

Foglight™ *for Oracle 5.7.5.1*

Monitoring Oracle Database Systems
User and Reference Guide



© 2015 Dell Inc.
ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Dell Inc.

The information in this document is provided in connection with Dell products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Dell products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, DELL ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL DELL BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF DELL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Dell makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Dell does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Dell Inc.
Attn: LEGAL Dept
5 Polaris Way
Aliso Viejo, CA 92656

Refer to our web site (software.dell.com) for regional and international office information.

Patents

Foglight™ is protected by U.S. Patents # 7,979,245 and 8,175,862. Additional Patents Pending.

Foglight™ for VMware is protected by U.S. Patents # 8,175,863 and 8,364,460. Additional Patents Pending.

Foglight™ for Hyper-V is protected by U.S. Patents # 8,175,863 and 8,364,460. Additional Patents Pending.

Foglight™ for Resource Optimization is protected by U.S. Patents # 8,171,201 and 8,332,571.


For more information, go to <http://software.dell.com/legal/patents.aspx>.


Trademarks


Dell, the Dell logo, and Foglight, IntelliProfile, PerformaSure, and Tag and Follow are trademarks of Dell Inc. "Apache HTTP Server", Apache, "Apache Tomcat" and "Tomcat" are trademarks of the Apache Software Foundation. Google is a registered trademark of Google Inc. Chrome, Android, and Nexus are trademarks of Google Inc. Red Hat, JBoss, the JBoss logo, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the U.S. and other countries. CentOS is a trademark of Red Hat, Inc. in the U.S. and other countries. Microsoft, .NET, Active Directory, Internet Explorer, Hyper-V, SharePoint, SQL Server, Windows, Windows Vista and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. AIX, IBM, and WebSphere are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Sun, Oracle, Java, Oracle Solaris, and WebLogic are trademarks or registered trademarks of Oracle and/or its affiliates in the United States and other countries. SPARC is a registered trademark of SPARC International, Inc. in the United States and other countries. Products bearing the SPARC trademarks are based on an architecture developed by Oracle Corporation. OpenLDAP is a registered trademark of the OpenLDAP Foundation. HP is a registered trademark that belongs to Hewlett-Packard Development Company, L.P. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. MySQL is a registered trademark of MySQL AB in the United States, the European Union and other countries. Novell and eDirectory are registered trademarks of Novell, Inc., in the United States and other countries. VMware, ESX, ESXi, vSphere, vCenter, vMotion, and vCloud Director are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Sybase is a registered trademark of Sybase, Inc. The X Window System and UNIX are registered trademarks of The Open Group. Mozilla and Firefox are registered trademarks of the Mozilla Foundation. "Eclipse", "Eclipse Foundation Member", "EclipseCon", "Eclipse Summit", "Built on Eclipse", "Eclipse Ready", "Eclipse Incubation", and "Eclipse Proposals" are trademarks of Eclipse Foundation, Inc. IOS is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. Apple, iPad, iPhone, Xcode, Mac OS, and Safari are trademarks of Apple Inc., registered in the U.S. and other countries. Ubuntu is a registered trademark of Canonical Ltd. Symantec and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. YAST is a registered trademark of SUSE LLC in the United States and other countries. Citrix and XenDesktop are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. AlertSite and DéjàClick are either trademarks or registered trademarks of Boca Internet Technologies, Inc. Samsung, Galaxy S, and Galaxy Note are registered trademarks of Samsung Electronics America, Inc. and/or its related entities. MOTOROLA is a registered trademarks of Motorola Trademark Holdings, LLC. The Trademark BlackBerry Bold is owned by Research In Motion Limited and is registered in the United States and may be pending or registered in other countries. Dell is not endorsed, sponsored, affiliated with or otherwise authorized by Research In Motion Limited. Other trademarks and trade names may be used in this document to refer to either the entities claiming the

marks and names of their products. Dell disclaims any proprietary interest in the marks and names of others.

Legend

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

 **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

 **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Foglight *for Oracle* User and Reference Guide
Updated - June 2015
Software Version - 5.7.5.1

Contents

Installing and Configuring Agents	7
Installation	7
Installing and Configuring Agents	7
Upgrading to the Current Version	7
Overview	8
Running the Upgrade Wizard	8
Migrating Agents	10
About Monitoring Extensions	10
SQL Performance Investigator Extension	10
Operating System Extension	11
VMWare Extension	11
Configuring the On Demand Port on the Agent Manager Concentrator	11
Installing and Monitoring a Single Oracle Instance	13
Discovering Multiple Instances to Monitor	14
Loading a TNS file	16
Adding Agents for Monitoring Exadata Environments	18
Adding ASM Instances	20
Configuring Multiple Instances for Monitoring using Silent Installation	21
Using Foglight <i>for Oracle</i>	26
The Foglight <i>for Oracle</i> Instance Homepage	26
Viewing the Databases Dashboard	26
Assigning Instances to Users	28
Selecting an Instance to Monitor	28
Foglight <i>for Oracle</i> Overview Page	33
SQL Performance Investigator	35
Reviewing Overall Activity	38
Reviewing RAC Historic Activity	41
Reviewing Instance Historic Activity	45
Viewing the Foglight <i>for Oracle</i> Real-time Summary Page	48
Oracle Activity Drilldown	53
Pluggable Databases Drilldown	61
Storage Drilldown	62
Reviewing Configuration Settings	73
Reviewing the Alert and Error Logs	76
Reviewing ASM Instances	78
Reviewing Exadata-related Information	80
Administering Foglight <i>for Oracle</i>	84
Opening the Databases Administration Dashboard	84
Reviewing the Administration Settings	84
Defining Connection Details	85
Defining Data Collection and Storage Options	89
Customizing Alarms for Foglight <i>for Oracle</i> Rules	90

Defining Retention Policies	99
Defining the Collection Settings	100
Configuring the Alert Log Panel Display	102
Configuring the List of Listeners per Instance	103
Configuring User-defined Collections	104
Configuring the On-demand Data Port	107
Changing the Monitoring Mode of the RAC One Node Agents	107
Administering SQL Performance Investigator	108
Reporting	109
Generating Reports	109
Generating Reports for a Foglight <i>for Oracle</i> RAC or Instance	109
Studying the Various Reports	109
Reference	111
Overview of Collections	111
Alert Log	111
Archive Destination Miscellaneous	112
Archive Logs Storage	113
ASM Diskgroups	115
ASM Diskgroups Balance	116
ASM Disks	116
ASM Instance Total	118
ASM Operations	118
Backup Config	119
Backup Jobs	120
Backup Type Status	120
Cluster Latency	121
Cluster Misses	123
Datafile Fragmentation	124
Datafile IO Activity	125
Datafiles Storage	126
Dataguard Destination Status	127
Dataguard Redo Apply	128
Dataguard Status	129
Dump File	130
Event Wait	130
Exadata Cell Flashcache	131
Exadata Cell General	132
Exadata Cell InfiniBand	133
Exadata Cell InfiniBand Ports	134
Exadata Cell InfiniBand Switches	135
Exadata Cell Statistics	135
Exadata Cell General	136
Filesystem Disk	137
Flashback Area	138
History Locks Tree Instance	139
History Locks Tree RAC	140

Host Properties	141
Initialization Parameters	142
Instance General	143
Invalid Objects	144
Invalid Objects List	145
IO Activity	146
Jobs Queue	150
Latches Activity	151
Locks Activity	151
Locks Breakdown	152
Locks Instance	153
Locks Tree	154
Locks Tree RAC	155
OS and Instance Statistics	156
OS File System IO Summary	161
OS General Network Summary	162
Parallel	162
PGA Statistics	164
Pluggable Databases Information	165
Pluggable Databases State	165
Redo Archive Performance	166
Redo Instances Miscellaneous	167
Redo Logs Member Storage	167
Redo Log Storage	168
Sessions Summary	169
Servers	170
Tablespaces Free Space Fragmentation	171
Tablespaces Storage	171
Top Blocked Objects	173
Top Sessions	174
Top SQLs	176
Top SQLs Full Text	178
Top SQLs Text	178
Total Archive Storage	179
Total Database Storage	179
Total Instance Waits	180
Undo Activity	180
Total Summary	181
Usability	182
Glossary	184
About Dell	192
Contacting Dell	192
Technical support resources	192

Installing and Configuring Agents

Welcome to the *Foglight for Oracle User and Reference Guide*. This guide provides configuration instructions, conceptual information and instructions on how to use the Foglight *for Oracle* cartridge. It describes the dashboards included with the cartridge and how they are used for collecting monitoring data from the entire relational database management system, as well as the cartridge's interaction with and support of additional services and modules, such as replication and virtualization.

This guide is intended for any user who wants to know more about monitoring Oracle instances through Foglight *for Oracle*, and the steps required for discovering and configuring the Oracle agent. It is also meant for those users who want to learn about the methods used for configuring and applying user-defined settings.

Installation

Information regarding the installation of Foglight *for Oracle*, its pre-requisites and permission can be found in the *Foglight for Oracle Getting Started Guide*. You can view this guide on Dell's website at <http://documents.software.dell.com/foglight-for-oracle-cartridge/5.7.5.1/foglight-for-oracle-getting-started-guide/>.

For further information to help you determine your environment's hardware requirements, refer to the Foglight *for Oracle* Hardware Sizing Guide, which is available at <http://documents.software.dell.com/foglight-for-oracle-cartridge/5.7.5.1/foglight-for-oracle-hardware-sizing-guide/>.

Installing and Configuring Agents

Foglight *for Oracle* monitors the Oracle Server database activity by connecting to and querying the Oracle database. The agents provided monitor the Oracle database system. The dashboards included with the cartridge provide a visual representation of the status of the major components of the Oracle agents. They allow you to determine any potential bottleneck in database performance.

Upgrading to the Current Version

Starting to work with a Foglight *for Oracle* cartridge requires upgrading to the current version of both the cartridge and the Foglight Agent Manager that runs the cartridge.

This chapter contains instructions for using the upgrade wizard.

- ❗ **IMPORTANT:** Foglight *for Oracle* does not support the upgrade method of placing the new version under the folder `FGLHOME/upgrade/cartridge` folder.
- ❗ **IMPORTANT:** If upgrading to the current version of Foglight *for Oracle* in a Federation architecture, the upgrade should be applied first to the stand-alone Management Servers (the Federated Children) and then to the central Management Server (the Federation Master). If the Federation Master is upgraded first, this server will display incorrect information regarding the number of instances being monitored, as reported under the Status Summary section. All the same, this information will be displayed correctly on the Federated Children.

Overview

Any upgrade of Foglight *for Oracle* to the latest version, which is not met by similar upgrade of the database cartridge components, is detected upon opening the Databases dashboard, when the upgrade wizard starts automatically.

Selecting the check box at the bottom left of the screen prevents this wizard from appearing when entering the Databases dashboard. Nevertheless, if several database cartridge components still require upgrade, the need to upgrade them is indicated in the dashboard. The indication is evident both in the caption *Upgrade required*, which is displayed in red to the right of the requested instance, and in the button Upgrade, which appears only if upgrade is required.

NOTE: Even though instances whose components need to be upgraded appear in the Databases table, such instances cannot be accessed by clicking them, until they have been fully upgraded.

Running the Upgrade Wizard

To upgrade the requested instances:

- 1 Click Upgrade.

The screen that appears now is the same screen that appears by default upon entering the Databases dashboard when one or more database cartridge components require upgrading. While the need for upgrade is indicated for both Foglight for Oracle and Foglight *for Oracle*, the upgrade is carried out separately for each database cartridge type.

- 2 Click Foglight *for Oracle*.


A Welcome page displays What's New in this release, including bug fixes. At the bottom of the page, is a link to the Rule Modification Overview page, where you can review the list of rules that have been modified and are in the following states:

- Removed – The rule is obsolete and was deleted from the Management Server.
- New – The rule is new in this release. To review its definition, click View new rule.
- Updated – The rule was updated in this release. To review the updated rule, click View updated rule. If a rule targeted for update was modified by a user, the user's modified rule is copied and disabled before the updated rule is installed. To review the modified rule, click View user-updated rule.

TIP: To avoid having to repeat rule modifications after an upgrade, do not edit rules with the DBO prefix in the Rule Management dashboard. Use the Alarms view on the Databases Administration dashboard instead. When you make your edits on the Alarms view, the edits are saved separately and applied over the predefined rules. For instructions, see [Customizing Alarms for Foglight for Oracle Rules](#) on page 90.

- 3 If SQL Performance Investigator is *not* configured, proceed to [Step 5](#).
- 4 If SQL Performance Investigator is configured, the SQL PI repository must be upgraded prior the agents upgrade.
- 5 Select the Agent Managers to upgrade.

Figure 1. Select the Agent Managers to install.



NOTE: The Other Agents column in the table indicates whether the specified Foglight Agent Manager runs agents of other cartridge types. If a value is displayed in this column, go to Dashboards > Administration > Agents > Agent Status to view which other agents run under the specified Foglight Agent Manager. If the other agents appear in the list below, upgrading Foglight Agent Manager to the latest version is safe. For agents of any other cartridge type, consult the documentation of the respective cartridge types.

- DB_DB2_*
- DB_SQL_Server_*
- DB_Oracle_*
- Sybase_MDA
- UnixAgent
- WindowsAgent

6 Click Upgrade the selected Foglight Agent Managers.

A progress bar appears.

7 After the upgrade is complete, click Next.

The next screen is used for granting privileges to users that were detected as possibly requiring additional privileges to ensure full functionality.

8 Select the agents displayed in the table and click Validate connectivity.

9 If the status row of one or more agents displays the status *Insufficient privileges*, complete the following steps:

- a Select all agents that require privileges updates, and click Grant privileges.
- b In the Grant Database Privileges dialog box, enter a SYSADMIN user name and password.
- c To view the script used for granting the privileges, click View script.

IMPORTANT: The script for granting privileges is invoked using a popup. To view the script, ensure that pop ups are not blocked on the page.

- d Click Grant database privileges.
- 10 Foglight Performance Analysis is no longer supported and has been replaced by the SQL Performance Investigator. All Instances that are configured to use Foglight Performance Analysis will be deprecated and replaced with the SQL Performance Investigator capabilities.

During this process you are prompted to select a dedicated Agent Manager on which to install an external database repository. The SQL PI configuration will take into account the resources needed for monitoring with SQL PI and will prompt you with further information.

All information regarding the resource allocations needed for SQL PI monitoring can be found in the Oracle Server cartridge *Hardware Sizing Guide*.

Migrating Agents

The SQL Performance Investigator extension is only supported on Windows 64-bit and Linux 64 (Red-Hat 6, SUSE 11 and CentOS 6) . As a result, you may need to migrate Oracle agents from an Agent Manager running on Linux to a Windows Agent Manager.

In addition, the migration tool allows you to perform resource balancing between the agent managers running Oracle agents with and without the SQL PI extension installed.

To migrate an agent to a different Agent Manager:

- 1 Open the migration wizard by navigating to Dashboards > Databases> Support Dashboard.
- 2 Select the Agent Manager that will be the destination host for the migrated Oracle Agents. Click Next.
- 3 On the Oracle Agents for Migration pane, the left-hand column displays a list of available Agent Managers. Select an Agent Manager to view a list of the agents installed on it.
- 4 From the list, select the agents that you want to migrate and move them to the Agent Manager Destination column.
- 5 If required, select agents and set the Windows credentials.
- 6 Click Next to migrate the agents.
- 7 When migration has completed, click Finish to exit the Wizard.

About Monitoring Extensions

During the installation process you can choose to install and configure one or more of the monitoring extensions. The monitoring extensions provide a more in-depth analysis of the monitored instance and the environment it is running on, creating a whole and unified status.

SQL Performance Investigator Extension

SQL Performance Investigator allows you to rapidly identify bottlenecks, anomalies, and application trends by focusing on top resource consumers and providing multi-dimensional SQL domain drilldowns. SQL PI allows you to:

- Monitor real-time Oracle database performance at a glance
- Gather and diagnose historical views
- Identify and anticipate performance issues
- Analyze and optimize execution plan changes
- Compare day-to-day values to identify anomalies and application changes

① | **NOTE:** SQL Performance Investigator requires a license. If you are using a trial version and would like to request pricing, contact <https://software.dell.com/register/57891>.

① | **NOTE:** SQL PI requires a repository database that is installed automatically on the Agent Manager.

Operating System Extension

Monitoring the operating system allows you to identify resource consumption and provides a full view of the server health. An Infrastructure agent, which is created automatically as part of the monitoring process, monitors the operating system.

① | **NOTE:** The Operating System extension is enabled by default.


VMWare Extension

Monitoring the VMWare system allows you to identify resource consumption and provides a full view of the data center and ESX health when the server is part of a VMWare environment.

Configuring the On Demand Port on the Agent Manager Concentrator

If a firewall is installed between the Management Server and the Agent Manager, the bidirectional RMI connection shown in the figure above will not be possible because RMI requests require direct connection.

This issue can be resolved by creating an RMI proxy server on the concentrator – an Agent Manager instance that works similarly to an HTTP proxy. This RMI proxy server can be configured to accept RMI requests from the Management Server and direct them to the Agent Manager instances (called downstream instances). For additional details about the concentrator, see *Configuring an Agent Manager Instance as a Concentrator* in the *Agent Manager Installation Guide*.



A concentrator agent, created specifically for Foglight *for Oracle*, allows controlling the RMI server so that the server is initialized (starts collecting data) when the agent starts, and stops when data retrieval by the agent is brought to an end. The agent, which is created manually by the user, requires configuring the following parameters:

- A list of the downstream hosts that are connected through the concentrator
- The RMI port number

Each time the Foglight Management Server sends an RMI request, the list of downstream hosts submitted by the concentrator agent is being searched.

- Inclusion of the target RMI server host name within this list indicates that a Foglight Agent Manager concentrator exists, in which case a connection is made to the proxy RMI server.
- If the list does not contain the target host name, the connection is made directly to the target RMI server.

To create the concentrator agent:

- 1 Go to Dashboards > Administration > Agents > Agent Status.
- 2 Click Create Agent.
- 3 Select the agent type DB_Oracle_Concentrator.
- 4 Enter a name of your choice in the Instance Name field.
- 5 Click Create.
- 6 After the creation process is completed successfully, click OK.

To edit the concentrator agent connection details:

- 1 Select the newly created concentrator agent.
- 2 Click Edit Properties.
- 3 Click Modify properties for this agent only.
- 4 Enter the number of the proxy RMI server port, or accept the default number.
- 5 Enter the concentrator Foglight Agent Manager host name (optional).
- 6 Click Edit to edit the list of downstream Agent Managers.
- 7 Click Add Row.
- 8 Enter the name of the requested Foglight Agent Manager host.

ⓘ IMPORTANT: When editing the properties of the Foglight Agent Manager concentrator (Dashboards > Administration > Agents > Agent Status), the name of the Foglight Agent Manager should be entered in the Downstream FglAMs section exactly as it appears in the topology, under Home > Agents > All agents > <Agent name> > RMI data > Fglam host property. The <Agent name> parameter refers to any of the Foglight for Oracle agents that reside on the selected Foglight Agent Manager.

- 9 Repeat [Step 7](#) to [Step 8](#) as many times as required.
- 10 Click Save Changes.
- 11 Click Save at the lower right corner of the screen.
- 12 Click Back to Agent Status.
- 13 Select the newly created concentrator agent.
- 14 Click Activate.
- 15 Click OK to complete the process.

ⓘ IMPORTANT: In an environment that includes a Foglight Agent Manager concentrator, when upgrading the environment where the concentrator does not contain any regular agents, the concentrator is not upgraded. Therefore, after the upgrade process takes place the Foglight Agent Manager concentrator has to be deployed manually.

Installing and Monitoring a Single Oracle Instance

Enabling the Foglight Management Server to monitor Oracle instances requires the creation of the Foglight agents that monitor these instances and ensuring that these agents communicate properly with the Foglight Management Server.

Foglight *for Oracle* provides a graphic, intuitive method for creating and configuring multiple agents, which can be used instead of Foglight's default method for creating agents and editing their properties using the Agent Administration dashboard. Foglight *for Oracle* allows you to run a wizard that provides a common entry point for adding and discovering all database instances within a user-specified range, and then configuring these instances for monitoring.

- ① **IMPORTANT:** When running Foglight *for Oracle* in a Federation architecture, neither the creation nor the administration of agents can be accomplished from the central Foglight Management Server (the Federation Master). These two tasks should be carried out from the stand-alone Management Servers (the Federated Children).

To run the instance installation wizard:

- 1 On the navigation panel, click Homes > Databases.
- 2 Click Monitor > Oracle in the upper left corner of the Databases View.

The Monitor Oracle Instance dialog box appears.

- ① **NOTE:** If a user-defined database group is currently selected, the databases table's title displays the name of this group instead of All; however, all newly discovered or created databases are added to the general (All) group of databases.

- 3 Choose the agent manager on which the agent will be running. The default is the agent manager with the least agents installed.
 - a Click the Agent Manager Host *<agent manager>* link located at the top of the Oracle Connection section.

A dialog box appears with a list of all agent managers connected to the Foglight management server.
 - b Select the appropriate host name and click Set.

- ① **IMPORTANT:** You have the option set this host as the default for all future installations.

- 4 Use the Connection Details section to select the requested method for entering connection details:
 - TNS – used for selecting a database from the tnsnames.ora file by clicking Load file.
 - Host – Manually enter Host, port and either service name or SID.
- 5 Use the Login Credentials section to specify the username and password to be used for monitoring the Oracle instance.
- 6 Select an Alarm Sensitivity Level to determine what level of alarms the system will store and display for this instance. The default is Normal. Alarm levels include:
 - Essential – Store and display only critical or fatal alarms.
 - Normal – Store and display most alarms – essential and best practices; only critical and fatal statistical alarms.
 - Tuning – Store and display all Oracle alarms sent to Foglight.
- 7 **Optional** – In the Monitoring Extensions pane, click the SQL PI monitoring extension. You are prompted to choose the Agent Manager on which the SQL PI repository will be installed.
- 8 **Optional** – In the Monitoring Extensions pane, click the Operating System link.

To configure the extension, choose the connection details of the host on which the Oracle instance is running:

- Local user – using the same credentials that were used to launch the Agent Manager.
- Windows – the user name (in the domain\user name format) and password, used for logging in to a Windows-based monitored host.
- SSH (DSA) – using the SSH private key, generated via the DSA algorithm.

