View the performance chart for **IOPS**. The x-axis represents time and the y-axis represents IOPS for Reads, Writes and Redo Writes for both Database and Storage. Note that for Oracle12 PDB, the redo statistics can be seen at the CDB level only.
11. View the performance chart for **Throughput**. The x-axis represents time and the y-axis represents Throughput for Reads, Writes and Redo Writes for both Database and Storage.
12. View the performance chart for **I/O Size kb**. The x-axis represents time and the y-axis represents I/O size for reads, writes and redo writes for both database and storage. Information about the average IO size for reads, writes and redo writes (for both database and storage).
13. View the performance chart for **IOPS** for the Storage Back-end Activity and Tier Capacity. The x-axis represents time and the y-axis represents IOPS for Reads or Writes (displayed with or without Density) for EFD, FC or SATA drives.
14. View the performance chart for **Throughput** for the Storage Back-end Activity and Tier Capacity. The x-axis represents time and the y-axis represents Throughput for Reads or Writes (displayed with or without Density) for EFD, FC or SATA drives.
15. View the performance chart for **Tier Capacity** for the Storage Back-end Activity and Tier Capacity. The x-axis represents time and the y-axis represents Tier Capacity for Reads or Writes (displayed with or without Density) for EFD, FC or SATA drives.
16. Optional: Select a chart and click ⊟ to close a chart.
17. Optional: Select a chart and click ⊞ to re-open a chart.
18. Optional: Select a chart and click to export a performance chart as one of the following:
    - JPEG image
    - PNG image
    - PDF document
    - SVG Vector image

Viewing the Analytics Page

The Analytics page provides a means to analyze the I/O characteristics of specific business processes running in the database. Oracle wait events represent the various event types that Oracle sessions are waiting for. This information is helpful in the analysis of database bottlenecks. Wait classes are a way of grouping events to logical groups for high level analysis.

Procedure

1. Click the **DASHBOARD** tab to view the full list of monitored databases and their associated attributes.

2. Click on a database name instance to view the **PERFORMANCE** tab.
3. Click the **ANALYTICS** tab.
4. Set the time range for the monitoring activity.
5. For an Oracle database, select the **IO Wait Classes** context and view the performance chart for **Wait Events/Classes over Time**. The x-axis represents time and the y-axis represents the following wait classes:
   - **Single Block Read**—Wait class associated with Single Block read operations.
   - **Multi-Block Read**—Wait class associated with Multi-Block read operations.
   - **Direct I/O**—Wait Class that represents an I/O operation where the data is asynchronously read from the database files. It comprises of the following wait events: direct path read temp, direct path write temp, direct path write, and direct path read.
   - **Other I/O**—Wait Class that comprises of I/O events where the type of I/O (for example, random or sequential scan) is unknown.
   - **System I/O**—Wait Class associated with system I/O operations.
   - **Commit**—Wait Class associated with database commit.
6. For an Oracle database, select the **All Wait Classes** context and view the performance chart for **Wait Events/Classes over Time**. The x-axis represents time and the y-axis represents the following wait classes:
   - **Application**—Waits resulting from user application code (for example, lock waits caused by row level locking or explicit lock commands).
   - **Commit**—This wait class only comprises one wait event - wait for redo log write confirmation after a commit (that is, 'log file sync').
   - **Concurrency**—Waits for internal database resources (for example, latches).
   - **Configuration**—Waits caused by inadequate configuration of database or instance resources (for example, undersized log file sizes, shared pool size).
   - **CPU + Wait for CPU**—Sessions that are consuming CPU or waiting in the CPU queue.
   - **System I/O**—Waits for background process I/O (for example, DBWR wait for 'db file parallel write').
   - **User I/O**—Waits for user I/O (for example 'db file sequential read').
   - **Other**—Waits which should not typically occur on a system (for example, 'wait for EMON to spawn').
   - **Direct I/O**—Wait class that represents an I/O operation where the data is asynchronously read from the database files. It comprises of the following wait events: direct path read temp, direct path write temp, direct path write, and direct path read.
   - **Single Block I/O**—Wait class associated with Single Block I/O operations.
   - **Commit**—Wait class associated with database commit.
   - **Multi-Block I/O**—Wait class associated with Multi-Block I/O operations.
   - **Other I/O**—Wait class that comprises of I/O events where the type of I/O (for example, random or sequential scan) is unknown.
7. For an Oracle database, select the **IO Wait Events** context , and view the performance chart for **Wait Events/Classes over Time**. Wait events represent the exact wait type the Oracle process is waiting on. Every wait event belongs to a class of wait event. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class.

8. For an Oracle database, select the **All Wait Events** context , and view the performance chart for **Wait Events/Classes over Time**. Wait events represent the exact wait type the Oracle process is waiting on. Every wait event belongs to a class of wait event. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class.
9. For a MS SQL server instance, select the **Wait Classes** context and view the performance chart for **Wait Events/Classes over Time**. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database where waiting for each wait event/class. The SQL server wait classes are listed below:
   - **Buffer I/O**
   - **Buffer Latch**
   - **CPU**
   - **Latch**
   - **Locks**
   - **Logging**
   - **Memory**
   - **Network I/O**
   - **Waiting for CPU**
   - **Other**
10. For a MS SQL server instance, select the **Wait Events** context , and view the performance chart for **Wait Events/Classes over Time**. Wait events represent the exact wait type the MS SQL server process is waiting on. Every wait event belongs to a class of wait event. The x-axis represents the date and the y-axis represents the accumulative time the sessions in the database where waiting for each wait event/class.
    - **Running**
    - **PEGEIOLATCH_SH**
    - **Runnable**
    - **LCK_M_U**
    - **LATCH_EX**
    - **Null event**
    - **LCK_M_IX**
    - **PREEMPTIVE_OS_DECRYPTMESSAGE**
    - **LCK_M_X**
    - **OTHERS**
    - **PREEMPTIVE_OS_ENCRYPTMESSAGE**
11. Create a filtered view for the Oracle databases by using the following filters:
    - **Hint**—Name of an associated hint.
    - **Owner**—Name of the schema in Oracle that stores the objects.
    - **Object**—Object name.
    - **Program Like**—the executable name.
    - **Login User**—Database user name.
    - **Machine**—Machine name on which the session logged.

- **Node**—Node name.
- **Module**—Module name.
- **Service**—Service name.

12. Create a filtered view for the MS SQL server instances by using the following filters:
    - **DB Name**—Database name.
    - **Schema**—Schema name.
    - **Object**—Object name.
    - **Host**—Host name.
    - **Program Like**——the executable name.
    - **NT Domain**—NT domain name.
    - **NT User**—NT user name.
    - **Login name**—Login name.
    - **Login DB**—Login database.
13. Optional: Click **Apply** to apply the filter.
14. Optional: Click **Clear** to clear the applied filter.
15. Optional: Click to export the performance chart as one of the following:
    - JPEG image
    - PNG image
    - PDF document
    - SVG Vector image
16. View the properties of the Oracle database objects.
    - **Owner**—Owner of object, for example, SYSTEM.
    - **Object Name**—In the case of a partition object, the partition name is displayed.
    - **Sub-Object Name**—In the case where the object is partitioned, the sub-object name contains the partition name.
    - **Object Type**—Type of object, for example, table or index.
    - **Size**—Size of the object in MB.
    - **IO Wait Time**—I/O Wait time in seconds.
    - **IO Type**—I/O Type is one of the following:
        - Single Block Read
        - Multi Block Read
        - Direct I/O
        - System I/O
        - Commit I/O
        - Other I/O
17. View the properties of the MS SQL server instance objects.
    - **Database name**—Name of the database.
    - **Schema name**—Name of the schema.
    - **Object Name**—In the case of a partition object, the partition name is displayed.

- **Partition**—partition name.
- **Type**—Type of object, for example, table or index.
- **File Group**—Associated file group.
- **Data Compression**—Associated data compression.
- **IO Wait Time**—I/O Wait time in seconds.