To enable OS monitoring using a DSA key (UNIX users only):

- a Create one or more Oracle agents using the Oracle Monitoring Installer wizard (without OS monitoring).
 - b Navigate to the Foglight Management Server credentials view (Dashboards > Administration > Cartridges > Credentials > Manage Credentials).
 - c Add credential of type DSA key and link this credential to the agents created in [Step a](#). For details, see "Adding and deleting credentials" in the *Foglight Administration and Configuration Guide*.
 - d Go to Agent Status screen (Administration > Agents > Agent Status)
 - e Select the agent or agents created in [Step a](#).
 - f Click Edit Properties.
 - g Select Modify the private properties for this agent.
 - h Under the Agent Management section, set Enable OS collections to True.
- SSH (RSA) – using the SSH private key, generated via the RSA algorithm.
 - SSH (login credentials) – the user name and password used for logging in to a UNIX-based monitored host.
- 9 Optional – In the Monitoring Extensions pane, click Collect VM statistics.

To configure the extension, select the connection details of the vCenter or ESX on which the Oracle instance is running:

- The name of IP address of the vCenter server that hosts the Oracle instance virtual machine, or the name of its parent ESX server
- The port number used by the vCenter server system or by ESX server for listening to the connections from the vSphere Client (default: 43).
- The name and password of the user that has the privileges required for connecting to the vCenter server or ESX server and retrieving information.

10 Click Monitor.

IMPORTANT: If the monitoring verification fails click the message that is displayed on the Status column and resolve the issue according to the instructions that appear in the dialog box. For example, (for example: insufficient privileges, incorrect credentials or an Agent Manager that reached its full monitoring capacity).

11 When the installation completes successfully, the Monitoring Initialized Successfully dialog box appears. Click Add another Database or Finish to exit.

Discovering Multiple Instances to Monitor

To discover additional instances to monitor:

- 1 In the upper left corner of the Databases View, click Monitor > Oracle.
The Monitor Oracle Instance dialog box appears.

- 2 Click the [Use this option to discover multiple Oracle instances link](#).

The Select an Agent Manager Dialog box appears.

- 3 Choose an Agent Manager host validate and click Next.

IMPORTANT: At the bottom of the pane you can select a default agent manager for all Future installations.

- 4 Click Add instances and select one of the following methods for adding instances:

- Load a TNS file. See [Loading a TNS file](#) on page 16
- Add manually. See [Adding CRS Agents for Monitoring Clusterware](#) on page 16

- 5 Select the check boxes beside the instances whose credentials are to be configured.

- 6 Click Set credentials on the toolbar to provide the login credentials of the instances

- 7 Follow the steps described in [Installing and Monitoring a Single Oracle Instance](#) on page 13, beginning with [Step 5](#) to enter the monitoring credentials and enable the monitoring extensions for this instance.

NOTE: To use stored credentials when enabling the Operating System monitoring extension, click the Operating System > Select from stored credentials link to open the Stored Credentials dialog box. Here you can review the login credentials and authentication methods used for logging in to Foglight.

Foglight stores encrypted credentials in lockboxes, which may be password-protected for added security. Database agents store all user login credentials in a default lockbox called DB-Agent Lockbox. If credentials have already been entered in another lockbox, use the Lockbox list to select from that lockbox.

- 8 Click Validate.

If during the connectivity validation process several RACs with a single active node are detected, a dialog box appears asking you to identify the RACs that are not RAC One Node.

- 9 Select the check boxes near the multi-node RACs and click OK. After the validation process completes, all instances that passed the validation process successfully are displayed in the table with the status Validated.

- 10 Select the validated instances and click Next.

- 11 Use the Monitor Oracle Instances screen to review the table that displays all instances that passed the connectivity validation successfully.

- To add other instances or modify the connection details of existing instances, click Previous.
- To confirm the selection and start monitoring the instances, click Next.

After clicking Next, the table displays all instances that were configured for monitoring.

At the end of this process, the Status column of the instance table displays either the status Monitored, for the instances that connected successfully to the database or a status that indicates failure of the connectivity verification process and the reason for the failure. Click this text to view a dialog box that allows changing the credentials or to grant privileges, depending on the message that appears on the Status column:

- If the message is *Insufficient database privileges*, this issue can be resolved using the Insufficient Database Privileges dialog box that appears. For details, go to [Step 12](#).
- For all other messages, the Database Connection Failed dialog box appears. For details, go to [Step 13](#).

- 12 Click Insufficient Database Privileges.

The Insufficient Database Privileges dialog box appears.

This dialog box allows you to specify a SYSAdmin (System Administrator) user with sufficient privileges.

Enter a SYSDBA user and password, and then click Grant Privileges. Alternatively, click the View script link, to the right of the Grant privileges button, to grant privileges manually by means of a script.

13 Click the message text.

A dialog box that shows the message details appears.

- 1 Click Show details to view the error's description.
- 2 Click Set credentials to display the dialog box used for entering the credentials, and enter the required changes.

Loading a TNS file

To load a TNS file:

1 Click Load TNS file.


The Select a file dialog box opens.

- 2 Click Choose file and browse for the requested file.
- 3 Click Open.
- 4 Click OK.

The Oracle Instances dialog box appears, displaying the Oracle instances that were discovered in the provided tnsnames file.

Adding CRS Agents for Monitoring Clusterware

Receiving alarms about the CRS software, which is used for monitoring Oracle Clusterware, requires creating CRS agents.

 **NOTE:** The Monitor More Clusterwares wizard supports the monitoring of a single CRS cluster. Configuring multiple CRS clusters for monitoring requires re-running the wizard for each cluster.

To add agents for monitoring Clusterware:

1 Click the Oracle tile.

2 Go to the CRS tab.

3 Click Monitor more Clusterwares at the upper left corner above of the Databases table.

The Monitor More Clusterwares wizard opens.


4 In the Configure CRS Agent screen, use the CRS agent name field to set a user-defined name for the agent. Alternatively, accept the default name provided for the CRS agent.

5 Select the requested platform for the CRS nodes: either UNIX or Windows.

6 Click Next.

7 Use the Configure Nodes for Monitoring screen to add the requested CRS nodes, using either of the following methods:

- Adding the nodes manually – go to [Step 14](#) on page 17.

 **IMPORTANT:** Adding the nodes manually requires specifying the CRS_HOME.

- Discovering the nodes – proceed to [Step 8](#).

8 Click Add Nodes, and from the menu select Discover by cluster node.

The Discover Nodes dialog box opens.

9 Use the Node Details section to enter the following details:

- Node name – the name of one of the nodes, through which all other nodes in the cluster are to be discovered.

- CRS_HOME (optional) – the path to the CRS Home directory.

If this field is left empty, the installer automatically detects the path (recommended).

10 Use the OS Credentials section to enter the OS credentials of the requested user.

IMPORTANT: Ensure providing a user that has permissions to execute CRS utilities.

- c To select existing credentials using the Foglight lockbox mechanism, click Select from stored credentials...

The Stored Credentials dialog box that appears now allows reviewing the login credentials and authentication methods used for logging in to Foglight.

- d Foglight stores encrypted credentials in lockboxes, which may be password-protected for added security. The default lockbox is DB-Agent Lockbox. If credentials have already been entered in another lockbox, use the Lockbox list to select from that lockbox

11 Fill in the other details (optional) and click Apply.

12 The wizard tries to connect to the node specified, and through this node to all other nodes that take part in the CRS cluster. All nodes are validated for OS connectivity and the permissions of the provided user are verified.


After this process completes, nodes that were successfully validated appear in the table with the status Connection verified, while other nodes display the relevant error message (for example: Failed to establish connection to host).

13 Click the error message to view the error. To handle the issues indicated in the message, click Set credentials, enter the correct details and click Apply, then proceed to [Step 19](#) on page 18.

IMPORTANT: If the error results from specifying the wrong host, click Remove and repeat the node addition process.

14 Start the process of manually adding cluster nodes for monitoring by entering the CRS_HOME of the requested cluster.

Click the CRS_HOME link at the bottom of the screen.



The Edit CRS Home dialog box opens.

15 Enter the path of the requested CRS_HOME, then click OK.

16 Click Add nodes.

17 Select the option Add manually.

The Add CRS Nodes dialog box opens. This dialog box is similar to the Discover Nodes dialog box, with the exception of the CRS_HOME parameter that has already been configured.

18 Follow [Step 9](#) to [Step 13](#), then proceed to [Step 19](#).

19 Click Next.

The agent is now created, based only on nodes that have been validated.

After the agent's creation completes, monitoring starts automatically, and the *Summary* screen appears, displaying the newly created agent with all its components.

20 Click Finish to return to the Databases dashboard.

The CRS tab now displays data about all newly monitored nodes.

Modifying the Clusterware Monitoring Settings

After the CRS agents have been configured for monitoring, their monitoring settings can be modified using the Agent Settings button.

IMPORTANT: Unlike Foglight for Oracle instances, the agent properties of CRS agents cannot be modified using the Databases Administration dashboard, but only through the default Foglight tools (see *Foglight Administration and Configuration Guide*). Clicking Agent settings and selecting the Administration option only allows re-running the wizard used for configuring the instances for monitoring, and changing the options defined using this wizard.

To modify the CRS monitoring settings:

- 1 Select the requested CRS agents.
- 2 Click Agent settings.
- 3 Select Administration to re-run the wizard used for configuring the agents for monitoring, and change the options defined using this wizard.
- 4 Select the requested option for changing the agent's monitoring:
 - Start monitoring
 - Stop monitoring
 - Remove from monitoring

Adding Agents for Monitoring Exadata Environments

To add agents for monitoring Exadata environments:

- 1 Click the Oracle tile.
- 2 Go to the Exadata tab.
- 3 Click Monitor more Exadata environments at the upper left corner above of the Databases table.

The Monitor More Exadata Environments wizard opens.

On the first screen of this wizard, *Monitor Exadata*, use the Exadata agent name field to set a user-defined name for the agent. Alternatively, accept the default name provided for the Exadata agent.

- 4 Click Next.

Use the next screen, *Monitor InfiniBand Switches*, to configure the InfiniBand switches for monitoring, either by adding the switches manually or by discovering by switch name. Configuring the switches for monitoring allows retrieving InfiniBand topology information, as well as automatically detecting the Exadata cells.

IMPORTANT: This screen is optional; if you prefer not to monitor the InfiniBand switches, you can skip this stage by clicking Next and proceeding directly to configure the Exadata cells for monitoring. In such a case, go to step 13.

- 5 Click Add switches.
- 6 From the menu that appears now, choose the requested option:

- Add manually – go to [Step 7](#) on page 19.
- Discover by switch name – go to [Step 9](#) on page 19.

7 Select the option Add manually.

The Add InfiniBand Switch dialog box opens.

Use this dialog box to add all requested InfiniBand switches that are monitored by the same credentials (user name and password), by entering a comma-separated list of all the names of the hosts on which these switches reside.


To add switches that would be monitored by a different credentials set, click **Apply** and repeat the process with the other credentials.

To edit the credentials of an existing switch, use either of the following methods:

- For a single switch, click **Set credentials** on the switch's Status row.
- For multiple switches, select the check boxes near the requested switch and click **Set credentials**.

The Set credentials dialog box opens. This dialog box is identical to the Add InfiniBand Switch dialog box, except for the host name that cannot be modified after it was set.

To edit a single switch name, click the Switch name column. In the Edit Switch Name dialog box that appears, use the Switch name field to enter a new name for the switch.

 **IMPORTANT:** The switch name specified must be taken from the list displayed by running the `ibswitches` command.

8 After adding all requested switches, click **Next** to proceed to [Step 12](#).

9 Select the option Discover by switch name.

The Discover InfiniBand Switches and Cells dialog box opens.

10 Use this dialog box to enter the host name and login credentials of a single InfiniBand switch, to be used for the discovery of all switches and cells.

11 After adding all requested details, click **Apply** and then click **Next** to proceed to [Step 12](#).

12 The following screen, *Monitor Exadata Cells*, contains a table that lists all of the Exadata Cells that were automatically discovered when the InfiniBand switches were configured for monitoring. If the stage of configuring InfiniBand switches was skipped, this table is empty and cells need to be added manually.

13 Click **Add**.

The Add Cells dialog box opens.

Use this dialog box to add all of the requested Exadata cells that are monitored by the same credentials (user name and password), by entering a comma-separated list of all the names of the hosts on which these cells reside.

To add cells that would be monitored by a different credentials set, click **Apply** and repeat the process with the other credentials.

To edit the credentials of an existing cell, use either of the following methods:

- For a single cell, click anywhere on the switch's Status row.
- For multiple cells, select the check boxes near the requested cells and click **Set credentials**.

The Set credentials dialog box opens. This dialog box is identical to the Add Cells dialog box, except for the host name that cannot be modified after it was set.

14 Click **Apply**.

After the connectivity and credentials were successfully validated, all cells appear with the status **Connection Verified**.

15 Click **Next**.

The agent is now created, based only on switches and cells that have been validated.

After the agent's creation completes, monitoring starts automatically, and the *Summary* screen appears, displaying the newly created agent with all its components.

- 16 Click Finish to return to the Databases dashboard.

The Exadata tab now displays all newly monitored cells, along with partial monitoring data. To view in-depth data about the Exadata environment, click the Home page button on the Exadata tab to access the Exadata drilldown. For details, see [Reviewing Exadata-related Information](#) on page 80.

Modifying the Exadata Monitoring Settings

After the Exadata agents have been configured for monitoring, their monitoring settings can be modified using the Agent Settings button.

- IMPORTANT:** Unlike Foglight for Oracle instances, the agent properties of ASM instances cannot be modified using the Databases Administration dashboard, but only through the default Foglight tools (see *Foglight Administration and Configuration Guide*). Clicking Agent settings and selecting the Administration option only allows re-running the wizard used for configuring the agents for monitoring, and changing the options defined using this wizard.

To modify the Exadata monitoring settings:

- 1 Select the requested Exadata switch.
- 2 Click Agent settings.
- 3 Select Administration to re-run the wizard used for configuring the agents for monitoring, and change the options defined using this wizard.
- 4 Select the requested option for changing the switch's monitoring:
 - Start monitoring
 - Stop monitoring
 - Remove from monitoring

Adding ASM Instances

The process of adding ASM instances is almost identical to the process of adding Oracle instances as described in:

- [Installing and Monitoring a Single Oracle Instance](#) on page 13.
- [Discovering Multiple Instances to Monitor](#) on page 14.

To add ASM instances go to the ASM tab. Click Monitor more ASM instances at the upper left corner of the Databases table.


If the installer detected that the ASM instance is part of an ASM cluster, a dialog box appears asking you to verify the connection details.

After ASM instances are configured, the ASM tab will provide access to a more detailed ASM Dashboard. For details, see [Reviewing ASM Instances](#) on page 78

Configuring Multiple Instances for Monitoring using Silent Installation

- ① **NOTE:** The instructions below refer to adding multiple Oracle agents. A similar process exists for adding multiple ASM instances. For more information, see [Configuring Multiple Instances for ASM](#) on page 23.

To add multiple agents using the silent installation through a command line interface:

- 1 In the navigation panel, under Dashboards, click Administration > Cartridges > Components for Download.
- 2 Click the Installer  icon beside DB_Oracle_CLI_Installer.
- 3 Follow your browser's prompts to save the ZIP file.
- 4 Extract the ZIP file to a directory of your choice.

The ZIP file contains the following files:

- *README.txt* file
 - *oracle_cli_installer.groovy* – a groovy script file that runs the silent installation. This file should be copied to the <FMS_HOME>/bin directory.
 - *oracle_silent_installer_input_template.csv* – a template file that should serve as the basis for inserting contents into the input CSV file. This file can be copied to any folder of your choice, provided that the path indicated by the <oracle_instances_file_name> parameter (see below) points to the selected path. For details about the contents of this file, see [Contents of the Input CSV File](#) on page 22
- 5 Copy the *oracle_cli_installer.groovy* file to the <FMS_HOME>/bin directory.
 - 6 Go to the command line and execute the command:

```
<FMS_HOME>/bin/fglcmd -srv <fms_host_name> -port <fms_url_port> -usr  
<fms_user_name> -pwd <fms_user_password> -cmd script:run -f  
oracle_cli_installer.groovy fglam_name <fglam_name>  
oracle_instances_file_name <oracle_instances_file_name> lockbox_name  
<lockbox_name> lockbox_password <lockbox_password>
```

The descriptions of the flags and parameters in this command are as follows:

- <FMS_HOME>: The Foglight Management Server installation directory.
 - <fms_host_name>: The name of the host where the Management Server is installed.
 - <fms_url_port>: The Management Server HTTP port.
 - <fms_user_name>: The user name used for connecting to the Management Server.
 - <fms_user_password>: The password of the specified user.
 - <clusters_file_name>: The full path and the name of the input CSV file.
 - <lockbox_name>: Optional – specifies the name of an existing lockbox that would be used.
 - <lockbox_password>: Optional – the selected lockbox's password.
- 7 After the installation process completes, review the CSV files that are created in the same folder where the input file resides:
 - A file with the input file's name and *_status* suffix (for example: if the input file is named *input*, this file is named *input_status*).

This file indicates the monitoring situation of all Oracle instances listed in the input CSV file. If monitoring failed for one or more of the instances, the file specifies the error message.

The `_status` file includes the name of the monitored Foglight *for Oracle* agent, the result of the monitoring validation process - `MONITORED`, `FAILED`, or `AGENT EXISTS` - and the error message, in case the agent creation failed.

IMPORTANT: If all instances listed in the input CSV file passed the monitoring validation successfully, no additional file is created.

- A file with input file's name and `_new` suffix (for example: if the input file is named `input`, this file is named `input_new`).

This file, which is only created if one or more instances, nodes or RACs failed the monitoring validation, contains all details about these instances, RACs, or nodes.

Using this CSV file, the user only needs to fix the errors and re-run the script, this time by specifying this file as the input file, that is: the original file name with suffix `_new`. In the example used here, the new file should be named `input_new.csv`.

Contents of the Input CSV File


The input CSV file contains the following fields, which are used as columns in the resulting file:

Name	Description
RAC Agent Name	The name of the RAC agent. This name must be unique per agent. <ul style="list-style-type: none"> • If the instance belongs to a RAC, this field has to contain the name of the RAC agent. • If this RAC runs multiple nodes, all instances must have the same RAC agent name. • If the instance does not belong to a RAC, leave this field empty.
Instance Type	The type of the instance to be monitored. The possible values are: <code>SINGLE</code> , <code>NODE</code> , <code>RAC ONE NODE</code> . <ul style="list-style-type: none"> • <code>SINGLE</code>: the instance does not belong to a RAC. • <code>NODE</code>: the instance belongs to a RAC that has more than one instance. • <code>RAC ONE NODE</code>: the instance belongs to a RAC with a single node.
Agent Name	The requested agent name. This name must be unique per agent. If an agent with the given name already exists, the newly provided agent will not be created.
Host	The host that is serving the Oracle instance or RAC.
Port	The TNS listener port.
Connection Method	The requested connection method. The possible values are either <code>SERVICE_NAME</code> or <code>SID</code> .
SID/Service Name	The name of the SID or service name used for establishing connection to the instance.
User name	The name of the user to be used for connecting to the instance. NOTE: The user name for monitoring must exist prior to running the silent installation.
Password	The password of the database user name.
Administrator User Name (Optional)	The user name of a user that can grant privileges to the regular database user; required only if the regular user does not have sufficient privileges.
Administrator Password (Optional)	The password of the database administrator user name mentioned above.

Name	Description
Enable OS Monitoring	Specifies whether to create an Infrastructure cartridge agent for the given host name; the possible values are <i>True</i> or <i>False</i> .
OS Platform (Optional)	The type of the OS platform where the Oracle instance runs; the possible values are either UNIX or WINDOWS.
Authentication (Optional)	OS authentication type. This field can have one of the following values: <ul style="list-style-type: none"> LOGIN_PASSWORD: user/password authentication RSA_KEY: RSA-based public key authentication DSA_KEY: DSA-based public key authentication CLIENT_LOGIN: uses the Foglight Agent Manager's login WINDOWS_ACCOUNT: Windows authentication, in the domain\user and password convention
OS User Name (Optional)	The user name required for connecting to the OS.
OS Password (Optional)	The password of the OS user specified above.
SSH Port (Optional)	The port that is used for listening to SSH connections on the monitored hosts; if the monitored host is Windows-based, leave this field empty.
Private Key File Path (Optional)	The full path of the private key file; relevant only to the following credential types: RSA_KEY, DSA_KEY.
Passphrase (Optional)	The passphrase of the OS user that is used in public key authentication; relevant only to the following credential types: RSA_KEY, DSA_KEY.
Use sudo (Optional)	Used for indicating that the OS user specified above is not an administrator user, but can run certain commands that require administrative privileges as root. The possible values are <i>TRUE</i> or <i>FALSE</i> . if the monitored host is Windows-based, leave this field empty to indicate <i>FALSE</i> .
Enable VMware Collection (Optional)	Indicates whether to monitor VMware metrics; The possible values are <i>TRUE</i> or <i>FALSE</i> .
VMware Host (Optional)	The VMware host name.
VMware Port (Optional)	The VMware port. If no value is inserted in this field, the port number to be used is the default number 443.
VMware User Name (Optional)	The user name required for connecting to the VMware.
VMware Password (Optional)	The password of the VMware user specified above.

Configuring Multiple Instances for ASM

To add multiple agents using the silent installation through a command line interface:

- 1 In the navigation panel, under Dashboards, click Administration > Cartridges > Components for Download.
- 2 Click the Installer  icon beside DB_Oracle_ASM_CLI_Installer.
- 3 Follow your browser's prompts to save the ZIP file.
- 4 Extract the ZIP file to a directory of your choice.

The ZIP file contains the following files:

- *README.txt* file
- *asm_cli_installer.groovy* – a groovy script file that runs the silent installation. This file should be copied to the <FMS_HOME>/bin directory.

- *asm_silent_installer_input_template.csv* – a template file that should serve as the basis for inserting contents into the input CSV file. This file can be copied to any folder of your choice, provided that the path indicated by the `<asm_instances_file_name>` parameter (see below) points to the selected path. For details about the contents of this file, see [Contents of the Input CSV File](#) on page 22
- 5 Copy the *asm_cli_installer.groovy* file to the `<FMS_HOME>/bin` directory.
 - 6 Go to the command line and execute the command:

```
<FMS_HOME>/bin/fglcmd -srv <fms_host_name> -port <fms_url_port> -usr
<fms_user_name> pwd <fms_user_password> -cmd script:run -f
asm_cli_installer.groovy fglam_name <fglam_name> asm_instances_file_name
<asm_instances_file_name> lockbox_name <lockbox_name> lockbox_password
<lockbox_password>
```

The descriptions of the flags and parameters in this command are as follows:

- `<FMS_HOME>`: The Foglight Management Server installation directory.
 - `<fms_host_name>`: The name of the host where the Management Server is installed.
 - `<fms_url_port>`: The Management Server HTTP port.
 - `<fms_user_name>`: The user name used for connecting to the Management Server.
 - `<fms_user_password>`: The password of the specified user.
 - `<fglam_name>`: The name of the selected Foglight Agent Manager where new agents are to be added.
 - `<clusters_file_name>`: The full path and the name of the input CSV file.
 - `<lockbox_name>`: Optional - specifies the name of an existing lockbox to be used.
 - `<lockbox_password>`: Optional - the selected lockbox's password
- 7 After the installation process completes, review the CSV files that are created in the same folder where the input file resides:

- A file with the input file's name and `_status` suffix (for example: if the input file is named *input*, this file is named *input_status*).