18. Optional: Select **Group by Object Name** to group database objects by object name.

    By default, each partition of an object is displayed in a separate row. Selecting "Group by Object name" will group all the partitions of an object to a single row, displaying total activity for that object. For MS SQLserver instances, the following properties display:

19. View the properties of the Symmetrix Devices. The following properties display:

    - **Host Device Name**—Name of host physical device.
    - **Symmetrix Device**—Note for VMware virtual devices there could be multiple Symmetrix devices associated with the host device.
    - **IO Wait Time**—I/O Wait time in seconds.
    - **IO Type**—I/O Type is one of the following:
      - Single Block Read
      - Multi Block Read
      - Direct I/O
      - System I/O
      - Commit I/O
      - Other I/O
      - DataStore—The VM DataStore where the host device resides. This is applicable for VMWare virtual devices only.
      - Host Device Size GB - The size of the host device in GB.

      You can also view Oracle database session contextual information. The following is the list of session context types:

      - Program Names
      - Login Users
      - Machines
      - Modules
      - Nodes
      - Services

      You can also view MS SQL server database session contextual information. The following is the list of session context types:

      - Hosts
      - Program names
      - NT Domain
      - NT User
      - Login Name
      - Login DB

20. View the properties of the different session types. The following properties display:
    - **Elapsed Time**—Total wait time for the session.
    - **IO Wait Time**—Total wait time for the session.
    - **IO Type**—I/O Type is one of the following:
        - Single Block Read
        - Multi Block Read
        - Direct I/O
        - Single Block Read
        - Commit I/O
        - Other I/O
    - **Wait Classes**—Wait Classes type is one of the following:
        - **Administrative**—Waits resulting from DBA commands that cause users to wait (for example, an index rebuild).
        - **Application**—Waits resulting from user application code (for example, lock waits caused by row level locking or explicit lock commands).
        - **Cluster**—Waits related to Real Application Cluster resources (for example, global cache resources such as 'gc cr block busy').
        - **Commit**—This wait class only comprises one wait event - wait for redo log write confirmation after a commit (that is, 'log file sync').
        - **Concurrency**—Waits for internal database resources (for example, latches).
        - **Configuration**—Waits caused by inadequate configuration of database or instance resources (for example, undersized log file sizes, shared pool size).
        - **Network**—Waits related to network messaging (for example, 'SQL*Net more data to dblink').
        - **CPU + Wait for CPU**—Sessions that are consuming CPU or waiting in the CPU queue.
        - **Other**—Waits which should not typically occur on a system (for example, 'wait for EMON to spawn').
        - **Scheduler**—Resource Manager related waits (for example, 'resmgr: become active').
        - **System I/O**—Waits for background process I/O (for example, DBWR wait for 'db file parallel write').
        - **User I/O**—Waits for user I/O (for example 'db file sequential read').
21. View a report in a dialog containing the details of sessions that are accessing specific database objects by selecting the check box beside a specific database object and clicking **Details**.
22. View a report in a dialog containing the details of sessions that are accessing specific devices by selecting the check box beside a specific device and clicking **Details**.
23. View a report in a dialog containing the details of database objects and devices that are associated with specific sessions by selecting the check box beside a specific session and clicking **Details**.
24. Select one or more database objects and click **Add to Hint** to add the object(s) to database hint.