This file indicates the monitoring situation of all Oracle instances listed in the input CSV file. If monitoring failed for one or more of the instances, the file specifies the error message.

The `_status` file includes the name of the monitored Foglight *for Oracle* agent, the result of the monitoring validation process - *MONITORED*, *FAILED*, or *AGENT EXISTS* - and the error message, in case the agent creation failed.

① IMPORTANT: If all instances listed in the input CSV file passed the monitoring validation successfully, no additional file is created.

- A file with input file's name and `_new` suffix (for example: if the input file is named *input*, this file is named *input_new*).

This file, which is only created if one or more instances, nodes or RACs failed the monitoring validation, contains all details about these instances, RACs, or nodes.

Using this CSV file, the user only needs to fix the errors and re-run the script, this time by specifying this file as the input file, that is: the original file name with suffix `_new`. In the example used here, the new file should be named *input_new.csv*.

Contents of the Input CSV File

The input CSV file contains the following fields, which are used as columns in the resulting file:

Name	Description
RAC Agent Name	The name of the RAC agent. This name must be unique per agent. <ul style="list-style-type: none"> If the instance belongs to a RAC, this field has to contain the name of the RAC agent. If this RAC runs multiple nodes, all instances must have the same RAC agent name. If the instance does not belong to a RAC, leave this field empty.
Agent Name	The requested agent name. This name must be unique per agent. If an agent with the given name already exists, the newly provided agent will not be created.
Host	The host that is serving the ASM instance.
Port	The TNS listener port.
SID	The name of the SID used for establishing connection to the instance. NOTE: The SID name must be preceded by \\. For example: \+ASM1.
User name	The name of the user to be used for connecting to the ASM instance. Can be either ASMSNMP or SYS.
User Type	Connection privileges level to be used for monitoring the ASM instance. The possible values are either "SYSASM" or "SYSDBA".
Password	The password of the database user name.
Enable OS Monitoring	Specifies whether to create an Infrastructure Cartridge agent for the given host name; the possible values are <i>True</i> or <i>False</i> .
OS Platform (Optional)	The type of the OS platform where the Oracle instance runs; the possible values are either UNIX or WINDOWS.
Authentication (Optional)	OS authentication type. This field can have one of the following values: <ul style="list-style-type: none"> LOGIN_PASSWORD: user/password authentication RSA_KEY: RSA-based public key authentication DSA_KEY: DSA-based public key authentication CLIENT_LOGIN: uses the Foglight Agent Manager's login WINDOWS_ACCOUNT: Windows authentication, in the domain\user and password convention
OS User Name (Optional)	The user name to be used for connecting to the OS and monitoring it.
OS Password (Optional)	The password of the OS user name specified above.
SSH Port (Optional)	The port that is used for listening to incoming SSH connections on the monitored hosts; if the monitored host is Windows-based, leave this field empty.
Private Key File Path (Optional)	The full path of the private key file; relevant only to the following credential types: RSA_KEY, DSA_KEY.
Passphrase (Optional)	The passphrase of the OS user that is used in public key authentication; relevant only to the following credential types: RSA_KEY, DSA_KEY.
Use sudo (Optional)	Used for indicating that the OS user specified above is not an administrator user, but can run certain commands that require administrative privileges as root. The possible values are <i>TRUE</i> or <i>FALSE</i> . if the monitored host is Windows-based, leave this field empty to indicate <i>FALSE</i> .

Using Foglight *for Oracle*

Foglight *for Oracle* monitors the Oracle database activity by connecting to and querying the Oracle database. The agents provided monitor the Oracle database system. The dashboards included with the cartridge provide a visual representation of the status of the major components of the Oracle agents. They allow you to determine any potential bottleneck in database performance.

The Foglight *for Oracle* Instance Homepage

The Foglight *for Oracle* Instance home page provides a quick and intuitive representation of Oracle instance architecture and activity. Using this home page allows efficient, real-time diagnosis and resolution of Oracle operation and performance issues.

The home page is displayed in several modes:

- **Overview** – the Overview page provides various detailed views that are accessible using the toolbar. In addition, this page contains links to the Real-time and History home pages.
- **History** – the History page allows tracking the instance activity during the time range selected at the upper right corner of the screen.
- **Real-time** – the Real-time page allows using various indicators for monitoring the current activity of the selected instance.

The Real-time page home page comprises several panes, each of them relating to areas of the Oracle architecture through performance metrics. These panes are connected by a series of graphical flows that represent the logical process flows between the areas of Oracle architecture. The various graphical elements and metrics change color as their values move through a range of thresholds, and the data flows' values are updated in real time to illustrate the rate at which data is moving through the system.


Viewing the Databases Dashboard

NOTE: The following section describes the various components of the Databases dashboard. To learn how to customize this dashboard's display and focus on a requested instance, see [Selecting an Instance to Monitor](#) on page 28.

The Instance home page is launched by accessing the Databases dashboard.

To access the Databases dashboard:

- 1 Ensure that the navigation panel on the left is open.

To open the navigation panel, click the right-facing arrow  on the left.

- 2 On the navigation panel, click Homes > Databases.

The Databases dashboard provides an at-a-glance view of the monitored environment, with all of the currently monitored database types.

The Databases dashboard comprises the sections listed below:

- Database cartridge type tiles – each tile represents a database type (SQL Server, Oracle, Sybase, DB2, or All) and displays the number of instances for each database type, along with a breakdown according to the instance health state severity (Normal, Warning, Critical, Fatal).

When selecting the Oracle Tile the following advanced monitoring can be performed and reviewed:

- ASM – ASM-related data, which was previously retrieved through the Foglight *for Oracle* agent, is only enabled if the Foglight ASM instance is configured and running. For details about configuring and running this instance, see [Adding ASM Instances](#) on page 20. For details about the monitoring information displayed on the ASM tab, see [Reviewing ASM Instances](#) on page 78.
- CRS – used for configuring the monitoring of the Clusterware software that manages the Oracle RACs. For details about configuring Clusterware for monitoring, see [Adding CRS Agents for Monitoring Clusterware](#) on page 16.
- Exadata – used for configuring the monitoring of the Exadata environment. After the Exadata agents are added, data retrieved from Exadata is displayed on this tab, as well as on the Exadata drilldown.
 - For details about configuring the Exadata environment for monitoring, see [Adding Agents for Monitoring Exadata Environments](#) on page 18.
 - For details about the monitoring information displayed on the Exadata tab, see [Viewing Data about the Exadata Cells](#) on page 32.
 - For details about the drilldown that displays data retrieved from Exadata, see [Reviewing Exadata-related Information](#) on page 80.
- Status section – includes the following components:
 - Status summary – a color-coded bar, which provides a visual representation of the summarized health condition of all instances listed in the Database Group table.

The status summary bar provides a graphic representation of the monitored environment's current state, broken down to the number of instances and their current health state: Fatal, Critical, Warning, Normal or Unknown.
 - Database-specific health summary – when the database group All is selected, this section displays all of the currently monitored instances for each database type, divided by their health state (for example: four Oracle instances, three of which have the health state Warning and one is indicated with the health state Fatal). When a user-defined database group is selected, this section displays data only about the agents included within the selected group.

NOTE: The Status Indicators section can be used for filtering the Databases Group table to display only instances that meet a criterion set in this section. For details, see [Filtering the Display by Severity](#) on page 29.
- The Monitor button – Use this button to add instances to be monitored.
- The Configure Alarms button – Takes you directly to the Administration > Alarms page. On the Alarms page you can configure alarm settings and specify alarm sensitivity levels. Sensitivity levels control which alarms are enabled by default.
- The Settings button – Use this button to do one of the following:
 - Access the User Management settings, allowing you to restrict which instances specific users are allowed to view. This makes it easier for users to find information about only the instances they are interested in. For details, see [Assigning Instances to Users](#) on page 28.
 - After selecting one or more instances of the same database type, use this button to set options for collecting, storing, and displaying data for the selected instances, as well as configuring the connection to SQL Performance Investigator (if installed). For details, see [Administering Foglight for Oracle](#) on page 84.
- Currently selected database group table – a list of all monitored databases within the database group that is currently selected in the Databases section. For details, see [Using the Currently Selected Database Group Table](#) on page 31.

- The Open the Quick View button – provides the ability to open a quick view of the instance at the bottom of the screen.
- The Select dashboards (🔗) button – provides direct link to several drilldowns and panels, thereby saving the need to navigate to these locations through the Overview drilldown.

❗ **IMPORTANT:** If the display is filtered to show only Oracle instances the Databases table is renamed to Oracle Instances, and a new tab titled Reporting Services is shown. For details about this tab, see the following section.

Assigning Instances to Users

The User Management screen allows you to assign specific instances to users. As a result, when users view their Database dashboard, they will only see the instances which have been assigned to them. This makes it easier to find the data which is relevant to them.

The User Management screen only displays users who are not Administrative users.

To assign instances to a specific user:

- 1 On the main Database dashboard, click Settings and select Users from the menu.
The User Management pane appears.
- 2 Click Enable User Management feature. The table will display user names of non-Administrative users only, log in information, and instance information.
- 3 To assign one or more instances to a specific user, locate the user's name in the table and click Assign Instances in the right-hand column.
- 4 In the Assign Instances view, select the instances which should be assigned to this user and click the > button to add the selected instances to the Assigned Instances column.
- 5 Click Set and Close to save your choices and exit.

Selecting an Instance to Monitor

Because the Databases dashboard displays by default all of the currently monitored databases, it is advisable to customize the dashboard's view to display only the relevant instances, and then preview such instances to decide which to monitor.

To select an instance to monitor:

- 1 Display only instances relevant for your needs, using one of the following methods:
 - Filter by database type – if multiple database types are being monitored, click the database type tile that represents the requested type (in this case, Oracle).
 - Filter by severity – use the status indicators to display only instances of a specific database type that share a specific health state severity. For further details, see [Filtering the Display by Severity](#) on page 29.
 - Create user-defined groups – use the Databases area to create groups that contain only the databases that need to be monitored for a specific need. For further details, see [Creating User-defined Database Groups](#) on page 29.
 - Assign Instances to specific users – When accessing a Database dashboard, non-administrative users will view only instances which have been assigned to them. For details, see [Assigning Instances to Users](#) on page 28.
- 2 View the selected instances' severity level, using the status indicators.
- 3 Click the requested instance to view a cue card with a preview of the instance's most significant performance indicators.

- 4 Click Home Page on the cue card to launch a full-screen view of the requested instance.

IMPORTANT: Launching a full-screen view of the requested database by clicking the Home Page link can be carried out only for instances monitored in Foglight *for Oracle* mode.

Filtering the Display by Severity

The Status Indicators section can be used for filtering the Databases Group table to display only instances that meet a criterion set in this section.

To filter the display by severity:

- 1 Click a certain type of severity within a specific database type (for example: Oracle instances whose severity level is Critical).

The selected filter is shown in the table's title.


- 2 To select another filter, click the required status indicator (for example: Oracle > Normal severity). To display the entire list of databases, click Clear filters.

Creating User-defined Database Groups

Use the Databases area of the navigation pane to create, edit, and remove groups of database instances. Click Edit Group to add or remove database instances from the sub-group.

To create a database group:

- 1 Select the parent database instance group, *Database Group*, under Databases in the navigation pane.



- 2 Click the **+** button.
The Add Group dialog box opens.
- 3 Type a name for the group in the Name field.
- 4 Type a description for the group (optional).
- 5 Select an instance in the Available column, and click [>] to move the instance to the Selected column.
Alternatively, click [>>] to move all of the databases to the Selected column.
- 6 Click OK.

The group name appears in the Database Groups list.

The database instances are listed in the Databases pane.

After adding one or more user-defined groups, it is possible to add sub-groups to these groups.

To add sub-groups to a selected group:

- 1 Position the cursor on the requested group's name.

- 2 Repeat [Step 2](#) to [Step 6](#).

① **NOTE:** The Add Group dialog box displays the name of the parent group in the Parent field. Before proceeding, ensure that the sub-group is added to the requested group.

To remove a database group:

- 1 Select the group to be removed.

- 2 Click the **X** button.

The Remove Groups dialog box appears, asking approval to remove the selected group with all its sub groups.

- 3 Click Remove.

The group is removed from the Database Groups list.

To edit a database sub-group:

- 1 Select the requested group.

- 2 Click the  button.

The Edit Group dialog box opens.

- 3 To add an instance to the group, select the requested instance in the Available column and click [>] to move the selected instance to the Selected column.


- 4 To remove an instance from the sub-group, select the requested instance in the Selected column, and click [<].

The selected instance moves to the Available column.

- 5 Click OK.

Components Shared by All Foglight *for Oracle* Screens

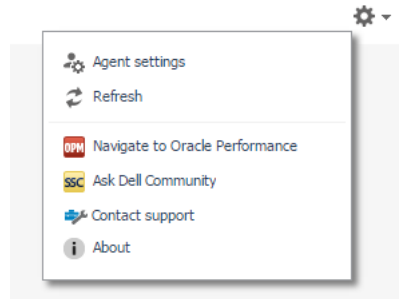
The tables' Customizer button and the In-context action button are common to all Foglight *for Oracle* screens, as described below.



- The Customizer button – use this button, which is found on the upper right corner of each table, to create a custom filter for the table. The filter is created by clicking Customizer and specifying the criteria that the various values should meet in order to be displayed in the table (having a specific name, exceeding a certain size, and so on). For detailed instructions and examples, see *Foglight User Guide > Working with Dashboards > Working with Tables*.
- In-context actions button – found on the upper right corner of all screens, and provides the following options:
 - Agent settings – when using panels whose configuration is set using the Databases Administration dashboard, clicking this button opens the relevant screen in the Databases Administration dashboard (for example, clicking this button from within the Databases > Locks panel opens the Locks screen).
 - Refresh – clicking this button retrieves the data acquired using the most recent sampling.
 - Useful links – used for carrying out further investigation using the following external sources:
 - Navigate to Oracle Performance – a central knowledge base about all of Dell Software's database monitoring products.
 - Ask Dell Community – Forum for Foglight products.

- Contact support – central knowledge base about all of Dell Software’s products.
- About – Information regard the versions being used.


Figure 1. In-context action buttons



Using the Currently Selected Database Group Table

The currently selected database group table displays the columns listed below.


Name	Description
Instance	
Sev	The highest severity alarm of the Oracle database (whether a database instance, RAC or RAC node), which determines the database’s health state.
Name	The name of the Oracle instance, RAC or RAC node. In addition, this row includes the Go to Home Page button.
Version	The version number of the Oracle database instance.
Up Since	The date and time that the Oracle database instance (or, in a RAC, the instance that was the earliest to restart) was last restarted.
Workload	The workload (average active sessions) for the Oracle database instance or RAC. When holding the cursor over the workload graph, the dwell displays a chart that shows the workload history over the specified time range.
Alarms	The number of warning, critical, and fatal alarms for the Oracle database instance/RAC. When holding the cursor over one of the alarm counts, the dwell displays the most recent alarms invoked for this database, sorted by severity. Clicking this field displays the Alarms list, which is listed by severity order. See the Foglight Online Help, Monitoring System-Wide Alarms for details on the alarm information.
System Utilization	
Host	The name of the computer that is hosting the database instance or RAC.
CPU Load (%)	The overall operating system CPU usage by all processes (including CPU usage by the database).
Memory (%)	The percentage, within the total memory, of memory consumed by all operating system processes (including the database). This value includes both RAM resident memory and swapped memory.
Disk (% Busy)	The percentage of time the busiest device spent serving system-wide input/output requests. This metric serves as a measure for the system I/O load.

Name	Description
Monitoring Status	
Agent	<p>The operational status of the monitoring agent.</p> <p>When the agent instance is running, the State icon () is green, and holding the cursor over the icon displays the message Collecting Data.</p> <p>When the agent instance is running but not collecting data, holding the cursor over the icon displays one of the following status messages listed below.</p> <ul style="list-style-type: none"> • Starting • Stopped • Stopping • Unknown <p>When SQL PI is configured the SQL PI icon is displayed.</p>
OS	<p>The status of the OS data retrieval.</p> <p>NOTE: While the Foglight <i>for Oracle</i> agent retrieves OS data, several OS metrics cannot be retrieved without the Infrastructure cartridge agent. IF OS monitoring was disabled during the creation of the Foglight <i>for Oracle</i> agent through the Oracle Monitoring Installer wizard, the wizard cannot automatically create the Infrastructure cartridge agent; However, this agent can be created manually using the Dashboards > Hosts > Hosts dashboard. For details, see <i>Adding a Monitored Host</i> in the <i>Managing the Infrastructure Cartridge User and Reference Guide</i>.</p> <p>NOTE: If the Infrastructure agent was created, but OS database is collected only by the Foglight <i>for Oracle</i> agent, click the OS status to OS data collection through the Infrastructure agent.</p>

Viewing Data about the Exadata Cells

The table displayed in the Exadata tab contains the columns listed below.

Name	Description
Exadata	
Sev	The highest severity alarm invoked within all cells monitored by the selected agent, which determines the agent's health state.
Name	The name of the Exadata agent.
Version	The version number of the Oracle Linux used for running the Exadata storage server.
Online Cells	The number of cells monitored by the selected agent that are currently online.
Temperature Status	The temperature status of all cells monitored by the selected agent, as provided by the Baseboard Management Controller (BMC); indicates whether the temperature is normal. Any deviation from the normal temperature value, which is determined by Oracle, invokes an alarm indicated using an orange color on this column.
Total Read Rate (Mb/S)	The rate (megabytes/second) at which data was read from all celldisks during the specified time range.
Total Written Rate (Mb/S)	The rate (megabytes/second) at which data was written on all celldisks during the specified time range.

Name	Description
Alarms	<p>The number of warning, critical, and fatal alarms for the celldisks.</p> <p>When holding the cursor over one of the alarm counts, the dwell displays the most recent alarms invoked for these celldisks, sorted by severity.</p> <p>Clicking this field displays the Alarms list, which is listed by severity order. See the <i>Foglight Online Help</i>, “Monitoring System-Wide Alarms” topic for details on the alarm information.</p>
System Utilization	
CPU Load (%)	The overall operating system CPU usage by all processes (including CPU usage by the database).
Memory (%)	The percentage, within the total memory, of memory consumed by all operating system processes (including the database). This value includes both RAM resident memory and swapped memory.
Monitoring Status	
	Indicator whether data is being collected.

Foglight *for Oracle* Overview Page

The Foglight *for Oracle* Overview page provides various detailed views that are accessible using the toolbar. In addition, this page contains links to the Real-time and History home pages.

The Overview page contents depends on whether a RAC (Real Application Cluster) or an Oracle database instance is selected.

Breakdown and Baseline Chart Formats

Several charts can display information in either of the following formats:

- Baseline format – where a selected individual metric is displayed as a single line.
- Breakdown format – a representation of actual activity of a metric or a set of metrics, compared with the typical behavior of these metrics for the selected time range. A breakdown display can also present a single metric divided into various components. For example, space utilization is broken down by the various components that occupy the space.

Home Page Toolbar

The Foglight *for Oracle* Overview page provides the toolbar buttons listed below.

Name	Description
Summary	<p>Opens the Summary drilldown, which contains the following panels:</p> <ul style="list-style-type: none"> • Overview page – see SQL Performance Investigator on page 35. • History Summary – see Reviewing Instance Historic Activity on page 45. • Real-time Summary – see Viewing the Foglight for Oracle Real-time Summary Page on page 48.
Performance	<p>Opens the Performance drilldown, which displays performance-related data at the following levels:</p> <ul style="list-style-type: none"> • High-level data – available for all instances. • In-depth data, delivered by means of the Cluster view or Instance view tree – available only for instances or RACs that have Performance Analysis installed and configured.
Activity	<p>Opens the Oracle Activity Drilldown, which contains the following panels:</p> <ul style="list-style-type: none"> • Usability – see Reviewing Instance Usability on page 54 and Reviewing Cluster Usability on page 55. • Locks – see Reviewing Locks Activity at the Cluster Level on page 55 and Reviewing Locks Activity at the Instance Level on page 57. • Sessions – see Viewing Detailed Sessions Data on page 58. • Overall Datafiles I/O – see Viewing Datafiles I/O Activity at the Cluster Level on page 60 and Viewing Datafiles I/O Activity at the Instance Level on page 60. • User-defined Collections – see Viewing User-Defined Collections on page 61.
Pluggable Databases	<p><i>Oracle instance 12C - pluggable only</i> – Opens the Pluggable Databases Drilldown, which contains the following components:</p> <ul style="list-style-type: none"> • Pluggable Databases List table – see Pluggable Databases List table on page 61. • Overview tab – see Overview tab on page 62. • Activity within Instance – shows the selected pluggable databases’ activity as part of the entire instance activity. For details, see Activity within Instance/Cluster tab on page 62.
Storage	<p>Opens the Storage Drilldown, which contains the following panels:</p> <ul style="list-style-type: none"> • Tablespaces and Datafiles – see Tablespaces and Datafiles Panel on page 63. • Redo Logs – see Redo Logs Panel on page 69. • Archive Logs – see Archive Logs Panel on page 70. • Invalid Objects – see Invalid Objects Panel on page 72. • Database Backup – see Database Backup Panel on page 72. <p>The Storage panels are common for both RAC and all of the instances.</p>
Configuration	<p>Opens the Configuration drilldown, which contains the panels listed below.</p> <ul style="list-style-type: none"> • SGA – see Reviewing SGA Settings at the RAC Level on page 73. • Initialization Parameters – see Reviewing the Initialization Parameters on page 75.

Name	Description
	Opens the Alerts drilldown, which contains the following panels:
Alerts	<ul style="list-style-type: none"> • Error Log – see Reviewing the Error Log on page 76. • Alert log – see Reviewing the Alert Log on page 77.
	In RACs that have Exadata configured, opens the Exadata drilldown, which contains the following panels:
Exadata	<ul style="list-style-type: none"> • Summary • Performance • InfiniBand Topology

SQL Performance Investigator

The SQL Performance page provides the ability to investigate the activity and resource consumption of a selected instance. There are two levels of metrics available:

- Without SQL PI configured – basic History and lock analysis data.
- With SQL PI configured – The ability to perform a more in-depth analysis and investigation of the Instance activity and resource consumption by adding a dimension view of the activity, change tracking analysis, execution plan analysis and compare toll. Performance Tree

The performance tree provides iterative (up to three levels) access to any of the key dimensions associated with Oracle database activity, based on the OLAP multidimensional model and an instance view of the instance activity. Domain nodes offer a hierarchical view of all types of Oracle activity characteristics.

Selecting a dimension from the tree determines what subset of activity is displayed. Iterative drill-down into domains of interest provides increasingly refined focus and diagnosis.

For example, to begin the investigation by first identifying the most active DB User, follow the steps described below:

- 1 Select the DB Users node, to display the most active database users in the selected time range. That is, the database users who consumed the highest amount of the selected resource.
- 2 Select the first user, to focus the entire window on that user's activity.
- 3 Identify the most demanding SQL statement that this specific user has executed, by expanding the user node and then selecting the SQL statement dimension node.
This displays the most active SQL statements executed by this user.
- 4 Select a specific SQL statement to focus the entire window on the selected statement's activity.
- 5 Select Client Machines under the selected SQL Statement, to view the computers on which the statement was run.

In a similar manner, such iterative drilldowns can be carried out into any Oracle dimension of interest, to gain a complete understanding of the causes of its behavior.

The default Oracle dimensions are as follows:

- PL/SQL Blocks – the blocks that comprise the PL/SQL code used
- SQL Statements – the executed SQL queries
- Programs – names of the programs that connected to Oracle and executed the SQL statements
- OS Users – Operating system users running the client program.
- DB Users – Oracle login names used for logging in to Oracle.
- Client Machines – The machines on which the client executable (connected to Oracle) is running.

- Actions – the name of the currently executing action, as set by calling the DBMS_APPLICATION_INFO.SET_ACTION procedure
- Modules – the name of the currently executing module as set by calling the DBMS_APPLICATION_INFO.SET_MODULE procedure
- Command Types – Executed SQL command type (for example, INSERT and SELECT).
- Consumer Groups – displays a collection of users with similar requirements for resource consumption
- Services – displays Oracle's internal services
- Sessions – Presents the top sessions which consumed the highest active time during the selected time frame.
- Locked Objects – Displays the objects that experienced locks, the duration of the lock and the type of the lock. The object view can be sorted by selecting a dimension in the performance tree, for example: by selecting a database name only the locked object that occurred on that database will be displayed.
- Files – displays the Data/Temp files which consumed the highest I/O Wait time during the selected time frame.
- Object I/O – displays the objects (e.g. Tables/Indexes) which were associated with the highest I/O Wait time during the selected time frame.
- Disks – displays the disks which were associated with the highest I/O Wait time during the selected time frame.

Viewing Historical Metrics

The History section view is divided into two sections that are correlated to each other:

- Resource consumption charts – This section displays data in five different charts:
 - Workload chart – Displays the instance resource activity over the selected time frame by emphasizing the resources by colors.
 - Baseline chart – Displays the instance workload compared to the baseline over time.
 - Breakdown chart – Activity of the instance by second.
 - Resource Breakdown Pie chart – Displays the resource breakdown usage by % of the total instance activity.
 - All wait events pop up – Displays details of the wait events that the instance is waiting on during the selected time range.
- Overview section- Displays a graphical representation of the metrics highlighted in the Workload related Metrics table below.
 - Workload related Metrics - A table that displays a variety of resource consumption metrics which can give an in-depth of the instance activity, each resource holds its default metrics.

Selecting each dimension in the performance tree together with a specific resource effects the data displayed for each Level.

- For example, by selecting the Lock resource the Instance view dimension will present only locks related data, the SQL Statements dimension will present only the statements that were experiencing locks and DB users the were experiencing locks and so on through all the dimensions and resources.

Blocking History

The Lock Analysis displays all locks that took place within the selected time range.

The lock analysis feature is integrated as part of the performance tree and it displays all the lock trees including further details for both the blocker and the blocked session including:

- Lock event start date

- Session Identifier - [Sid,SERIAL#]
- Session Identifier of the blocker
- Locked Object Name
- Status
- Lock Duration
- Program
- DB User
- SQL Text
- Client Machine

Viewing Change Tracking

The change tracking tool is an integrated monitoring mechanism. It periodically tracks changes in environments and activity that can potentially influence system performance and enables the user to view correlation between occurrences of changes and Oracle's activity and behavior patterns.

Use the Categories filter mechanism to refine the set of displayed change tracking occurrences. These categories are displayed on the right hand side of the pane and include:

- Execution plan – SQL statements whose execution plans have changed. Unsuitable execution plans can result in SQL performance degradation.
- Oracle configuration – instance-level configuration, such as init.ora parameters, location of files (data, log, and control), size and status
- Oracle schema – any schema object, such as tables and indexes and their partitions, clusters, constraints, views, and materialized views.
- System configuration – hardware and operating system configuration, such as OS global parameters, amount of RAM, and number of CPU units.
- User-defined – used for documenting changes inserted manually by the user, which can affect performance. Recording this activity can assist in determining whether a change in performance can be directly associated with this event

Viewing Execution Plans

This view presents the execution plan of a selected SQL and the cost of it. A Historical execution plan can be generated and will present any two views of the execution plan.

There are two ways to access the Execution Plan History dialog:

- From a Change Tracking pane by clicking on a row that displays Execution Plan change.
- From History by selecting the statement or batch and by clicking Analyze in the top of the table.

The execution plan dashboard displays the following content.

Top Bar

Resolving Date – Displays all the execution date and time of the selected statement.

Type – Displays the type of the execution plan.

Plan Hash Value – Displays the plan hash value of the execution plan

Plan analysis section

The execution plan is contains 3 tabs:

- Plan details – Displays details of each operation in the execution plan.
- Operation Analysis – Displays statistics aggregated per each operation in the execution plan.
- Object Analysis - Displays statistics aggregated per each object in the execution plan. Each object is clickable and displays a popup with important performance-related information.

By clicking the Generate Plan button the execution plan is generated and can be viewed. By clicking on the Compare Plan button a pop up will appear that allows comparing all available execution plans of the selected SQL statement.

Comparing Performance

Use the SQL PI Compare where differences in period activity illustrate underlying performance and monitoring issues. It helps you determine whether a comparison occurrence is an isolated incident or a sign of a potentially significant performance problem.

Compare can be access from the Performance tree.

The comparison can be of whole instances or selected dimension breakdowns (such as user or SQL). Use Compare to address questions such as:

- What caused a specific activity?
- What were the resource and load demands of today's instance activity compared with that of a previous day?
- Is an instance imposing different load levels now than previously?
- How do we explain the difference in a SQL performance compared with a previous period?
- What are the differences in program characteristics over two periods that caused different performance?
- After identifying the different resource usage, you can use compare to identify what caused this difference: Usage pattern? Different SQL behavior? Environmental problem?

The upper panes graphically displays the Workload or any other selected resource.

The middle Activity section displays the dimensions which were significant elements of the difference. Expanding the individual lines displays the dimension members which caused the difference and the composing metrics whose differences exceed the specified threshold.

Use the set of performance related metrics (defined by the chosen resource) to help you explain the difference in resource consumption; for example, a rise in I/O Wait might be explained by a rise in the quantity of physical reads.

Use the Comparison Parameters section in order to enter all the desired information for the cooperation. For example, time range, instance, dimension, time frames, resource.

Reviewing Overall Activity

The Oracle Overview page provides an overall view of the RAC or instance status for the selected time range, through the views and tools listed below.

- ① **NOTE:** The ASM (Automatic Storage Management) and archiving components are enabled only when being used by the monitored instances. Monitoring the ASM-related data is possible only when the ASM instance is installed and configured. For details, see [Using Foglight for Oracle](#) on page 26.

Monitoring Current Activity

The Real-time Status pane provides the indicators listed below for monitoring the current activity of either the selected instance/RAC node or the entire RAC, in Whole Cluster view.

- Sessions (no.) – displays the total number of currently running sessions, distributed to inactive sessions (that is, sessions that did not initiate an SQL call to Oracle), and active sessions.

- CPU (%) – displays the percentage of CPU that was consumed on the Oracle host by all processes (Total) or by Oracle processes only (Oracle).
- Memory (%) – displays the percentage of RAM that was consumed on the Oracle host by all processes (Total) or by Oracle processes only (Oracle)
- Up Since – identifies the date and time when the instance/RAC node was last started. If the database is not running this area reads *Database is down*, and a popup is available, displaying the database status.
 - ⓘ **IMPORTANT:** In Whole Cluster view, this indicator displays data regarding the most recently started RAC node.
- Response Time – the response time, in milliseconds, of the instance or RAC node.
 - ⓘ **IMPORTANT:** In Whole Cluster view, this indicator displays the maximum (slowest) response time among all of the cluster nodes.
- Blocked Sessions (no.) – the total number of sessions waiting for a lock held by another session.
- Avg. Virtual. Overhead – displays the percentage of CPU that is unavailable to this virtual machine because it is being utilized either by other virtual machines or by VMware itself.
 - ⓘ **IMPORTANT:** The virtualization overhead metric holds value only for virtual hosts running on VMware ESX servers, and had their VMware connection profile defined using the Databases Administration dashboard. For more information, see [Defining the VMware Connection Profile](#) on page 87.

Clicking the average virtualization overhead indicator displays the Virtualization Resource Utilization popup. This popup shows the CPU usage by both the virtual machine and the ESX host, as well as the memory sharing and swapping between all powered-on virtual machines on the host.

Monitoring Current Storage

The Real-time Storage pane enables monitoring of the current storage utilization by means of the following indicators:

- Total – displays the total amount of datafile storage available (used and free), in megabytes
 - Free – displays the amount of free storage as a percentage
 - Autoextensible – indicates if at least one datafile in the database is autoextensible:
 - No – none of the datafiles in the database are autoextensible
 - Yes – at least one datafile in the database is autoextensible
 - Archive Space – if archiving is configured, displays the following indicators of storage-related information about the most critical archive destination, that is, the archive destination with shortest time to failure:
 - Cylinder – shows the amount of used archive destination memory
 - Time to Failure (hours) – indicates the estimated time to failure in hours, based upon the rate of archiving
- ⓘ **IMPORTANT:** If the archive files reside on a disk that is managed by ASM, the Archive Space cylinder is disabled and the indication ASM agent is required appears. To retrieve archive data, the ASM instance should be installed and configured. For instructions, see [Using Foglight for Oracle](#) on page 26.

Tracking Throughput

The Real-time Throughput section enables tracking of the amount of activity running on the instance, using the speedometer charts displayed in the following sections:

- Logons – the rate per second of sessions that were initiated during the specified time range
- Executions – rate per second of statement executions during the specified time range

- Physical Data Processed – the total size (MB) of data processed through physical read and write operations during the specified time range

Monitoring Historical Activity

The History pane indicators, described in the following sections, provide information for tracking the instance activity during the time range selected at the upper right corner of the screen:

- [Workload Chart](#) on page 40
- [Top 3 busiest resources](#) on page 40

Workload Chart

The Workload chart displays the workload of the average active sessions, plotted over time.

The workload can be plotted as either:

- Baseline – workload plotted over time
- Breakdown – the amount of time the database spent waiting for various wait events, plotted over time

To display the wait events breakdown, click the arrow to the right of the Breakdown link, as shown below.

Top 3 busiest resources

The Top 3 busiest resources section displays the three resources whose wait time constituted the highest percentage of the total active time during the specified time range.

Monitoring Alarms


The Alarms pane displays color-coded indicators of the monitored host status.

The alarms are divided into the following types:



- Baseline deviations – invoked for several of the metrics for which a baseline is configured
- Thresholds – displays alarms about metrics or workloads, whose configured threshold has been exceeded


When metrics or workload resources behave normally, without any deviations, the alarm representation is displayed in gray. When a configured baseline or threshold has been exceeded, alarms are displayed in colors that represent levels of severity, ranging from yellow to red. See the Foglight Online Help for details on the alarm information.

The following indicators apply to all alarm types, including thresholds.

 **NOTE:** The numbers displayed inside the indicators show the number of alarms encountered for each alarm type.

NOTE:


Indicator	Description
 1 Yellow (Warning)	The metric (or workload resource) has exceeded the configured threshold or deviated from the baseline.
 2 Orange (Critical)	The metric (or workload resource) has significantly exceeded a threshold or deviated greatly from the baseline.

Indicator	Description
 Red (Fatal)	A severe connectivity issue was encountered, thereby triggering a fatal alarm. Such an issue can occur in one or more of the components that take place in the monitoring process; for example, if Oracle database is not responding, the listener is down, or the number of connections reached the number of sessions defined in the <i>init.ora</i> file.

Enabling or disabling baseline calculation

The Enable/Disable Baseline command under the Databases Technical Support dashboard allows enabling or disabling the baseline calculation for selected metrics.

To enable or disable baseline calculation:

- 1 Go to Dashboards > Databases > Support Dashboard > Databases Technical Support.
- 2 Click  near the Enable/Disable Baseline command.
- 3 Select the Oracle tab.
- 4 Select the requested metrics under the Metric Name column.
- 5 Select the requested command, either Enable Baseline or Disable Baseline.

Generating Reports

The Reports section can be used for generating the following types of reports:

- Reports about various aspects of the selected instance or RAC performance. These are identical to the reports generated using the Dashboards > Report command. For information on how to generate the various reports, as well as a brief description of each report, see on page 83.
- Executive summary report – a one-click report-generating feature, designed to deliver the critical parameters displayed on the Overview Page in a useful configurable format to key decision-makers. If the monitored instance is part of a multitenant architecture, this report displays a grid showing data for each pluggable database.

Reviewing RAC Historic Activity

The History Summary page in Whole Cluster view contains the following panes:

- Cluster Overview
- Interconnect/GCS
- Storage
- I/O Activity
- Resource Utilization
- Sessions

Monitoring the Cluster's Usability and Load Balance

The Cluster Overview pane displays the cluster's overall functioning and load balance.

Cluster Section

- Availability – the percentage of RAC nodes that are available (up) and running.
Clicking the number on the icon displays a popup that shows the percentage of cluster nodes that were available (up) during the specified time range, plotted over time.
Click this indicator to drill down to the Activity > Cluster Usability panel. For details, see [Reviewing Cluster Usability](#) on page 55.
- Response Time – the maximum (slowest) response time among all of the cluster nodes, in milliseconds.
Holding the cursor over the number on the icon displays a dwell that shows the average response time, that is, the average time (in milliseconds) that elapses from the moment a dummy query (either `select 1 from dual` or a user-defined query) is submitted, until the application indicates that the query was executed. This value is plotted over time.
Click this indicator to drill down to the Activity > Cluster Usability panel. For details, see [Reviewing Cluster Usability](#) on page 55.

Balance Criterion Section

The Balance Criterion section contains the following charts:

- Values aggregated to the cluster level
- Bar chart values for each cluster node

These charts display data according to the value selected from the Balance Criterion list.

The available values are:

- Workload – in average active sessions
- Availability – average availability, in percent
- Response Time – maximum response time, in milliseconds
- Logical Reads – average number of logical reads, in blocks
- CPU Load – average CPU load, in percent
- Interconnect Load – average interconnect load, in blocks
- Cluster Wait Time – average cluster wait time

Balance Section

The Balance section contains the balance bars listed below, which display the balance rate of RAC nodes for each metric.

- Workload (%) – the extent, in percent, to which workload (average active sessions) is balanced across all RAC nodes.
Ideally, all instances should equally share the wait time caused by active sessions workload; however, this metric can indicate a situation where a specific instance takes a much larger share, thereby possibly overloading the system.
- I/C Load (%) – the extent, in percent, to which the Interconnect workload-related wait events are balanced across all RAC nodes.
Ideally, all instances should equally share the wait time caused by Interconnect workload; however, this metric can indicate a situation where a specific instance takes a much larger share, thereby possibly overloading the system.
- Cluster Wait (%) – the extent, in percent, to which Cluster wait is balanced across all RAC nodes.
Ideally, all instances should equally share the wait time caused by cluster wait; however, this metric can indicate a situation where a specific instance takes a much larger share, thereby possibly overloading the system

Click the number on the icon to display a popup that shows the percentage of cluster balance between all cluster instances. This value is plotted over time.

- Logical Reads (%) – the extent, in percent, to which Logical Reads are balanced across all RAC nodes.

Ideally, all instances should equally share the wait time caused by logical reads; however, this metric can indicate a situation where a specific instance takes a much larger share, thereby possibly overloading the system.

- CPU Load (%) – the extent, in percent, to which CPU usage is balanced across all RAC nodes.

Ideally, all instances should equally share the wait time caused by CPU usage; however, an overly low value of this metric can indicate a situation where a specific instance takes a much larger share, thereby possibly overloading the system

Monitoring Interconnect, Cluster, and Global Cache

The Interconnect/GCS pane contains the indicators listed below, which are used for monitoring the global cache functioning. These indicators also demonstrate the effect of the global cache functioning over the interconnect load and the cluster latency and overhead.

- Cluster Latency – the average cluster latency value, in milliseconds.

Cluster latency is the time that has elapsed since the request for the block was submitted until the block was received by the requesting instance.

- Cluster Overhead – the percentage of cluster wait events within the overall active time. This value summarizes the cluster-related activity of all cluster instances.
- Global Cache – the total amount of global cache in megabytes. This value represents the amount of space allocated to Global Cache across all instances in the cluster.
- Cache Miss Rate – the maximum cache miss rate in percent. This value represents the frequency at which a data block is requested from a remote instance, after the user session failed to find this block in the buffer cache.
- I/C Load – the total Interconnect Load across all cluster nodes, that is, the aggregated number of requests for blocks that were not satisfied from the local instance, and therefore had to make a request for the block from the remote instance.

Monitoring Disk Storage

The Storage pane contains the components listed below, which display storage status. The information shown in this pane is aggregated to the RAC level.

Datafiles-related Indicators

- Cylinder – displays the total amount of used and free datafile storage.

Clicking the cylinder displays a popup, which shows the total space usage by the database, plotted over time.

This popup can also be used to drill down for further investigation, using the Tablespace Storage drilldown. For further details, see [Tablespaces page](#) on page 63.

- Total – displays the total amount of datafile storage available (used and free), in megabytes.
- Free – displays the amount of free RAC storage, in percent.
- Autoextensible – indicates if at least one datafile in the database is autoextensible:
 - No – all of the datafiles in the database are not autoextensible.
 - Yes – at least one datafile in the database is autoextensible.

Archive Destinations-related Indicators

If archiving is configured, the following indicators under the Archive Space (MB) section display storage-related information about the most critical archive destination, that is: the archive destination with shortest time to failure.

This section contains the following indicators:

- Cylinder – shows the amount of used archive destination memory.
- Time to Failure (hours) – indicates the estimated time to failure in hours, based upon the rate of archiving.

IMPORTANT: If the archive files reside on a disk that is managed by ASM, the Archive Space cylinder is disabled and the indication ASM agent is required appears. To retrieve archive data, the ASM instance should be installed and configured. For instructions, see “Adding ASM Instances” on page 34.

Indicators of Storage-related ASM Information

If Automatic Storage Management (ASM) is configured, and the ASM instance is installed and configured, the following indicators display storage-related information about ASM:

- Cylinder – displays the total amount of allocated and free ASM storage memory.
Clicking the cylinder displays a popup, which shows the space usage of storage managed by ASM, plotted over time.
This popup can also be used to drill down for further investigation, using the Tablespace Storage drilldown. For further details, see [Tablespaces page](#) on page 63.
- Allocated (megabytes) – the total amount of ASM memory available (allocated and free), in megabytes.
- Free (percent) – the amount of free RAC ASM storage memory, in percent.

Monitoring the Cluster’s I/O Activity

The Average I/O activity pane uses the indicators listed below to display I/O status information, aggregated to the RAC level.

- Block Read Service Time – the maximum amount of service time taken for a block read, in milliseconds.
Click the number on the icon to display a popup that shows the amount of time required for reading a single block into memory from a datafile on each instance, plotted over time.
This indicator can also be used for further investigation by means of the IO Activity for RAC drilldown.
- Physical Reads – the average number of physical reads, in blocks per second.
Click the number on the icon to display a popup that shows the rate at which data blocks are read from disk on all of the cluster nodes. This value is plotted over time.
This indicator can also be used for further investigation by means of the IO Activity for RAC drilldown.
- Physical Writes – the average number of physical writes, in blocks per second.
Click the number on the icon to display a popup that shows the rate at which modified blocks are written from the SGA to disk by the DBWR processes on all of the cluster nodes. This value is plotted over time.
This indicator can also be used for further investigation by means of the IO Activity for RAC drilldown.
- Redo Writes – the average number of redo writes, in writes per second.
Click the number on the icon to display a popup that shows the rate at which redo log entries are written to the redo log files by the LGWR processes on all of the cluster nodes. This value is plotted over time.
This indicator can also be used for further investigation by means of the Redo and Archive Performance for RAC drilldown.

The following indicators display data only if Automatic Storage Management (ASM) is configured, and the ASM instance is installed and configured (see [Using Foglight for Oracle](#) on page 26):

- ASM Reads – the total number of reads carried out by ASM disk groups.

Click the number on the icon to display a popup that shows the total reads carried out by the ASM disk groups, plotted over time.

- ASM Writes – the total number of writes carried out by ASM disk groups.

Click the number on the icon to display a popup that shows the total reads carried out by the ASM disk groups, plotted over time.

Monitoring Resource Utilization

Use the Resource Utilization pane to perform the tasks described in the following sections:

- Monitoring the CPU Load
- Monitoring the RAM

Monitoring the CPU Load

The indicators listed below are used for monitoring the CPU load.

- CPU (%) – the average CPU load (percentage), during the specified time range, of all CPU units that host the RAC nodes. This indicator displays the distribution between Oracle-incurred load and total CPU load.
- Number of CPUs – the average number of CPU units on all of the RAC nodes' hosts.

Clicking the number on the icon displays a popup that shows the total CPU usage on the cluster nodes, plotted over time.

Monitoring the RAM

The indicators listed below are used for monitoring the RAM consumption.

- Memory (%) – the average memory consumption (percentage), during the specified time range, of all CPU units that host the RAC nodes. This indicator displays the distribution between Oracle-incurred memory consumption and total memory consumption.
- Total RAM – the total amount of RAM configured for the RAC database nodes, in megabytes.
- Free RAM – displays the percentage of RAM not in use by the cluster nodes.

Monitoring Session Activity

Use the Sessions section to monitor the session activity that took place during the specified time range, by comparing between the following indicators:

- Avg. Active Sessions – the average number of active sessions that experienced resource contention during the selected time range.
- Avg. Connected Sessions – the average number of sessions that connected to the database during the selected time range.

Reviewing Instance Historic Activity

The Oracle Database instance History Summary page provides a summary view of the instance health during the specified time range, and can be used to drill down even further for carrying out root cause analysis of various performance issues.

NOTE: Archive and Data Guard data is available only if these modules have been configured. ASM data is available only if storage management is carried out by ASM and the ASM instance is installed and configured.