Adding hints from Analytics screen

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click on a database instance to open the **Performance** tab.
2. Click the **Analytics** tab.
3. Select one, more than one, or all database objects and click **Add to Hint** to open the Hint wizard.
4. Click **New Hint** and click **Next**.
5. Type the hint name.
6. Select the **Hint Priority**. Possible values are 1 through 3, with 1 being the highest.
   - Priority 1—Simulates a diamond service level. It marks all extents as active including the inactive ones and move data to EFD ahead of time. The primary use case for this priority is a periodic process such as end of month process where the assumption is that most of the data is inactive during the month.
   - Priority 2—Simulates a Platinum (OLTP) service level. It marks all extents as active (even the inactive ones) but there is no immediate promotion. The primary use case is ensuring that a given process receives better response time than other processes in the storage group even if the data was inactive.
   - Priority 3—Simulates a Platinum (DSS) service level change response time target without marking the inactive extents. The primary use case is to ensure that a given process receives better response time than other processes in the storage, yet it is less powerful than priority 2 as it only promotes active data.
7. Select the **Hint Type**. Possible values are:
   - **One Time**—Type a value for the start date and time and the end date and time.
   - **Ongoing**—Hints are applied for the selected objects on an ongoing 24/7 basis.
   - **Recurrence** — Select the recurrence pattern.
     - **Daily**—Occurs on a daily basis.
     - **Weekly**—Select one or more days of the week.
     - **Monthly**—Type the day of the month.
8. Type the values for the recurrence range—the start date and end date.
9. Optional: Select **Group by Object Name** to group the objects by object name.
10. Optional: Add more database objects to the hint by clicking + and typing owner and object name or partition.

    For partition objects, if you select a recurrence hint when **Group by Object Name** is selected and you select a date partition table, you can set a pattern to determine which partition would be hinted for every iteration by using the partition access

pattern column in the grid. The partition access pattern determines the number of recent partitions that would be sent. For example:
For a monthly partition table, access pattern set to 3, the hint on April 1st will send January, February and March partitions (assuming that the April partition is still empty or very small). Note that it would skip very small partitions in order to avoid “future” partitions that have no data. To send all partitions, set access pattern field to 0.

11. Optional: Select one or more Database objects and click to remove the objects from the hint.
12. Click **Finish**.

Results

The hint has been added and its status is Enabled.

Adding hints from Hinting screen

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **Hinting** tab.
2. Select a database.

   Only databases, where hints are supported, are displayed.
3. Type the hint name.
4. Select the **Hint Priority**. Possible values are 1 through 3, with 1 being the highest.
5. Select the **Hint Type**. Possible values are:
   - **One Time** — Type a value for the start date and time and the end date and time.
   - **Ongoing** —Hints are applied for the selected objects on an ongoing 24/7 basis.
   - **Recurrence** — Select the recurrence pattern.
     - **Daily**—Occurs on a daily basis.
     - **Weekly**—Select one or more days of the week.
     - **Monthly**—Type the day of the month.
6. Type the values for the recurrence range—the start date and end date.
7. Optional: Select **Group by Object Name** to group the objects by object name.
8. Optional: Add more database objects to the hint by clicking + and typing owner and object name or partition.

   If any object is added to the list, you cannot change both the database and **Group by Object Name** selections.
9. Optional: Select one or more database objects and click to remove the objects from the hint.

10. Click Finish.

Results

The hint has been added and its status is Enabled.

Adding database objects to an existing hint

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.

A hint helps to assure the best response time for mission critical processes or database objects. You can also add database objects to existing hints when editing a hint (see Editing hints on page 29).

Procedure

1. Click on a database instance to open the **Performance** tab.
2. Click the **Analytics** tab.
3. Select one, more than one, or all database objects and click **Add to Hint** to open the Hint wizard.
4. Click **Add to an Existing Hint (not grouped by object name)**.
5. Optional: Select **Include expired hints** to include expired hints in the existing hint selection (see next step).
6. Select an existing hint from the list.
7. Click **Next**.
8. Optional: Select **Group by Object Name** to group the objects by object name.
9. Optional: Add more database objects to the hint by clicking +, typing owner and object name or partition, and clicking **Add**.
10. Optional: Select one or more Database objects and click trash to remove the objects from the hint.
11. Click **Finish**.

Editing hints

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **HINTING** tab to view the list of hints.
2. Select the radio button to the left of the hint that you wish to edit and click **Edit** to open the Hint wizard.