This dashboard contains the following sections:

- Workload
- Storage
- Cache Hit Ratio
- Data Guard
- Availability
- OS Performance

Monitoring Workload

Use the Workload pane to track the workload of the average active sessions, using the Workload chart, which displays the workload of the average active sessions, plotted over time.

The workload can be plotted as either:

- Baseline – workload plotted over time.
- Breakdown – the amount of time the database spent waiting for various wait events, plotted over time.

NOTE: The chart and the indicators below do not display all wait event types, but only the top wait events, that is, the wait events that consumed the highest amount of active time during the selected time range.

Monitoring Disk Storage

The Disk Storage panel contains the following components:

- Storage chart – displays the following indicators regarding the datafiles storage:
 - Incremental change growth – the growth of the size occupied by the database, plotted over the specified time range.
 - Database allocated size
- Avg. Consumption – the average rate (MB/hour) of disk space allocated to the database
- Total Allocated – the total amount of space allocated to the database
- ASM – indicates whether the database is monitored using automatic storage management (ASM).

IMPORTANT: Starting from version 5.6.4.32x, the ASM-related data that was previously retrieved through the Foglight for Oracle agent is only enabled if the Foglight ASM instance is configured and running. For instructions about installing the ASM instance, see [Using Foglight for Oracle](#) on page 26.

Viewing the Instance Cache Hit Ratio

The Cache Hit Ratio section displays the percentage of file read operations that were satisfied by the file system cache without requiring any physical I/O. The value of this metric should be as high as possible.

The cache hit ratio is displayed for the following cache memory types:

- PGA – the percentage of work areas executed by the system, which processed the entire input data exclusively using the PGA memory
- Buffer cache – the percentage of block requests, which found the block in the buffer cache, out of the entire block requests (including requests that required disk reads)
- Library cache – the percentage of queries that found the parsed execution plan in the library cache (soft parsing)

The information is provided by both a text and an icon, which displays real-time behavior, color-coded for severity.


Clicking the icon displays a popup that shows the cache hit ratio in a chart form, plotted over the selected time range.

This popup can also be used for further investigation, by means of the Performance drilldown. For further details, see [SQL Performance Investigator](#) on page 35.

Monitoring the Data Guard Service

The Data Guard section enables the monitoring of Oracle Data Guard on primary databases, by displaying the gap between the last archive log on the primary database and last applied archive log on the standby database.

Because Oracle Data Guard supports physical standby databases that use Redo Apply technology, use this section to monitor the rate of applying redo log files on standby database instance.

 | **NOTE:** The Data Guard section displays data only if the Data Guard service is configured.

Monitoring the Instance's Availability

The indicators included in the Instance pane provide various details about the instance.

All of the indicators display real-time behavior, color-coded for severity. Clicking each indicator displays a popup that shows the metric's value, plotted over time.

These indicators are:

- Sessions – displays a chart showing the total number of sessions, distributed between the following session types:
 - Avg. Active Sessions – the average number of active sessions that experienced resource contention during the selected time range.
 - Avg. Connected Sessions – the average number of sessions that connected to the database during the selected time range.
- Availability – the average database instance availability, in percent, that is, the percentage of time that the instance was available during the selected time range. This value is displayed as a number and a bar gauge.

Click this indicator to drill down to the Activity > Instance Usability panel. For details, see [Reviewing Instance Usability](#) on page 54.

- Average Response Time – the average response time, in milliseconds, that elapsed from the moment a dummy query (either `select 1 from dual` or a user-defined query) is submitted, until the application indicated that the query was executed. This value is plotted over time.

Click this indicator to display a popup that shows the actual response time, plotted over the specified time range. This popup can also be used to drill down to the Activity > Instance Usability panel. For details, see [Reviewing Instance Usability](#) on page 54.

Monitoring the Host Activity

The indicators displayed on the Host Activity pane provide various details about the host activity, that is, the host's CPU usage and memory consumption, affected by all of the processes rather than Oracle processes only, as well as the system I/O load.

All of the indicators display real-time behavior, color-coded for severity. Clicking each indicator displays a popup that shows the metric's value, plotted over time.

These indicators are:

- Average CPU (%) – The overall operating system CPU usage, including Oracle

- Average Memory (%) – The percentage of memory consumed by all of the instance’s OS processes (including Oracle), compared with the maximum amount of memory the instance could use. This metric includes both resident and swapped RAM.
- Disk Utilization (%) – the percentage of time the busiest device spent serving system-wide I/O requests. This metric serves as a measure for the system I/O load.

Viewing the Foglight *for Oracle* Real-time Summary Page

The following image displays the main elements of the Foglight *for Oracle* real-time summary page. The table that follows the image identifies each element, and provides a link to display more information, corresponding to the logical dataflow within Foglight *for Oracle* Instance Homepage.

The main groups (panes) of gathered icons and gauges, which represent the main activity areas in the Oracle connection process, are as follows:

- Instance identification – used for identifying the instance, its type, and its properties. See [Identifying the Instance](#) on page 48.
- Components representing instance data flow – the main activity area in the Foglight *for Oracle* Instance Homepage includes several panes and flows that represent the data flow in the Oracle operation. The dataflow is represented in a top-down design, that is, from the session to the physical disk storage.

Identifying the Instance

Use the Instance identification indicators to identify the instance, its type and its operation period.

These indicators are:

- Instance name – identifies the instance.
- Specified time range – indicates the period of time for which data is being displayed (by default: last 60 minutes).
- Instance pane – contains the following indicators:
 - DB Type – identifies whether the monitored database’s type is SQL Server, Sybase, DB2 or Oracle.
 - DB Version – identifies the Oracle version number, along with the most recently installed service packs.
 - Up Since – identifies the date and time when the instance was last started.
 - ① **NOTE:** The pane’s initial view only shows the date; to display the time as well, hover the mouse or click the date text.
 - OS Version – identifies the operating system’s version number, along with the most recently installed builds and service packs.
 - ① **NOTE:** The pane’s initial view only shows the OS name; to display the build and service packs as well, hover the mouse or click the OS name’s text.

Monitoring the Instance

The indicators included in the Instance pane provide various details about the instance. All of the indicators display real-time behavior, color-coded for severity.

- ① **NOTE:** Clicking the Response Time, CPU (%) and Memory (%) indicators displays a popup that shows the metric’s value, plotted over time.

- **Response Time** – the time (in milliseconds) that elapses from the moment a query, which is supposed to represent the general workload, is submitted, until the application indicates that the query was executed.

As the response time is usually the starting point for investigation, the Response parameter leads to the SQL Instance Summary panel in the SQL Activity drilldown,

- **CPU (%) and Memory (%)** – the average CPU load and memory consumption (percentage), during the specified time range, of all CPU units that host the Oracle instance. This indicator displays the share of Oracle-incurred CPU load and memory consumption within the total figure.

Clicking the number on both icons displays a popup that shows the total CPU usage or memory consumption on the host, plotted over time.

- **Number of CPUs** – the number of CPU units on the instance’s host.
- **Total RAM** – the total amount (in megabytes) of the host’s physical memory.
- **Free RAM** – the total amount (in megabytes) of physical memory available to the applications

Monitoring General Sessions’ Statistics

The Sessions pane monitors all session types, that is: system, user, and Oracle sessions. Use this pane to view the response time compared with the number of sessions and the instance’s level of activity. A high response time value may result from a long queue, that is: an overly high percentage of active users.

Long queues can indicate one of the following issues:

- **A massive workload** – the system handles more users than it was initially designed to do.
- **A bottleneck** – lack of system resources prevents users from carrying out their transactions, resulting in wait events and an increasing number of sessions that remain active for prolonged periods.

The Sessions pane displays the total number of sessions, distributed according to the following distinctions:

- Active vs. inactive sessions
- System (background) vs. User (foreground) sessions

The parameters in this pane lead to the Activity > Sessions panel. For details, see [Viewing Detailed Sessions Data](#) on page 58.

The Client applications represented graphically in the Sessions pane communicate with Oracle by sending and receiving network packets and by submitting SQL statements for execution by Oracle. The flows, described in the following table, help indicate performance issues if their values are too low.

Flow	Description
Executions	The rate per second of calls (both user and recursive) that executed SQL statements.
Received packets	The rate at which network packets are being received by Oracle from Client applications. When this icon’s color is yellow, clicking it displays text that describes the deviations that triggered this display, along with a representation of the rate at which Oracle is encountering network packet errors.
Sent packets	The rate at which network packets are being sent from Oracle to Client applications. When this icon’s color is yellow, clicking it displays text that describes the deviations that triggered this display, along with a representation of the rate at which Oracle is encountering network packet errors.

Monitoring the Data Guard Service

Use the Data Guard section for monitoring Oracle Data Guard on primary databases. This section displays the gap between the last archive log on the primary database and last applied archive log on the standby database.

Oracle Data Guard can be configured as either physical or logical standby. For details about the preparations required for monitoring each of these setups, see [Monitoring the Dataguard physical standby](#) on page 50 and [Working with logical standby \(SQL Apply\)](#) on page 50.

As Oracle Data Guard supports physical standby databases that use Redo Apply technology, use this section to monitor the rate of applying redo log files on standby database instance.

IMPORTANT: The Data Guard section displays data only if the Data Guard service is configured.

NOTE: Starting from Oracle 11g, Oracle provides the Active Data Guard option, which allows the standby to be read-only mode and at the same time to receive changes received from the primary node.

Monitoring the Dataguard physical standby

To prepare the physical standby:

- 1 Ensure that the Standby Database is in receiving mode.
- 2 Use sqlplus to connect to the Primary Database as SYS user.
- 3 Create the standby Foglight schema owner.
- 4 Run the privileges granting script for granting privileges to the created standby user, still from primary database. This script can be downloaded from the Connection Details screen.

If the Data Guard mode being used is Active Data Guard, proceed to [Step 6](#). Otherwise proceed to the next step.

- 5 Switch to read-only mode on the standby database.

IMPORTANT: Switching to read-only is only for the standby user availability check; there is no need to stay in that mode after the backup.

- 6 Ensure that the Foglight schema owner is available on the standby database.
- 7 Open a browser and run the Oracle Monitoring Installer wizard.
- 8 Select Add a New Instance.
- 9 Specify the standby database credentials and connection details (the user that was created in [Step 4](#)).
- 10 Click Verify Connectivity

After connectivity verification is completed successfully, start monitoring the physical standby database.

Working with logical standby (SQL Apply)

To prepare the logical standby:

- 1 Use sqlplus to connect to the Logical standby as SYS user.
- 2 Create the standby Foglight schema owner
- 3 Open a browser and run the Oracle Monitoring Installer wizard.
- 4 Continue the Oracle Monitoring Installer process as usual, using the Logical standby user credentials and connection details

After connectivity verification is completed successfully, start monitoring the logical standby database.

Monitoring the Workload

The Workload chart displays the workload of the average active sessions, plotted over time.

The workload can be plotted as either:

- Baseline – workload plotted over time.
- Breakdown – the amount of time the database spent waiting for various wait events, plotted over time.

To display the wait events breakdown, click the arrow to the right of the Breakdown link, as shown below.

Monitoring Process Activity

The Process Activity pane enables the monitoring of the processes ran by the Oracle instance, using performance indicators such as the total number of blocked processes and the rate per second of parses that are carried out during the most recent sample.

The Process Activity pane displays the parameters described below:

Parameter	Description
Recent Executions	The number of statements executed during the most recent sample.
Blocked Processes	<p>The average number of blocked processes that were collected during the specified time range.</p> <p>Blocked processes are processes that are waiting for another process to release a resource that the process is currently locking. Such processes can sometimes result in bottlenecks.</p> <p>The Blocked Process indicator changes its color when one or more processes become blocked.</p>
Parses	Total number of parse calls. This figure is the total of both hard parse calls (calls that require loading the SQL source code into RAM for parsing) and soft parse calls.
Concurrent Transactions	The number of currently running transactions.

For further details, see [Oracle Activity Drilldown](#) on page 53.

To communicate with the Oracle memory, the Oracle processes use logical reads and soft and hard parses, represented as flows in the homepage.

Monitoring Memory Activity

The Memory Activity session provides a view on the Oracle instance's memory utilization, as described in the following table.

Section	Metric	Description
SGA	Total (MB)	The actual size currently allocated to the System Global Area (SGA) memory.
	Buffer Cache (MB)	The amount of memory currently allocated to the buffer cache. The buffer cache, which is used for caching frequently accessed data blocks, is normally the largest memory cache used by Oracle.
	Shared Pool (MB)	The size of the shared pool. The shared pool caches SQL statements, PL/SQL programs, object definitions, and session memory for MTS sessions. Performance can improve when the shared pool is sized properly, thereby avoiding re-parsing and reloading.
	Buffer Cache Hit (%)	The percentage of block requests, which found the block in the buffer cache, out of the entire block requests (including requests that required disk reads).
Library Cache Hit (%)	The effectiveness of the library cache, computed by comparing the number of queries submitted to Oracle server for execution, which found the parsed execution plan in the library cache (soft parsing), with the entire number of queries submitted during the specified time range (including queries that required the parsing of a new statement, also known as “hard parsing”).	
PGA	Total (MB)	The actual size currently allocated to the Program Global Area (PGA) memory.
	PGA Hit (%)	The percentage PGA cache hit ratio. When the value of this metric is 100%, it indicates that all work areas executed by the system processed the entire input data using exclusively the PGA memory, thereby making optimal use of this memory. Any lower value indicates that at least one work area required one or more extra passes over the input data.

Monitoring Disk Storage

The Storage pane contains the components listed below, which display storage status. The information shown in this pane is aggregated to the RAC level.

Datafiles-related Indicators

- **Cylinder** – displays the total amount of used and free datafile storage.
Click the cylinder to drill down. A popup appears, displaying the total space usage by the database, plotted over time.
This popup can also be used for further investigation by means of the Tablespace Storage drilldown. For further details, see [Tablespaces page](#) on page 63.
- **Total** – displays the total amount of datafile storage available (used and free), in megabytes.
- **Free** – displays the amount of free RAC storage, in percent.
- **Autoextensible** – indicates if at least one datafile in the database is autoextensible:
 - **No** – all of the datafiles in the database are not autoextensible.
 - **Yes** – at least one datafile in the database is autoextensible.

Archive Destinations-related Indicators

If archiving is configured, the following indicators display storage-related information about the most critical archive destination, that is: the archive destination with shortest time to failure.

This section contains the following indicators:

- Cylinder – shows the amount of used archive destination memory.
- Time to Failure (hours) – indicates the estimated time to failure in hours, based upon the rate of archiving.

Indicators of Storage-related ASM Information

If storage is managed using Automatic Storage Management (ASM), and the ASM instance is installed and configured, the following indicators display storage-related information about ASM:

- Cylinder – displays the total amount of allocated and free ASM storage memory.
- Allocated (megabytes) – the total amount of ASM memory available (allocated and free), in megabytes.
- Free (percent) – the amount of free RAC ASM storage memory, in percent.

Monitoring Disk Activity

The Disk Activity section displays the following indicators:

- Unarchived – if archiving is configured, indicates the amount of Redo Logs waiting to be archived, in percent.

The popup that appears when clicking the Unarchived flow can also be used for drilling down further to investigate, using the “Redo Logs” drilldown. For further details, see [Redo Logs Panel](#) on page 69.

- ASM I/O – a textual indication of the rate of ASM memory I/O operations. If storage is managed using Automatic Storage Management (ASM), and the ASM instance is installed and configured, this field displays real-time behavior, color-coded for severity.

Tracking Physical I/O Activity

The real-time summary page contains flows and graphic images that are used for tracking physical I/O activity, as listed below.

All of the indicators display real-time behavior, color-coded for severity. Clicking each indicator displays a popup, which shows the metric’s value, plotted over time.

- Physical Writes – displays the rate per second of physical writes, that is: the rate at which modified blocks are written from the SGA to disk by the DBWR processes.
- Physical Reads – displays the rate per second at which data blocks are read from disk, when a connection requests a page that does not already exist in the buffer cache.
- Redo Writes – displays the rate per second of redo writes, that is: the rate at which redo log entries are written to the redo log files by the LGWR processes.

Oracle Activity Drilldown

Use the Oracle Activity drilldown to carry out the operations described in the following sections:

- Reviewing usability and availability – using the Usability panel, either for a single instance or a RAC. For details, see [Reviewing Instance Usability](#) on page 54 and [Reviewing Cluster Usability](#) on page 55.
- Viewing current and historical locks-related data – using the Locks panel, whose display changes when selecting a single instance or a RAC. For details, see [Reviewing Locks Activity at the Cluster Level](#) on page 55 and [Reviewing Locks Activity at the Instance Level](#) on page 57.

- Viewing information about sessions that experienced the highest contention of a specified contention criterion – using the Sessions panel. For details, see [Viewing Detailed Sessions Data](#) on page 58.
- Viewing I/O activity of all datafiles on the cluster or instance – using the panel Overall Datafiles I/O, whose display changes when selecting a single instance or a RAC. For details, see [Viewing Datafiles I/O Activity at the Cluster Level](#) on page 60 and [Viewing Datafiles I/O Activity at the Instance Level](#) on page 60.
- Viewing the user-defined collections, which were created from the Databases Administration dashboard using the User-defined Collections screen. For details, see [Viewing User-Defined Collections](#) on page 61.

Reviewing Instance Usability

The Instance Usability panel displays aggregated information regarding the instance’s usability and availability.

This panel supports the following tasks:

- Identifying instances. For details, see [The Instance Pane](#) on page 54.
- Viewing usability-related alarms. For details, see [The Alarms Pane](#) on page 54.
- Tracking the instance’s availability, plotted over the defined time range. For details, see [The Instance Availability Section](#) on page 54.
- Tracking the listener’s availability, plotted over the defined time range. For details, see [The Listener Availability Section](#) on page 54.
- Tracking the instance’s connection time, plotted over the defined time range. For details, see [The Connection Time Section](#) on page 55
- Tracking the instance’s response time, plotted over the defined time range. For details, see [The Response Time Section](#) on page 55.

The Instance Pane

The Instance pane contains the following indicators:

- DB Type – identifies monitored database’s type (for example: SQL Server or Oracle).
- Oracle Version – identifies the Oracle version number, along with the most recently installed patches.
- Up Since – identifies the date and time when the instance was last started. If the database is not running this area reads “Database is down”, and a popup is available, displaying the database status.

The Alarms Pane

The Alarms pane displays the usability-related outstanding alarms. These alarms alert the following issues per instance: response time, instance availability, instance connection availability, OS connection availability, instance connection time, and Listener status.

The Instance Availability Section

The Instance Availability section displays the instance availability in percentages, plotted over the time range. If the instance was unavailable once or more during this time, the relevant message describing the last occurrence of unavailability appears in the Error Message box.

The Listener Availability Section

The Listener Availability section displays the availability of all listeners configured for the instance, in percentages, plotted over the time range. If one or more of the listeners were unavailable during this time, a message describing this issue appears under the Listener Errors section.

The Connection Time Section

The Connection Time section displays the instance's connection time, in seconds, plotted over the time range.

The Response Time Section

The Response Time section displays the instance's response time, in milliseconds, plotted over the time range.

Reviewing Cluster Usability

Use the Cluster Usability panel to compare between the various RAC nodes, regarding the following indicators:

- Availability – measured in percentages
 - ⓘ **IMPORTANT:** The Availability section also displays the listener's availability during the specified time range.
- Connection Time – measured in milliseconds
- Response Time – measured in milliseconds

To investigate a single instance, select the requested instance from the list at the top of the panel. The display switches to instance usability. For details, see [Reviewing Instance Usability](#) on page 54.

Reviewing Locks Activity at the Cluster Level

The Activity Locks panel for a RAC comprises the following panes:

- Locks Activity
- Lock Tree (Current)
- Locks Historical Tree

Locks Activity

The Locks Activity pane displays a breakdown of all lock types and their general properties.

Name	Description
Lock Name	The name of the lock.
Gets	The total number of enqueue requests or enqueue conversions for this type of enqueue.
Gets Rate	The rate of lock gets.
Gets (percent)	The percentage of lock gets for this type of lock within the total number of locks.
Enqueue Waits	The total number of times an enqueue request or conversion resulted in a wait event.
Waits Rate	The number of wait events per second.
Wait (percent)	The percentage of wait events for this type of lock within the total number of locks.
Lock Wait	The total amount of time spent waiting for the enqueue or enqueue conversion.
Average Lock Duration	The average length of time for a lock, in milliseconds.

Name	Description
Lock Efficiency	The percentage of lock requests that did not escalate to a blocking lock.
Alarms	<p>The number of warning, critical, and fatal alarms that were invoked for this Oracle database instance.</p> <p>Holding the cursor over each alarm type displays a dwell that shows the most recent alarms invoked for this database instance, sorted by severity.</p> <p>Click each alarm type to drill down and manage the alarms. See the Foglight Online Help, <i>Monitoring System-Wide Alarms</i> for details on the alarm information.</p>

Lock Tree (Current)

The Lock Tree (Current) pane displays current data, which is available only in real-time mode.

Clicking any locking or locked session table row displays a popup that is used for viewing further details. For information see [Session Details Popup](#) on page 57.

Name	Description
PDB	In multitenant architectures, the name of the pluggable database where the lock took place.
Session	The identifier for a session that holds or acquires the lock.
Serial	The session's serial number. The combination of SID and serial number provides a unique identification of the session.
DB User	The database user name.
Object Wanted	The object requested for lock.
Lock Type	<p>The type of user or system lock.</p> <p>The locks on the user types are obtained by user applications. Any process that is blocking others is likely to be holding one of these locks. The user type locks are:</p> <ul style="list-style-type: none"> • TM – DML enqueue • TX – Transaction enqueue • UL – User supplied
Lock Mode	<p>The lock mode in which the session holds the lock. This row can have one of the following values:</p> <ul style="list-style-type: none"> • None • Null • Row-S (SS) • Row-X (SX) • Share • S/Row-X (SSX) • Exclusive
Lock Request Mode	<p>The lock mode in which the process requests the lock. This row can have one of the values listed below.</p> <ul style="list-style-type: none"> • None • Null • Row-S (SS) • Row-X (SX) • Share • S/Row-X (SSX) • Exclusive

Name	Description
Lock Duration	The duration of the lock in seconds.
Cross Instance	Defines whether the lock is between different cluster nodes: <ul style="list-style-type: none"> • YES – for global locks • NO – for single node locks
Instance	The name of the instance to which the blocking or blocked session belongs.
Kill Session	Used for terminating the selected session, by means of the Kill Session dialog box that is displayed upon clicking this column. NOTE: Killing a session requires having specific privileges. Please provide an Oracle user with these privileges. Alternatively, select the check box Log-in as SYSDBA to provide the credentials of a SYSDBA user.

Session Details Popup


Name	Description
SID	The session identifier
Serial	The session's serial number. The combination of SID and serial number provides a unique identification of the session.
Process	The process number
DB User	The database user name
OS	The operating system name
Program	The name of the program that is processing the SQL statements
Client	The Client computer identifier
Host	The host computer identifier
Current SQL	The SQL text of the SQL statement that is currently being processed in this session

Locks Historical Tree

The Locks Historical Tree pane displays the overall locks activity that took place within the selected time range, as well as locks that exceeded a predefined threshold.

This pane comprises the following sections:

- Lock Wait Events – a chart that displays all of the lock wait events, plotted over time. This chart can be viewed in either of the following views:
 - Baseline – displaying only lock wait events.
 - Breakdown – displaying both lock and non-lock wait events.
- Locks Exceeding a predefined threshold – displays locks that reached or exceeded the minimal duration, in seconds, which a lock should reach or exceed in order to be collected.

 **NOTE:** The minimal duration's length is defined using the Locks pane of the Collection Settings view.

Reviewing Locks Activity at the Instance Level

The Activity Locks panel, which displays lock summary information for an Oracle instance, comprises the following panes:

- Locks Wait Events chart

The content of this pane is similar to that of the chart displayed under the Activity > Locks > Locks Historical Tree pane at the Whole Cluster mode. For more information, see [Locks Historical Tree](#) on page 57.

- Locks Activity pane – identical to the Locks Activity pane in the Whole Cluster mode. For more information, see [Locks Activity](#) on page 55.
- Locks Tree (Current) pane – identical to the Locks Tree (Current) pane in the Whole Cluster mode. For more information, see [Lock Tree \(Current\)](#) on page 56.
- Locks Historical Tree – see [Locks Historical Tree](#) on page 58.

Locks Historical Tree


The Locks Historical Tree pane displays the locks activity and the locks tree summary, plotted over time.


This pane includes the sections listed below.

- Locks Activity – this section, which is identical to the Locks Activity section in the Locks Overview pane, displays the lock wait (average active sessions), which can be plotted as either:
 - Baseline – lock wait time plotted over time.
 - Breakdown – lock wait and non-lock wait plotted over time.
- Locks Exceeding a predefined threshold – displays locks that reached or exceeded the minimal duration, in seconds, which a lock should reach or exceed in order to be collected. The minimal duration's length is defined using the Locks pane of the Collection Settings view. .

Viewing Detailed Sessions Data

The Sessions screen allows you to trace the activity of all currently connected sessions, as well as their resource consumption.

-  **NOTE:** The refresh rate of the data displayed on the Sessions panel can be set by selecting a value from the Refresh interval drop-down list, which appears on the panel's upper right side.



This panel can also be used to carry out the tasks described in the following sections:

- [Sessions Summary](#) on page 58
- [Sessions Details](#) on page 59
- [Session Drill Down](#) on page 59

Sessions Summary

The Current section in the top of the screen allows you to view current Key Performance Indicators of the monitored instance:

- Sessions - Displays the distribution of the currently connected sessions between active (sessions that are currently running a SQL statement) and inactive sessions.
- Throughput – Divided into 4 metrics which represent the throughput of the monitored instance:
 - Logons/s – The rate per second of sessions that were initiated during the current sampling period.

- Executions/s – The rate per second of statements executions during the current sampling period.
- Commits/s – The rate per second of commits during the current sampling period.
- Rollbacks/s – The rate per second of rollbacks during the current sampling period.
- Locks – Displays the total number of blocked sessions during the current sampling period.
- Host CPU Utilization – The host CPU utilization % during the current sampling period.

The Last 1 hour trend presents the activity on the instance during the last 1 Hour by presenting:

- Workload – Displays the general workload (average active sessions) during the specified time range.
- Breakdown – Displays the workload, distributed by the various wait event categories. The colors of the categories match the colors of the various resources displayed in the Resource Breakdown section.

By clicking the Top Wait Events link at the upper right corner of the section, The Active Wait Events popup appears, with the Resource drop-down list unfiltered and displaying the entire list of wait events

Sessions Details

The current sessions table allows you to monitor the currently running sessions and their individual resource consumption. The Session List section contains a table that lists all currently running sessions, displaying each session in a separate row.

The Sessions view in the table can be filtered by the Active only and Foreground Only check boxes:

- Active only – selected by default, presenting only the active sessions that are currently running, by clearing the check box, all inactive sessions will be presented as well.
- Foreground only – selected by default, presenting only the user sessions, by clearing the check box, the background sessions will be presented as well.

Session Drill Down

The various sections of the Session Drill Down page provide detailed information about the selected session upon selecting a specific session in the table.

Session Identification Section

The left section of the screen displays several parameters that provide general information about the selected session, as follows:

- SID – Session Session identifier. NOTE: Clicking on the Session ID enables you to switch sessions from a popup list of current sessions.
- Serial# – Session serial number.
- Logon Time – The time when the user logged on to the session.
- Status – Indicates whether the session is active or inactive.
- DB User – The Oracle login name for this session.
- OS User – The OS login name for this session.
- Program – The name of the program the user is running to access Oracle
- Machine – The name of the host the session is running from.
- Waiting On – The current wait event on which the session is waiting.

Reviewing Workload Details

The Workload chart displays the selected session workload (Average Active Sessions) during the specified time range (default 1 Hr.). The breakdown displays the workload, distributed by the various wait event categories. The colors of the categories match the colors of the various resources displayed in the Resource Breakdown section.

By clicking the Top Wait Events link at the upper right corner of the section, The Active Wait Events popup appears, with the Resource drop-down list unfiltered and displaying the entire list of wait events.

Reviewing SQL Summary

This tab provides details about the SQL statements that were executed by the selected session during the selected time range (up to 1 hour).

Reviewing Sessions Statistics

Additional statistics are displayed in the "Session Statistics" tab.

Reviewing Session Blocks

The Session Blocks pane displays all blocks held or requested by the selected session.

Viewing Datafiles I/O Activity at the Cluster Level

The Overall Datafiles I/O panel displays the I/O activity of all datafiles on the cluster. This panel includes the following sections:

- Total Physical I/O chart
- Datafiles I/O Summary table

Physical I/O chart


The Physical I/O chart displays the physical I/O activity within the cluster's datafiles, in either of the following methods:

- Baseline – displays the physical I/O activity along with the allowed range of typical observed behavior. Use the chart, in its baseline form, to display either a specific metric or all of the metrics.
Metric values and typical behavior are plotted over time.
- Breakdown – provides a color-coded display of the physical activity of all instances within the RAC. Select whether to display the total physical I/O, only physical reads or only physical writes.

Click the arrow to the right of the Physical I/O title to display data for the total physical activity, or only for physical reads or writes.

Datafiles I/O Summary Table

The Datafiles I/O Summary table displays all of the I/O activity details of all datafiles on the cluster. Clicking the datafile's name displays the Physical I/O Activity page, which presents the datafile's physical read and writes, plotted over the time range, along with the overall physical reads and writes.

 **IMPORTANT:** To view the specific datafile's storage data, click the link [Datafile Storage](#), on the upper right side of the Physical I/O Activity page. For details, see [Datafile-specific Page](#) on page 68.

Viewing Datafiles I/O Activity at the Instance Level

The Overall Datafiles I/O Activity panel displays the I/O activity of all datafiles on the instance. This panel includes the following sections:

- Total Physical I/O Chart
- Datafiles I/O Summary Table

Total Physical I/O Chart

The Total physical I/O chart provides a graphic representation of the physical I/O activity within the instance's datafiles. This representation is displayed in either of the following methods:

- Baseline – displays the total physical I/O activity, including both physical reads and writes.
- Breakdown – displays the physical activity, divided by physical reads and physical writes.

Datafiles I/O Summary Table

The Datafiles I/O Summary table displays all of the I/O activity details of all datafiles on the instance. Clicking the datafile's name displays the Physical I/O Activity page, which presents the datafile's physical read and writes, plotted over the selected time range, along with the instance physical reads and writes. In multitenant architectures, the PDB column identifies the pluggable databases that use the datafiles.

IMPORTANT: To view the specific datafile's storage data, click the link [Datafile Storage](#), on the upper right side of the Physical I/O Activity page. For details, see [Datafile-specific Page](#) on page 68.

Viewing User-Defined Collections

Use the User-defined Collections panel to see customized collections, which were created using the User-defined Collections view in the Databases Administration dashboard, either during the most recent sample (Last Sample) or plotted over the specified time range (Selected Period).

IMPORTANT: The User-defined Collections panel is used only for displaying the user-defined collections; any creation or management operation of these collections is carried out using the User-defined Collections view in the Databases Administration dashboard, at the whole cluster or single instance levels only (not from within a RAC node). For details, see [Configuring User-defined Collections](#) on page 104.

This panel comprises the following sections:

- The Collections column, on the left of the panel – displays all of the names of the existing user-defined collections.
- SQL text grid – displays the text of the user-defined SQL query.
By default, this grid displays the query's short text. To display the query's full text, click [View full text](#) (a toggle).
- The collection details table – displays all of the fields contained in the query.
Each of the table's columns shows the field's display name, while each row represents the records taken at each sample.

Pluggable Databases Drilldown

The Pluggable Databases drilldown contains the Pluggable Databases pane, which displays a table of all pluggable databases (PDB), as well as the container databases (CDB), which are monitored in the instance. For details, see [Pluggable Databases List table](#) on page 61.

Selecting a row in the table displays the details of the selected container in the following panes:

- Overview – shows general details about the selected database, such as name, availability, and workload. For details, see [Overview tab](#) on page 62.
- Activity within Instance/RAC – shows the activity of the instance as part of the overall instance or RAC activity. For details, see [Activity within Instance/Cluster tab](#) on page 62.

Pluggable Databases List table

The Pluggable Databases List table includes the following columns:

- Sev – indicates the pluggable database maximum severity level
- Name – the pluggable database name
- ID – the ID number of the pluggable database
- Up Since – displays the date and time when the pluggable database was last started
- Workload (% of Instance) – the percentage the pluggable database’s workload constitutes within the entire instance/RAC
- CPU (% of Instance) – the percentage the pluggable database’s CPU consumption constitutes within the entire instance/RAC
- I/O (% of Instance) – the percentage the pluggable database’s I/O activity constitutes within the entire instance/RAC
- Shares – the portion of the system’s CPU resources that is allocated to the selected pluggable database

Overview tab

The Overview tab contains the following sections, which display general information about the pluggable database selected in the [Pluggable Databases List table](#):

- Name – the name of the pluggable database
- Open mode – the pluggable database’s state; the possible values are: mount, read-write, read-only, and migrate
- Created on – the day on which the pluggable database was created
- Cloned from – if this pluggable database was cloned, the name of the originating pluggable database
- Availability – displays the selected pluggable database’s availability within the selected time range
- Storage – displays the pluggable database’s storage compared with the storage of the entire instance or RAC
- Workload – displays the pluggable database’s workload during the specified time range
- Workload Across Cluster – if the selected pluggable database is part of a RAC, displays the balance between the various pluggable databases within the RAC

Activity within Instance/Cluster tab

The Activity within Instance tab contains the following sections, which display information about the pluggable database selected in the [Pluggable Databases List table](#):

- Workload-related Metrics – displays a list of metrics for the pluggable database
- Workload – displays the pluggable database’s workload compared with the workload of the entire instance or RAC
- Selected metric – displays the activity of the metric selected in the Metric table, compared with the workload of the entire instance or RAC

Storage Drilldown

The Storage drilldown comprises the panels described in the following sections:

- [Tablespaces and Datafiles Panel](#) on page 63
- Redo Logs – for details, see [Redo Logs Panel](#) on page 69.

- Archive Logs – for details, see [Archive Logs Panel](#) on page 70.
- **IMPORTANT:** The layout of the panels mentioned above is common for both a RAC and database instance.
- Invalid Objects – for details, see [Invalid Objects Panel](#) on page 72.
- Database Backup – for details, see [Database Backup Panel](#) on page 72.

Tablespaces and Datafiles Panel

The Tablespaces and Datafiles panel comprises the sections described in the following sections:

- Tablespaces page – displayed by default. When switching to other sections of this panel, accessible by clicking the link [View all tablespaces](#) on the upper right side of the screen. For details, see [Tablespaces page](#) on page 63.
- Datafiles grid – accessed by clicking the link [View all datafiles](#) on the upper right side of the screen. For details, see [Datafiles Grid](#) on page 64.
- Tablespace-Specific page – accessed by clicking the name of the requested tablespace on the Tablespaces Summary table. For details, see [Tablespace-specific page](#) on page 65.
- Datafile-Specific page – accessed in either of the following methods:
 - Clicking the name of the requested datafile from the [Tablespace Details > Datafiles](#) pane, under the tablespace-specific page (see [Tablespace Details Section](#) on page 67).
 - Clicking the name of the requested datafile from the Datafiles grid.

For details, see [Datafile-specific Page](#) on page 68.

Tablespaces page

The Tablespaces page displays the database individual tablespace information, sorted by either percent of space usage or size.

- **IMPORTANT:** By default, all types of tablespaces are monitored. To exclude specific types of tablespaces from monitoring, see [Tablespaces Summary Table](#) on page 64.

The Tablespaces page contains the following components:

- Top 10 Tablespaces chart – a representation of the sizes of the tablespaces that have the highest used value, displayed according to the selected sort method. For details, see [Top 10 Tablespaces Chart](#) on page 63.
- Tablespaces Summary table – a full list of the tablespaces, with general parameters and individual alarms. For details, see [Tablespaces Summary Table](#) on page 64.

Top 10 Tablespaces Chart

The Top 10 Tablespaces bar chart displays the tablespaces that have the highest used value, sorted by:

- Size
- Percentage of space usage

Name	Description
Show top	Use the Show top options to select which metric the vertical axis displays for the top 10 tablespaces: <ul style="list-style-type: none"> • Size – size in megabytes. • % Space Usage – percent.
Percent (chart vertical axis)	The percentage of the database tablespace pages currently being utilized.
Megabytes (chart vertical axis)	The database tablespace utilized size in megabytes.
Tablespaces (chart horizontal axis)	The names of the top 10 most utilized tablespaces.
Tablespace bar	The bar chart representation of the utilized tablespaces. Click a specific tablespace representation to drill down further to the Tablespace-specific page. For details, see Tablespace-specific page on page 65.

Tablespaces Summary Table

The Tablespaces Summary table shows the status for all of the database tablespaces.

To view details on a specific tablespace, click the tablespace.

Name	Description
PDB	In multitenant architectures, identifies the pluggable database that uses the tablespace shown in the Tablespace column.
Tablespace	The name of the tablespace. Click the tablespace name to drill down further to the Tablespace-specific page. For details, see Tablespace-specific page on page 65.
Total Size	The total amount of allocated space (in megabytes) in this tablespace.
% Space Usage	The percentage of used allocated space in this tablespace.
Contents	The type of tablespaces content: <ul style="list-style-type: none"> • PERMANENT • TEMPORARY • UNDO
Number of Datafiles	The number of datafiles in the selected tablespace. To view a list of all datafiles of which the tablespace consists, and possibly drill-down to a datafile-specific page, click the tablespace name to access the Tablespace Details > Datafiles pane.
Estimated Until Full	Estimated number of days until tablespace is full. If the tablespace used space did not grow during the selected period, this field shows the indication <i>no growth</i> .
Autoextensible Datafiles	The sum of the Autoextensible OS Datafiles and Autoextensible ASM Datafiles metrics.
Alarms	The number of warning, critical, and fatal alarms for this Oracle database tablespace. For further details, see Alarms on page 31.

Datafiles Grid

The Datafiles grid displays a list, sorted by usage, of all datafiles used by the database.

Name	Description
PDB	In multitenant architectures, identifies the pluggable database that uses the datafile shown in the File Name column.
File Name	The file name of the datafile, including the full directory path. Click the file name to drill down further to the Datafile-specific panel. For details, see Datafile-specific Page on page 68.
Tablespace	The name of the tablespace that contains the selected datafile.
Size	The current size of the datafile in megabytes.
Used	The current used megabytes associated to the datafile, as a percentage of datafile's allocated size.
Autoextend	Indicates whether the datafile is auto-extensible.

Tablespace-specific page

The Tablespace-specific page charts used and free space, as well as the fragmentation level of the tablespace (to determine how much space can yet be utilized). The datafiles that make up the individual tablespace appear on the bottom pane of the dashboard.

This page comprises the following sections:

- Tablespace Properties section – for details, see [Tablespace Properties Section](#) on page 65.
 - Tablespace Properties alarms – for details, see [Tablespace Properties Alarms](#) on page 65.
 - Tablespace Properties table – for details, see [Tablespace Properties Table](#) on page 66.
- Space Utilization chart – for details, see [Space Utilization chart](#) on page 66.
- Free Space Fragmentation chart – for details, see [Free Space Fragmentation chart](#) on page 66.
- Tablespace Details section – for details, see [Tablespace Details Section](#) on page 67.
 - Datafiles table – for details, see [Datafiles Table](#) on page 67.
 - All Metrics table – for details, see [All Metrics Table](#) on page 67.

Tablespace Properties Section

The tablespace properties are provided using the following panes:

- Tablespace properties alarms
- Tablespace properties table

Tablespace Properties Alarms

Name	Description
Alarms	The number of warning, critical, and fatal alarms for the selected tablespace. For further details, see Alarms on page 31.

Tablespace Properties Table

Name	Description
Status	Indicates the current status of the tablespace – Online or Offline.
PDB	In multitenant architectures, shows the pluggable database that was listed beside the selected tablespace in the Tablespaces Summary table.
Extent Management	Indicates the method used for managing extents in this tablespace: <ul style="list-style-type: none">• DICTONARY – within the Oracle data dictionary• LOCAL – within the tablespace itself
Allocation Type	The type of extent allocation to be used for this tablespace. <ul style="list-style-type: none">• USER – allocates extents as per the storage clause provided with the object or defaulted at tablespace level.• UNIFORM – allocates the uniform sizes for the extents.• SYSTEM – automatically allocates system-generated extent sizes.
Force Logging	Indicates if the tablespace is in Force Logging mode, which logs all changes to all objects in the tablespace, except changes to temporary segments. The possible values are <i>Yes</i> and <i>No</i> .
Logging	Indicates if a redo log is to be generated for the tablespace. The possible values are <i>Logging</i> and <i>No</i> .
Segment Space Management	Indicates how free and used segments are managed. <ul style="list-style-type: none">• MANUAL – using free lists• AUTOMATIC – using bitmaps
Default Table Compression	Indicates the status of default table compression; either Enabled or Disabled.
Retention	Indicates if unexpired undo blocks are retained in the tablespace. <ul style="list-style-type: none">• NOT APPLY – used for tablespaces other than the undo tablespace.• GUARANTEE – applies only to undo tablespaces.• NOGUARANTEE – applies only to undo tablespaces.
Big File	Indicates if the tablespace is a BIGFILE tablespace (Yes) or a small file tablespace (No). BIGFILE tablespaces are supported only for locally managed tablespaces with automatic segment-space management.
Block Size	The size of the memory block used by this tablespace in megabytes.
Encrypted	Indicates whether the tablespace data is encrypted (Yes) or not (No).

Space Utilization chart

Name	Description
Total	Total amount of available memory for this tablespace, plotted over time.
Used	Amount of available memory that is used for this tablespace, plotted over time.

Free Space Fragmentation chart

The Free Space Fragmentation chart displays the results of the fragmentation analysis, whose purpose is to map the fragments by the extent to which they are usable, as shown below.

- !** **IMPORTANT:** As fragmentation analysis is highly resource-consuming, it is advisable not to carry out this operation in operation environments that experience heavy resource contention.

Name	Description
Fully Usable	The total sum of fragments that are bigger than, or equal to, the maximum extent in the tablespace.
Partly Usable	The total sum of fragments that are smaller than the maximum extent in the tablespace, but bigger than, or equal to, the minimum extent in the tablespace.
Unusable	The total sum of fragments that are smaller than the minimum extent in the tablespace.

Tablespace Details Section

The tablespace details are provided using the following panes:

- Datafiles table
- All Metrics table – for details, see [All Metrics Table](#) on page 67.

Datafiles Table

Name	Description
Datafile name	The datafile name, including the full directory path. Click the datafile name to drill down further to the Datafile-specific panel. For details, see Datafile-specific Page on page 68.
Size	The current size of the datafile in megabytes.
Filesystem Free	Available space on the file system as a percentage of file system size, rounded up to one decimal. If the filesystem resides on an Automatic Storage Management (ASM) disk, retrieving this information requires having the ASM instance installed and configured.
Alarms	The number of warning, critical, and fatal alarms for this Oracle database tablespace.

All Metrics Table

This table provides the following information for tablespace metrics.

- Metric – the metric name; can have one of the values listed in the table below.
- Value – current value

Metric Name	Description
Total	The total amount of space in the selected tablespace.
Free Space	The total amount of free space in the selected tablespace.
Used	The amount of space used in the selected tablespace.
Allocated Space Used Percentage	The percentage of used space for non auto-extensible files in the selected tablespace.
Free Space Percentage	The percentage of free space.
Minimum Fragment Size	The size of the smallest fragment in the selected tablespace.
Maximum Fragment Size	The size of the largest fragment in the selected tablespace.
Average Fragment Size	The average size of fragments in the selected tablespace.

Metric Name	Description
Recycle Free (MB)	The total amount of space used by the recycle bin within the allocated free space of the tablespace.
Recycle Free (%)	The percentage of recycle bin space which occupies the allocated free space of the tablespace.
Estimated Until Full	Estimated number of days until tablespace is full. If the tablespace used space did not grow during the selected period, this field shows the indication <i>no growth</i> .
Autoextensible OS Datafiles	Number of autoextensible OS datafiles (DF).
Autoextensible ASM Datafiles	If storage is managed using Automatic Storage Management (ASM), and the ASM instance is installed and configured, displays the number of autoextensible ASM datafiles (DF).
Manual OS Datafiles	Number of manual OS datafiles (DF).
Manual ASM Datafiles	If storage is managed using Automatic Storage Management (ASM), and the ASM instance is installed and configured, displays the number of manual ASM datafiles (DF).

Datafile-specific Page

The individual Datafile properties panel displays general properties of the selected datafile, as well as a representation of the used and free space that make up the specific datafile.

The Oracle database datafile-specific panel comprises the sections listed below.

- Datafile Properties
- Space Usage – for details, see [Space Usage](#) on page 69.
- Fragmentation – for details, see [Fragmentation](#) on page 69.

Datafile Properties

Name	Description
Alarms	The number of warning, critical, and fatal alarms for the selected datafile.

Property Label	Property Description
Datafile	The name of the datafile, including the full directory path.
Tablespace	The name of the tablespace to which the datafile belongs.
PDB	In multitenant architectures, the name of the pluggable database.
Size	The size of the datafile in megabytes.
Used	The used size of the datafile in megabytes.
% Used	The percentage of the datafile used.
Status	Indicates whether the datafile is available: <ul style="list-style-type: none"> • AVAILABLE • INVALID
Autoextend	Indicates whether the datafile is auto-extensible.