3. Optional: Modify the hint name.
4. Optional: Modify the **Hint Priority**. Possible values are 1 through 3, with 1 being the highest.
5. Optional: Modify the **Hint Type**. Possible values are:
   - **One Time** — Type a value for the start date and time and the end date and time.
   - **Ongoing** —Hints are applied for the selected objects on an ongoing 24/7 basis.
   - **Recurrence** — Select the recurrence pattern.
     - **Daily**—Occurs on a daily basis.
     - **Weekly**—Select one or more days of the week.
     - **Monthly**—Type the day of the month.
6. Optional: Modify the values for the recurrence range—the start date and end date.
7. Optional: Add more database objects to the hint by clicking + and typing owner and object name or partition.
8. Optional: Select one or more Database objects and click 🗑 to remove the objects from the hint.

   Note: there has to be at least one database object associated with the hint.
9. Click **Finish**.

Viewing hints

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **HINTING** tab to view the full list of hints and their associated properties.

   The following properties display for each hint:
   - **Hint Name**—Name of the hint.
   - **DB Name**—Name of the database associated with the hint.
   - **DB Host**—Name of the database host associated with the hint.
   - **Symmetrix ID**—Identity of the current Symmetrix that the database is running on.
   - **FAST Storage Groups**—Storage Groups associated with the database.
   - **Priority**—Priority associated with the hint.
   - **Start Date**—Start date and time.
   - **End Date**—End date and time.
   - **Next Run**—Next time the hint will be run.
   - **Occurence**—Occurrence of the hint—one time, ongoing or scheduled.
   - **Created By**—Name of the user that created the hint.

- **Status**— Status of the hint—Enabled, Disabled or Completed.
- **Last Sent Time**— Last time the hint was run.
- **Last Sent Status**— Status of the hint when it was last run.

2. Optional: Click to export the screen contents to Excel.
3. Optional: Create a filter for the list by typing any combination of DB Name, DBHost, Symmetrix ID, Hint Name, and Status.
4. Optional: Click **Apply** to apply the filter.
5. Optional: Click **Clear** to clear the applied filter.
6. Optional: Click **Refresh** to refresh the screen display.

Enabling hints

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.
- The hint is disabled.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **HINTING** tab to view the list of hints.
2. Select the radio button to the left of the hint that you wish to enable and click **Enable**.
3. Click **OK** to confirm the operation.

Disabling hints

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.
- The hint is enabled.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **HINTING** tab to view the list of hints.
2. Select the radio button to the left of the hint that you wish to disable and click **Disable**.
3. Click **OK** to confirm the operation.

Removing hints

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.
- To perform hinting operations, you must be a DSA Admin with hinting permission.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **HINTING** tab to view the list of hints.
2. Select the radio button to the left of the hint that you wish to remove and click **Remove**.
3. Click **OK** to confirm the operation.

Viewing hint logs

Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle database.

A hint helps to assure the best response time for mission critical processes or database objects.

Procedure

1. Click the **HINTING** tab.
2. Click **Logs** to view the full list of view the full list of hint logs and their associated properties.

   The following properties display for each hint log:

   - **Sent Time**—Time that the hint was sent to the storage system.
   - **Hint Name**—Name of the hint.
   - **DB Name**—Name of the database associated with the hint.
   - **DB Host**—Name of the database host associated with the hint.
   - **Priority**—Priority associated with the hint.
   - **Symmetrix ID**—Identity of the current Symmetrix that the database is running on.
   - **Start Date**—Start date and time.
   - **End Date**—End date and time.
   - **Status**—Status of the hint when it was last run.
   - **Description**—This lists the name of the user that created the hint.
   - **Size(GB)**—The size of the log.
   - **Hint ID**— Hint identity received from the storage system.

3. Optional: Click to export the screen contents to Excel.

4. Optional: Create a filter for the list by typing any combination of Hint Name, Start Time, End time, Status, DB Name, DB Host, and Symmetrix ID.
5. Optional: Click **Apply** to apply the filter.
6. Optional: Click **Clear** to clear the applied filter.
7. Optional: Click **Refresh** to refresh the screen display.