Space Usage

Details are provided for:

- File Space Usage
- Filesystem Space Usage

NOTE: For datafiles managed by ASM, the filesystem space usage cylinder is available only if the ASM instance is installed and configured. For autoextensible datafiles, the file space usage cylinder is not available.

Name	Description
Used (percent)	The percentage of the datafile being used.
Used (megabytes)	The number of megabytes being used.
Total (megabytes)	The total number of megabytes available.

Fragmentation

Use the Fragmentation section to view the file fragmentation status of used and free extents, using a fragmentation analysis triggered by clicking Analyze Fragmentation.

CAUTION: Fragmentation analysis is a resource-consuming operation. Carrying out this operation is strongly discouraged in production environment with heavy resource contention.

Name	Description
Fully Usable	The amount of memory that is not being used by a datafile and is viable; represented by the blue portion of the horizontal memory bar. This bar displays the file fragmentation snapshot of used and free extents.
Partly Usable	The amount of datafile memory that is partly usable; represented by the violet portion of the horizontal memory bar.
Unusable	The amount of tablespace memory that is unusable; represented by the pink portion of the horizontal memory bar.
Used	The amount of tablespace memory that is used; represented by the orange portion of the horizontal memory bar.

Redo Logs Panel

The Redo Logs panel, which displays details regarding all redo logs, either for the Oracle instance or for all nodes of an Oracle RAC, comprises the following sections:

- Redo Logs Storage Groups table – click on a group list row to see the list of members in that group.
- Members of Group [n] table – for details, see [Members of Group \[n\] Table](#) on page 70.

Redo Logs Storage Groups Table

The redo log records all changes made in datafiles.

Name	Description
Group #	The redo log group number.
Group Status	<p>The status of the redo logs indicates which logs are in use and whether there are any odd status codes, such as stale log indications or indications of corrupt redo logs.</p> <p>The log files can have one of the following status values:</p> <ul style="list-style-type: none"> • USED – indicates either that a log had just been added (never used), or that a RESETLOGS command has been issued. • CURRENT – indicates a valid log that is in use. • ACTIVE – indicates a valid log file that is not currently in use. • CLEARING – indicates a log is being re-created as an empty log due to database administrator action. • CLEARING CURRENT – indicates that a current log is being cleared of a closed thread. If a log stays in this status, it could indicate there is some failure in the log switch. • INACTIVE – indicates that the log is no longer needed for instance recovery but may be needed for media recovery.
Archived	Indicates if the redo logs are being archived.
Sequence #	The number assigned by Oracle to each redo log file whenever a log switch occurs and LGWR begins writing to the log file. When redo log files are archived, the archived log retains its log sequence number.
Size	The size of the log in megabytes.
First Change #	A system-allocated number (SCN), which identifies the first time a change was recorded in this log.
First Time	The time when the first SCN in the log was recorded.
Alarms	The number of warning, critical, and fatal alarms for the selected Oracle database tablespace.

Members of Group [n] Table

Name	Description
Member Name	The redo log group member name.
Member Status	<p>The redo log group member status. This status can have one of the following values</p> <ul style="list-style-type: none"> • INVALID – indicates that Oracle cannot access the redo log file. • STALE – indicates a log file that is either corrupt or incomplete. If the log file's group becomes the active group, the file becomes valid again. • DELETED – indicates an online redo log file that has been dropped from the database and deleted from the disk. • IN USE – identifies the currently active redo log file.

Archive Logs Panel

The Oracle database archive logs panel comprises the following sections:

- Archive Log Destinations table

- Archive Writes chart – displays the amount (in megabytes) of disk space being used per second for the archive log.

NOTE: If the instance is configured not to archive, the Archive Logs panel is blank, displaying only the message The database is not archiving.

Archive Log Destinations Table

When monitoring a RAC, the contents of the Archive Log Destinations table vary between Whole Cluster and Instance/RAC node view, as in whole cluster view the table also includes the column Instance Name, as described below.

Name	Description
Instance Name	Indicates which of the RAC nodes uses the specified destination archive file.
NOTE: The Instance Name column is displayed only in whole cluster view. From the Archive Destination column and forwards, the two tables are identical.	
Archive Destination	The name of the destination archive file.
Destination Directory	The name of the destination directory.
Destination Status	The current status of the destination in v\$archive_dest. This destination can have one of the following values: <ul style="list-style-type: none"> • VALID – destination was initialized and is available • INACTIVE – no destination information exists • DEFERRED – destination was manually disabled by the user • ERROR – an error occurred during open or copy • DISABLED – disabled after error • BAD PARAM – parameter has errors • ALTERNATE – the destination is in an alternate state • FULL – the destination’s quota size has been exceeded
Destination Total	The total size of the archive destination target.
Destination Free Space	The total amount of free space available for the archive destination target.
Destination Usable	The total amount of usable space available for the archive destination target.
NOTE: The Destination Total, Destination Free Space, and Destination Usable columns display information only if the archive is managed using ASM and the ASM instance is installed and configured.	
Time to Failure	The amount of time before the archive log destination becomes full.
Binding	The binding setting, which describes whether successful archival is required: <ul style="list-style-type: none"> • MANDATORY • OPTIONAL
Destination Type	The type of archived log destination – either PUBLIC or PRIVATE.
Transmit Mode	The mode in which data is transmitted across the network: <ul style="list-style-type: none"> • SYNCHRONOUS • ASYNCHRONOUS, or • PARALLELSYNC
Valid Type	The Redo log type or types that are valid for the destination: <ul style="list-style-type: none"> • ONLINE_LOGFILE • STANDBY_LOGFILE • ALL_LOGFILES

Invalid Objects Panel

The Invalid Objects panel displays all of the database's schemas and the invalid objects detected in each schema, as well as their distribution per schema.

This panel comprises the following sections:

- Invalid Objects Table
- Invalid Objects Distribution chart — see [Invalid Objects Distribution for the selected schema](#) on page 72.

Invalid Objects Table

The Invalid Objects table displays each schema in a separate row, together with invalid objects of the following types:

- Functions
- Packages
- Package bodies
- Procedures
- Types
- Views
- Other type objects

In addition, the table displays the total number of invalid objects and the alarms invoked as a result of the invalid objects detected in the schema.

- ① **IMPORTANT:** By default, all of the schemas are monitored by the Invalid Objects collection. A schema with no invalid objects is displayed with a zero value in all of its columns. To exclude specific schemas from being monitored, use the Exclude List pane under the Collection Settings view in the Databases Administration dashboard. For details, see [Invalid Objects Settings](#) on page 101.

For each object type (for example, Procedures) whose metric value contains data, that is: a number greater than zero, clicking the metric value displays a popup with a detailed list of all invalid objects for the specific object type in the selected schema.

Invalid Objects Distribution for the selected schema

The invalid objects distribution chart displays the distribution of the invalid objects for the schema selected in the Invalid Objects table.

Database Backup Panel


The Database Backup panel displays various details regarding the backup of the Oracle database and datafile.

This panel contains the sections listed below.

- Last Successful Backup — the exact time and date when the most recent full backup of the datafile and database took place. A warning icon appears to the left of this section's title if the time that elapsed from the last successful backup date exceeds a predefined threshold (by default, seven days).
 - ① **IMPORTANT:** The threshold for the number of days that elapsed since the last successful full backup of a specific database can be configured by modifying the settings of the Days Since Full Backup alarm, in the Threshold section of the Alarms view in the Databases Administration dashboard. For details, see "Modifying alarm threshold values" on page 179.
- Configuration Parameters — user-defined parameters that configure the backup operation, such as the device type and the backup jobs parallelism.
- Backup Jobs — a list of the backup operations that took place during the selected time range.

A backup operation can have one of the following statuses:

- Completed
- Failed
- Completed with errors
- Completed with warnings

Operations with a Completed status are indicated with a green check mark icon () , while a warning icon appears near each operation whose status is other than Completed. Each backup job is identified by both the Start Time column, which indicates whether the backup was carried on a database or a datafile, and the Job ID column.

Clicking the row displays a popup with the details of the backup job in the selected row.

Reviewing Configuration Settings

Use the Configuration drilldown to review the current configuration settings of the Oracle agent, as well as tracking the modifications made to these settings.

This drilldown comprises the following panels:

- SGA – displays the settings and utilization of the SGA memory, using different display for the RAC and instance levels. For details, see [Reviewing SGA Settings at the RAC Level](#) on page 73 and [Reviewing SGA Settings at the Instance Level](#) on page 74.
- Initialization parameters – displayed at both the instance and the RAC levels. For details, see [Reviewing the Initialization Parameters](#) on page 75.

Reviewing SGA Settings at the RAC Level

In Whole Cluster view, use the Memory drilldown to carry out the following operations:

- Reviewing SGA Settings
- Reviewing SGA Utilization
- Reviewing the SGA's Breakdown
- Reviewing the free space reserved for the various SGA pools

Reviewing SGA Settings

Use the SGA management section to review the current SGA settings for the entire RAC, using a row that contains the following indicators:

- Automatic SGA – indicates whether the SGA memory's automatic memory management is turned on or off.
- Max SGA Size – displays the maximum amount available for the SGA memory.
- SGA Utilization – displays the average amount of SGA memory, in megabytes, which was consumed during the selected time range
- Memory Target – displays the target amount of SGA memory, in megabytes

Reviewing SGA Utilization

The SGA Utilization section contains a chart that displays SGA-related information in a breakdown format, which provides a representation of the size the following components occupy within the SGA memory:

- ASM Buffer Cache (data is available only if the ASM instance is installed and configured)
- Buffer Cache

- Java Pool
- Large Pool
- Shared Pool
- Stream Pool

Reviewing the SGA's Breakdown

Use the SGA Size section to review the breakdown of the SGA within the RAC, by means of the following components:

- A chart – providing a representation of the size each of the RAC's nodes occupies within the SGA
- A table – containing the name of the RAC or the specific node, along with columns that indicate the SGA's breakdown between the RAC's nodes, using the following metrics:
 - Fixed size
 - Variable size
 - Database buffers
 - Redo buffers

Reviewing the free space reserved for the various SGA pools

The SGA Pool Free Space section displays, both graphically using a chart and in a table format, the free space reserved for the following types of SGA pools for each of the RAC's nodes:

- Java pool
- Large pool
- Shared pool

Reviewing SGA Settings at the Instance Level

Reviewing SGA Settings

Use the SGA management section to review the current SGA settings for the entire RAC, by means of a row that contains the following indicators:

- Automatic SGA – indicates whether the SGA memory's automatic memory management is turned on or off
- Max SGA Size – displays the maximum amount available for the SGA memory.
- SGA Utilization – displays the average amount of SGA memory, in megabytes, which was consumed during the selected time range
- Memory Target – displays the target amount of SGA memory, in megabytes

Reviewing SGA Utilization

The SGA Utilization section contains a chart that displays the instance's current SGA size in either of the following views:

- Baseline – total SGA space utilization (current SGA size)
- Breakdown – displays the size the following components occupy within the instance's SGA:
 - ASM Buffer Cache (data is available only if the ASM instance is installed and configured)
 - Buffer Cache
 - Java Pool
 - Large Pool

- Shared Pool
- Stream Pool

Reviewing the SGA's Breakdown

Use the SGA Size section to review the breakdown of the SGA within the instance, by means of the following components:

- A pie chart – providing a representation of the size each component described in the table occupies within the instance's SGA
- A table – shows the SGA's breakdown between the following components:
 - Fixed size
 - Variable size
 - Database buffers
 - Redo buffers

Reviewing the Free Space Reserved for the Various SGA Pools

The Free Space section displays, both graphically using a chart and in a table format, the free space reserved for the following types of SGA pools for each of the RAC's nodes:

- Shared pool
- Java pool
- Large pool

Reviewing the Initialization Parameters

The Initialization Parameters panel displays the entire list of initialization parameters, as defined by Oracle. These parameters define the characteristics of each Oracle instance when the instance is started.

Upon clicking anywhere on a row that displays a specific initialization parameter, a popup appears, listing the properties of the selected parameter:

- The parameter description, taken from Oracle's documentation (also appears under the Description column).
- Link to perform searching for further information.
- Change History table – displays the entire history (time and dates) when the Foglight *for Oracle* agent detected the changes made to this parameter since the instance was created.

The Initialization Parameters table varies between instance mode and RAC mode. For details about the RAC mode view, see [Reviewing the Initialization Parameters in RAC Mode](#) on page 76.


In Instance mode, the table displays the follow the columns listed below.

Name	Description
Parameter	Parameter name.
PDB	In multitenant architectures, the name of the pluggable database.
PDB is Modifiable	<i>Hidden by default.</i> In multitenant architectures, displays the list of pluggable databases that have been modified.
Current Value	The parameter's current value.
Default	Indicates whether the parameter's current value is its default value, as defined by Oracle.


Name	Description
Dynamic	Indicates whether the initialization parameter is dynamic, that is, can be modified using the ALTER SESSION, ALTER SYSTEM, or ALTER SYSTEM DEFERRED commands while an instance is running.
Modified	Indicates whether the parameter's value has been modified.
Deprecated	Indicates whether the parameter is deprecated, in which case specifying it in a parameter file will invoke a warning message when the instance starts.
Type	The parameter type, as defined by Oracle.
Description	The parameter description, taken from Oracle site.
PDB_modifiable	Indicates whether the pluggable database can be modified. NOTE: This column is not displayed by default.

Reviewing the Initialization Parameters in RAC Mode

In RAC mode, the Initialization Parameters panel can be accessed by selecting either the Whole Cluster or a specific instance view.

In the Whole Cluster view, instance-specific initialization parameters are displayed with the  icon to the left of the parameter name, and the names of the various instances are displayed in the Scope column.

The following columns in Initialization Parameters table differ in RAC mode:


Name	Description
Parameter	Parameter name. A  icon to the left of the name indicates that the parameter applies differently to specific instances.
Scope	Indicates whether the parameter applies to the entire RAC or varies between specific instances.
Current Value	The parameter's current value. This value can be either instance-specific or apply to the entire RAC.

Reviewing the Alert and Error Logs

The Alerts drilldown provides a view on the alert and error logs received for a selected Oracle agent during the specified time range.

This drilldown comprises the following panels:

- Error Log – see [Reviewing the Error Log](#) on page 76
- Alert Log – see [Reviewing the Alert Log](#) on page 77

 **NOTE:** Both the Error Log and Alert Log panels are displayed only at the instance level. Selecting the Error Log or Alert Log menu item from within a RAC, when the Whole Cluster view is selected, displays a message that prompts the selection of the requested instance.

Reviewing the Error Log

The Error Log panel shows error messages that were selected to be extracted and displayed from the Foglight agent error logs. The error messages appear with their exact date and time, as well as severity level.

Only a small number of the error log messages are displayed in this panel. The messages to be displayed, as well as their severity, are determined within the agent and cannot be modified by the user.

The type of agents whose error logs are used for the message display is determined according to the selected view, as listed below.

- DB_Oracle – if the selected view is a single instance.
- DB_Oracle_RAC_Instance – if the selected view is a RAC node.

Reviewing the Alert Log

By default, the Alert Log panel displays all of the alert log messages.

To display messages under a pre-defined name, category and severity, create or edit patterns for these messages in the Match List, which is found in the Alert Log view in the Databases Administration dashboard, and add the corresponding regular expressions to this list. To prevent messages from being displayed on the Alert Log panel, add these messages to the Ignore List.

- IMPORTANT:** Only messages that are explicitly defined in the Ignore List will not be displayed. Messages that were not added to either the Match or Ignore lists appear with the name Other, the type of Oracle alert messages, and the severity Information. Therefore, ensure that messages that do not need to be displayed are added to the Ignore List.

The Alert Log panel comprises the panes described in the following sections:

- [Alert Log Messages Chart](#) on page 77
- [Alert Log Messages Table](#) on page 77

Alert Log Messages Chart

The Alert Log Messages chart provides a graphic, color-coded representation of the various alert log messages, broken down by their severity level, as described below.

- Informational – displayed on the screen, but does not invoke alarms
- Warning
- Critical
- Fatal

The severity assigned for each alert, as well as the minimal severity level (Warning or higher) for which alarms are to be invoked as a result of the alert, are defined using the Alert Log view in the Databases Administration dashboard, accessible from the Agent Settings button. .

Alert Log Messages Table

The Alert Log Messages table provides the following details about each alert log:

- Timestamp – the exact time and date when the alert log message was created
- Severity – can have one of the values listed below
 - Informational
 - Warning
 - Critical
 - Fatal
- Message – the title of the Oracle error message
- Dump File – this field is populated by a link only if the error message also refers to a dump file. When the link is clicked, the Dump File popup appears, displaying the first 100 lines of the dump file contents.
- Name – the alert log name


- Type – the alert log type, such as Oracle errors or Oracle alert messages

NOTE: Both Name and Type refer to the parameters defined using the Alert Log view in the Databases Administration dashboard.

Reviewing ASM Instances

Starting in version 5.6.5.5, Foglight for Oracle provides a detailed dashboard for monitoring ASM performance. The ASM information is presented only after ASM instances are configured.

The ASM tab lists all ASM instances which have been configured. To access the detailed dashboard for that instance, click the home button.



The ASM dashboard provides information about the ASM cluster/instance, it includes general instance information, Disk groups and disks information, ASM Parameters and ASM Alert log information.

The ASM dashboard has three tabs:

- Summary – [Reviewing Summary Information](#) on page 78
- Parameters – [Reviewing Parameter Information](#) on page 80
- Alert log – [Reviewing Alert Log Information](#) on page 80

Reviewing Summary Information

The summary tab contains information about all the disk groups that are configured for the ASM, as well as the disks that are configured for each ASM disk group. This information includes the version number, status, total storage space, and the number of ASM clients.

The number of ASM clients is clickable and opens a pop-up that shows all the Oracle instances that are using a disk group managed by the ASM instance.

It also has a column named Monitored that gives an indication whether or not that Oracle instance is monitored by Foglight. If the Oracle instance is not monitored by Foglight, it is possible to monitor it immediately by clicking Monitor.

The Disk Groups table provides detailed information about each group. Selecting one or more disk groups will display additional information in the pane below, either in Map or Table View.

Each disk group in the top table or disk in the Map or Table view is clickable and opens a pop-up with additional information.

ISRSOL09-asm | Summary Parameters Alert Log

Version 11.1.0.6.0 Status STARTED Total Storage 200 GB ASM Clients 0

Disk Groups

Select all Select none

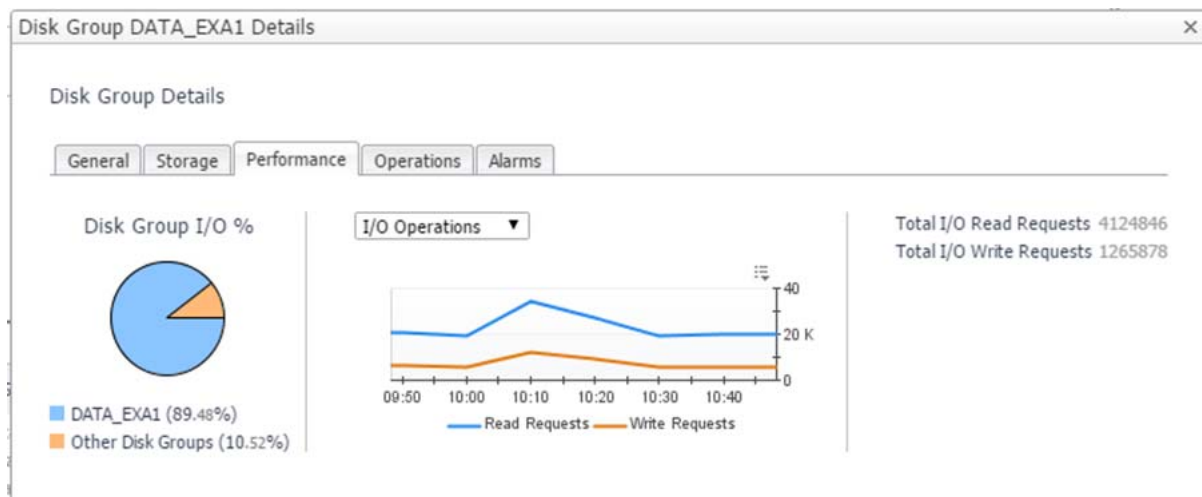
Severity	Number	Name	State	Total Size (GB)	% Space Usage	Avg Service Time (ms)	Number of Disks	Alert
✓	3	TESTDG	MOUNTED	40.00	2.6 %	2.63	1	
✓	1	DGROUP1	MOUNTED	79.98	100 %	99.65	2	
✓	2	DGROUP2	MOUNTED	80.00	100 %	99.91	2	

Views: Map | Table

Name	State	Size (GB)	% Space Usage	Path	Avg Service Time (ms)
DGROUP1					
DGROUP1_0000	NORMAL	39.98	100 %	/dev/rdisk/c113d0e0	99.67
DGROUP1_0001	NORMAL	40.00	100 %	/dev/rdisk/c113d0e1	99.62
DGROUP2					
DGROUP2_0000	NORMAL	40.00	100 %	/dev/rdisk/c113d0e3	99.91
DGROUP2_0001	NORMAL	40.00	100 %	/dev/rdisk/c113d0e4	99.91

For disk groups, the pop-up includes the following tabs:

- General – General Disk group information. For example, name, state number of disks, number of fail groups, allocation unit size.
- Storage – Disk group storage-related information
- Performance – Performance-related information
 - I/O Operations
 - Average Service Time
 - Disk Errors
- Operations – Currently running rebalance operations
- Alarms – List of alarms associated with this disk group



If you select a specific disk in the Map or Table view, the pop-up includes the following tabs:

- General – General Disk information (e.g. path, failure group, creation date)
- Storage – Disk storage-related information
- Performance – Disk performance-related information

Reviewing Parameter Information

The Parameters tab displays a list of all the ASM parameters and their values, as well as some other additional information for each parameter.

Initialization Parameters

Parameter ^	Current Value	Default	Dynamic	Modified	Type	
asm_diskgroups	FRA	<input type="radio"/>	<input checked="" type="checkbox"/>	FALSE	STRING	disk groups to mount automatically
asm_diskstring	/dev/oracleasm/disks	<input type="radio"/>	<input checked="" type="checkbox"/>	FALSE	STRING	disk set locations for discovery
asm_power_limit	8	<input type="radio"/>	<input checked="" type="checkbox"/>	SYSTEM_MOD	INTEGER	number of parallel relocations for d
asm_preferred_read_failure_groups		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	STRING	preferred read failure groups
audit_file_dest	/orav01/11.2.0.4/grid/rdbms...	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	STRING	Directory in which auditing files are
audit_sys_operations	FALSE	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	BOOLEAN	enable sys auditing
audit_syslog_level		<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	Syslog facility and level
background_core_dump	partial	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	Core Size for Background Processe
background_dump_dest	/orav01/oracle/diag/asm/+a...	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	STRING	Detached process dump directory
cluster_database	TRUE	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	BOOLEAN	if TRUE startup in cluster database
cluster_database_instances	4	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	INTEGER	number of instances to use for siz
cluster_interconnects		<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	interconnects for RAC use
core_dump_dest	/orav01/oracle/diag/asm/+a...	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	STRING	Core dump directory
cpu_count	2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	INTEGER	number of CPUs for this instance
db_cache_size	0	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	BIG INTEGER	Size of DEFAULT buffer pool for sta
db_ultra_safe	OFF	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	Sets defaults for other parameters
db_unique_name	+ASM	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	Database Unique Name
diagnostic_dest	/orav01/oracle	<input type="radio"/>	<input checked="" type="checkbox"/>	FALSE	STRING	diagnostic base directory
event		<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	debug event control - default null s
file_mapping	FALSE	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	FALSE	BOOLEAN	enable file mapping
filesystemio_options	none	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	IO operations on filesystem files
file		<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	PARAMETER FILE	include file in init.ora
instance_name	+ASM2	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	STRING	instance name supported by the in
instance_number	2	<input checked="" type="checkbox"/>	<input type="radio"/>	FALSE	INTEGER	instance number
instance_type	asm	<input type="radio"/>	<input type="radio"/>	FALSE	STRING	type of instance to be executed
large_pool_size	12582912	<input type="radio"/>	<input checked="" type="checkbox"/>	FALSE	BIG INTEGER	size in bytes of large pool

Reviewing Alert Log Information

The Alert log tab displays the entire alert log messages with the relevant severity, as defined in the Match List. Only messages that appear in the Ignore List will not be displayed in this screen.

Reviewing Exadata-related Information

Starting from version 5.6.5.330, Foglight *for Oracle* monitors Exadata database machines, both from the hardware and software aspects, using the Exadata drilldown.

The Exadata-related information is presented only in RACs where Exadata is configured. While Exadata appears as one of the drilldowns, clicking Exadata opens an Exadata dashboard.

The Exadata dashboard comprises the following drilldowns:

- Summary – see [Displaying Exadata-related Summarized Data](#) on page 80
- Performance – see [Reviewing the Exadata Performance](#) on page 82
- InfiniBand Topology – see [Reviewing the InfiniBand Topology](#) on page 83

Displaying Exadata-related Summarized Data


The Summary drilldown provides summarized data about the performance and resource utilization of the Exadata cells.

Summary


Cell Table

ID	Name	Status	Connection Availability	Up Since	Fans (Working/Total)	Power Supplies (Working/Total)	Temperature °C
exa1cel01ID01	exa1cel01	online		03/24/2013 22:21	12/12	2/2	21 °C
exa1cel02ID02	exa1cel02	online		03/24/2013 22:22	12/12	2/2	23 °C


Flash Cache Hit Ratio



Memory Utilization




CPU Busy Percent



Cell Performance

Name	Selected Cells	All Cells
Run Queue Length (Count)	0.40	0.45
InfiniBand Sent Rate (MB/s)	0.52	11.92
InfiniBand Received Rate (MB/s)	0.26	0.60

Cell Offload Efficiency



This drilldown comprises the following components:

- Cell Table – see [Cell table](#) on page 81.
- Performance and Resource Utilization section – see [Performance and resource utilization indicators](#) on page 81.

Cell table

The cell table provides the following details about each alert log:

- ID – displays the cell's ID
- Name – displays the cell's name
- Status – displays an Orange light if the cell is offline, or a green light if the cell is online
- Connect Availability – displays a green light if the cell's availability is 100%, otherwise displays a red light
- Up Since – displays the date and time when the instance was last started
- Fans (Working/Total) – displays a green light if the number of working fans equals the total number of fans, otherwise displays a red light
- Power Supplies (Working/Total) – displays a green light if the number of working units equals the total number of units, otherwise displays a red light
- Temperature – displays a green light if the temperature's status is Normal, otherwise displays a red light

Performance and resource utilization indicators

The performance and resource utilization indicators includes the following components.

Indicator	Description
Flashcache Hit Ratio (%)	<p>Displays the percentage of block requests resolved from the cell flash cache falls below a predefined threshold. This figure is displayed for the selected cells, as well as for all cells.</p> <p>Block retrieval is carried out significantly faster by accessing the flash cache, as this cache resides on high-performance flash SSD. Low flash cache hit ratio indicates that many blocks are being retrieved from the regular disks, possibly resulting in performance issues.</p>
Memory Utilization (%)	<p>Displays the memory utilization for the selected cells, as well as for all cells.</p> <p>Overly high memory utilization can result in excessive swapping, which is likely to lead to performance issues.</p>
CPU Busy Percent (%)	<p>Displays the CPU utilization for the selected cells, as well as for all cells.</p> <p>Overly high CPU utilization can result in excessive swapping, which is likely to lead to performance issues.</p>
Cell Performance	<p>Displays several key indicators for the cell's performance, as follows:</p> <ul style="list-style-type: none"> • Run Queue Length • InfiniBand MB Sent Rate • In fin i band MB Sent Rate
Cell Offload Efficiency	<p>The number of rows processed by the cells for all the queries divided by the number of rows forwarded by the cells to the DB node instances.</p> <p>Higher values of this metric indicate that the cells made good use of their CPUs and high I/O bandwidth to process high number of rows, extract the relevant ones and only forward those to the DB nodes.</p>

Reviewing the Exadata Performance

The Performance drilldown includes the following components:

- Cell table – see [Cell table](#) on page 81.
- Query efficiency statistics
- Celldisks statistics

Query efficiency statistics

Celldisk response time

The *Celldisk Response Time (ms)* section displays the following metrics on all of the celldisks, as well as on the celldisks selected in the cell table:

- Celldisks small writes latency – the average latency, in milliseconds, to write small blocks per request.
- Celldisks small reads latency – the average latency, in milliseconds, to read small blocks per request.

Flash cache statistics

Use the Flash Cache Statistics section to review the load on the flash cache and the number of I/O errors encountered there, by displaying the following metrics on all of the celldisks, as well as on the celldisks selected in the cell table:

- Flash cache I/O errors – the number of IO errors that were encountered on the FlashCache

- Data read from flashcache (MB) – the total number of megabytes that were read from the flash cache during the specified time range.
- Data written to flashcache (MB) – the total number of megabytes that were written to the flash cache during the specified time range.

Flash cache usage chart

The Flash Cache Usage (MB) chart displays the following metrics on all of the celldisks, as well as on the celldisks selected in the cell table:

- Flash cache usage – the total amount of space (in megabytes) used on the flash cache.
- Flash cache usage (keep objects) – the total amount of space (in megabytes) used only on the keep objects in the flash cache.

Flash cache request misses chart

The Flash Cache Request Misses (requests) displays the total number of read IO requests that did not find all data in Flash Cache and therefore had to perform physical read operations. This figure is displayed for all of the celldisks, as well as on the disks selected in the cell table, and is split into two charts:

- Total number of request misses of all types.
- Total number of request misses resulting from read IO requests for keep objects.

Celldisks statistics

The Celldisk Statistics section displays the rate (megabytes per second) of physical reads and writes that were carried out during the selected time range on all of the disks, as well as on the disks selected in the cell table.

Reviewing the InfiniBand Topology

Use the InfiniBand topology drilldown to monitor the InfiniBand switches and nodes.

This drilldown includes the following sections:


- Switches – includes a table that displays all of the InfiniBand switches, which are used for transferring information from the database to the storage cells and conversely. This section provides the following details:
 - Number of InfiniBand switches
 - Subnet manager on switch – identifies which of the switches is a master switch.
 - Link status – the InfiniBand Switch's current link status, as indicated during the last sample.
 - Health status – the InfiniBand Switch's current health status, as indicated by the success or failure of the environment test.
- Nodes – includes a table that displays all of the InfiniBand nodes. This section provides the following information for each node:
 - Node name
 - Node IP address
 - Node host channel adapter (HCA)
 - Node status
 - Node's ports on each of the InfiniBand switches
- Charts – display the rate (Mb/S) of data sent and received through the InfiniBand switch during the specified time range.

Administering Foglight *for Oracle*

You use the Databases Administration dashboard to set options for collecting, storing, and displaying data about monitored Oracle instances.

Opening the Databases Administration Dashboard


You can edit settings for one or more Oracle instances on the Databases > Administration dashboard.

 **NOTE:** If you attempt to select instances of more than one type of database, such as an SQL Server database and an Oracle database, an error message is displayed.

To open the Databases Administration dashboard:

- 1 In the navigation panel, under Homes, click Databases.
- 2 Select the row check boxes beside one or more Oracle instances.
- 3 Click Agent settings and then click Administration.

The Administration dashboard opens, containing settings for all the selected agents. Settings are broken down into categories, which are organized under an Oracle tree.

 **TIP:** The list of agents you selected can be found by clicking Selected Agents.

Reviewing the Administration Settings

Use the Databases Administration dashboard to set options for collecting, storing, and displaying data, which apply to all of the currently selected agents. Click a category of settings on the left (for example: Connection Details) to open a view containing related settings on the right.

The metrics defined on the Databases Administration dashboard apply to all of the agents that were selected before opening the Administration dashboard. As a result, the same unit of measure and aggregation value for display are enforced for all currently selected agents.

To view the full list of selected agents, click Selected Agents button at the upper right corner of the screen. To change the list of agents to which the metrics apply, exit the Databases Administration dashboard, select the required agents, and re-open the dashboard.

If the settings vary between the selected agents (for example: one agent uses the measurement unit kilobyte, while another uses megabyte), the fields that contain non-identical values are displayed as empty and marked

with an Inconsistent Values () icon.

Changes made to settings should be saved before selecting another category of settings.

To save changes made in an Administration dashboard view:

- 1 In the Database Administration dashboard, select a category of settings from the left-hand panel.
- 2 Make changes to settings as necessary.
- 3 Click Save changes at the bottom of the view.

If you attempt to exit the view without saving changes, a Warning dialog box prompts you to confirm your action.

Defining Connection Details

Use the Connection Details category to define global connection settings, which apply to all instances and hosts selected in the view. You can enable user-defined collections and set VMware connection details.

NOTE: The following sections instructs how to define the connection settings for monitored Oracle instances.

Defining the Connection Settings for the Monitored Oracle Instances

The Connection Details view contains a table that displays all the agents that were selected before entering the Databases Administration dashboard.

To define the connection settings for the requested agents:

- 1 Select the check boxes to the left of the agents for which uniform credentials are to be set. To cancel the selection, click **Select None** and select again.
- 2 Click **Set Credentials**.
The **Edit Instance Credentials** dialog box opens.
- 3 Use the **Connection Details** section to enter the TNS listener port, or accept the default port (1521).
- 4 Use the **Login Credentials** section to type the user name and the password used for connecting to the Oracle instance.
- 5 Select whether to enable OS monitoring, in order to retrieve OS-related data, such as CPU utilization and memory consumption. For instructions about configuring OS monitoring, see [Configuring OS monitoring credentials](#) on page 86.
IMPORTANT: When enabling OS monitoring, ensure that the selected OS user has the privileges required for monitoring the operating system. For details, see the Foglight for Oracle Release Notes.
- 6 Select whether to enable and set user-defined collections. For more information, see [Enabling and setting credentials for user-defined collections](#) on page 87.
- 7 If the selected agents reside on a VMware virtual host, click **Edit VMWare connection details** to define their connection details. For more information, see [Defining the VMware Connection Profile](#) on page 87.
- 8 Click **Validate** to proceed to the next stage of validating the instance's connectivity.

Validating Connectivity and Starting to Monitor the Instances

After setting the default credentials for the host, these newly created credentials can now be used by the wizard to attempt to log in to the instances.

To validate the instances' connectivity:

- 1 Click **Validate connectivity**.

The **Verifying Connectivity** progress bar appears.

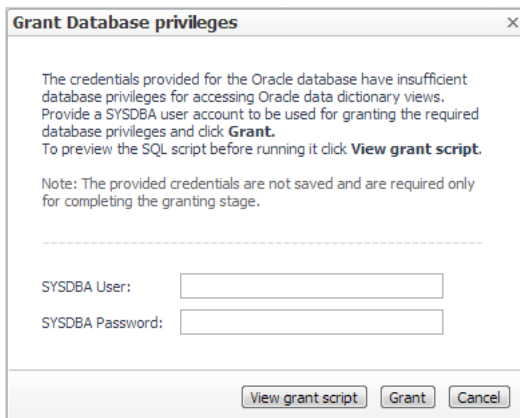
At the end of this process, any connectivity issues are listed in the **Status** column of the instance table. When the connection is successful, the **Status** column displays the status message *Validated*, which indicates that the instance connected successfully and the specified Oracle user has the required permissions.

If the connection failed verification, the **Status** column displays one of several connectivity status messages.

The messages, causes, and appropriate responses are:

- No valid RAC nodes were detected (for a RAC) – modify the RAC nodes’ credentials.
- Invalid username/password – check the credentials and try again.
- Process Error – see the code error under the Details column.
- Wrong Database Credentials – modify the login credentials.
- Database Connection Failed – the database to which the instance tries to connect is not running or cannot be accessed. Ensure clearing the check box near the database’s name.
- Insufficient Privileges – grant the user the privileges required for connecting to the database, by clicking the status Insufficient Privileges.

The Grant Database Privileges dialog box opens.



Use this dialog box to specify a SYSDBA (System Database Administrator) user with sufficient privileges.

Type a SYSDBA user and password, and then click Grant.

If the Sysdba credentials entered were incorrect, the column displays the status *Wrong Sysdba credentials*.

After correcting the mistakes that resulted in the connectivity failure, click again *Validate connectivity*.

- 2 Click *Validate connectivity* on the status bar.
- 3 Click *Save Changes*.

The Applying Modified Settings progress bar appears.

Configuring OS monitoring credentials

To configure the OS monitoring credentials:

- 1 Use the Host Credentials section to select whether to log in to the monitored host using the existing host connection details or to edit the host connection details.
- 2 If the option to edit the host connection details is selected, select whether to disable OS monitoring, which is enabled by default.

IMPORTANT: When enabling OS monitoring, ensure that the selected OS user has the privileges required for monitoring the operating system. For details, see the Foglight for Oracle Release Notes.

- 3 Specify the requested authentication type:
 - Local user – using the same credentials that were used for running the monitoring software (Agent Manager) on the Oracle host.
 - Windows – the user name (in the domain\user name format) and password, used for logging in to a Windows-based monitored host.
 - SSH (DSA) – using the SSH private key, generated via the DSA algorithm.

- SSH (RSA) – using the SSH private key, generated via the RSA algorithm.
 - SSH (login credentials) – the user name and password used for logging in to a UNIX-based monitored host.
- 4 If any option other than *Agent Manager local user* is selected, enter the required credentials.
- i** **IMPORTANT:** If either of the UNIX authentication types is selected, to run certain commands that require administrative privileges as root, without having to log in as administrator, select the Use sudo check box.
- i** **IMPORTANT:** It is possible to delete or edit existing credentials by clicking Select from stored credentials and entering the User Credentials dialog box. However, it is highly advisable not to make any changes to existing credentials, as such changes will affect all agents that currently use these credentials.


Enabling and setting credentials for user-defined collections

Use the user-defined collections feature to add customized collections to all of the currently selected agents, thereby providing for queries not included in Foglight *for Oracle*.

Use the Connection Details view to enable and configure the creation of user-defined collections.

To enable and configure the creation of user-defined collections:

- 1 Select the requested agents.
- 2 Click Set credentials.
The Edit Instance Credentials dialog box opens.
- 3 Click Edit Credentials for User-defined Collections.
The Edit Credentials for User-defined Collections dialog box opens.



- 4 Select the check box Enable user-defined collections.
- 5 Select whether to perform the collection using the existing agent credentials or by creating dedicated credentials.
- 6 If the user-defined collections are to be enabled using dedicated credentials, type the requested user name and password.
- 7 Select whether to connect as SYSDBA.
- 8 Click Set.

To add user-defined collections, go to the User-defined Collections view in the Databases Administration dashboard. For details, see [Configuring User-defined Collections](#) on page 104.

Defining the VMware Connection Profile

If you have Oracle instances that run on virtual hosts, monitoring such instances requires setting a dedicated connection profile, in order to connect to the requested VMware server.


- i** **NOTE:** The VMware connection profile can be defined only if OS monitoring is enabled.

Establishing such a connection is necessary in order to retrieve the Virtualization overhead data, that is, the percentage of CPU that is unavailable to this virtual machine because it is being utilized either by other virtual machines or by VMware itself. The Virtualization Overhead indicator is displayed in both the real-time and history summary pages.

To edit the VMWare Connection Profile:

- 1 Select the requested agents.
- 2 Click Set credentials.
- 3 In the Edit Credentials dialog box, click Edit VMWare connection details.

The Edit VMWare Credentials dialog box opens.



- 4 Select the check box Enable monitoring the VMWare CPU layer.
- 5 Enter the details required for monitoring the CPU distribution data, that is: vCenter or ESX host name, port, VMware user, and VMware password.
- 6 Click OK.


Defining the Connection Settings for the Monitored ASM Instances

The Connection Details view contains a table, which displays all of the agents that were selected before entering the Databases Administration dashboard.

To define the connection settings for the requested ASM agents:

- 1 Go to the ASM tab.
- 2 Select the check boxes to the left of the requested ASM agents.
- 3 Click Agent settings.
- 4 Click Configure ASM agents.
- 5 In the Connection Details view, select the check boxes to the left of the ASM agents for which uniform credentials are to be set. To cancel the selection, click Select None and select again
- 6 Click Set Credentials.

The dialog box Edit credentials of ASM instance opens.



This dialog box is identical to the one used for configuring ASM instances for monitoring.

Defining Data Collection and Storage Options

The Foglight *for Oracle* agent collects and stores data at all times, even when the browser window that displays the data is not active. Use the Collection Frequencies view in the Databases Administration dashboard to specify:

- Which collections are sampled and stored.
- The data collection values when sampling is carried out in offline, online, and real-time frequency modes.
- The collection frequency.

When a user is currently focusing on a screen, the sampling frequency for all of the collections associated with this screen automatically switches to Real-Time. The collection frequency setting determines the sampling frequencies of the other collections, which do not run in Real-Time mode).

The available collection frequencies, which can be selected from the field *The selected agent's current collection frequency is*, are:

- Low – all collections are running in Offline mode, regardless of whether a Client is connected
 - Normal – the collections' running mode (Online/Offline/Real-time) adjusts dynamically to the Client's connection status (disconnected/connected/focusing on a screen).
 - High – all collections are running in Online mode, regardless of whether a Client is connected.
- The Query timeout for on-demand collections.

This setting defines the number of seconds that a query for on-demand collections can run before it times out.

On-demand collections are collections whose data is retrieved not by predefined time intervals but upon entering a screen or clicking a button.

The default setting of this parameter is 60 seconds, but it can be modified by clicking the number that indicates the parameter's value in the field *Query timeout for on-demand collections*.

The *Collections* table contains the columns listed below.

Column	Description
Collection Enabled	Defines whether the selected collections are sampled and stored.
Collection Name	The name of the collection. This list is sorted in alphabetical order.
Offline Frequency (Sec)	Used for defining the collection interval, in seconds, in offline mode (for example: 300). Offline frequency refers to the longest interval possible for sampling the monitored instance.
Online Frequency (Sec)	Used for defining the collection interval, in seconds, in online mode (for example: 60).
Real-time Frequency (Sec)	Used for defining the collection interval, in seconds, in real-time mode (for example: 20). Only one collection can be sampled at real-time frequency at any given moment.
Query Timeout (Sec)	The amount of time, in seconds, that elapses before the query times out.

IMPORTANT: In the Normal collection frequency, when the browser window that displays Foglight for Oracle is active, the collection frequency mode in the active screen (for example: the Current Sessions panel in the Activity drilldown) switches to the fastest frequency possible - once every 20 seconds.

To modify the values of a specific collection:

- 1 Select the collection's row in the table.
- 2 Click Edit.
The Edit the Collection dialog box opens.
- 3 Select whether to enable the collection and storage of the selected collection.
- 4 Set the collection interval, in seconds, in offline frequency mode (if available).
- 5 Set the collection interval, in seconds, in online frequency mode (if available).
- 6 Set the collection interval, in seconds, in real-time frequency mode.
- 7 Set the query timeout, in seconds.
- 8 Click Set to apply these settings or Cancel to reject them.

Customizing Alarms for Foglight *for Oracle* Rules

Many Foglight *for Oracle* multiple-severity rules trigger alarms. To improve your monitoring experience, you can customize when alarms are triggered and whether they are reported. You can also set up email notifications.

This section covers the following topics:

- [Introducing the Alarms View](#)
- [Setting and Modifying Alarm Sensitivity Levels](#)
- [Modifying Alarm Settings](#)
- [Configuring Email Notifications](#)
- [Cloning Agent Settings](#)

- [Reviewing Rule Definitions](#)

Introducing the Alarms View

The Alarms view enables you to modify global settings and agent-specific settings for alarms.

To open the Alarms view:

- 1 Open the Administration dashboard as described in [Opening the Databases Administration Dashboard](#) on page 84.
- 2 Click Alarms.
The list of agents that you selected on the Databases dashboard is shown in the upper right corner of the view.
- 3 From the Alarms view, you can complete the following tasks:
 - [Setting and Modifying Alarm Sensitivity Levels](#)
 - [Modifying Alarm Settings](#)
 - [Configuring Email Notifications](#)
 - [Cloning Agent Settings](#)

Setting and Modifying Alarm Sensitivity Levels

Foglight *for Oracle* has three sensitivity levels that control which alarms are reported:


- Essential – Enables only the most important alarms.
- Normal – Enables a selection of alarms from different rule categories, including all Essential alarms.
- Tuning – Enables all alarms.

You can change the sensitivity level assigned to each agent. If a sensitivity level does not include all the alarms you want to track or includes too many alarms, you can view a list of multiple-severity rules and modify the sensitivity level that is mapped to each severity.

Changes made to a sensitivity level affect all agents that are assigned that sensitivity level. If you want to enable or disable alarms for the selected agents, see [Enabling or disabling alarms for selected agents](#) on page 93.

Setting the alarm sensitivity level by agent

Each agent has its own sensitivity level setting. The default is Normal.


-  **TIP:** Select a sensitivity level that is closest to what you want, then customize it as necessary. For instructions, see “Viewing and modifying alarms assigned to severity levels” on page 177.

To change the sensitivity level used by an agent:

- 1 In the Alarms view, click the Sensitivity Level tab.
- 2 Select the check boxes for agents you want to edit, and then click a sensitivity level button, such as Essential.
- 3 Click Save changes.

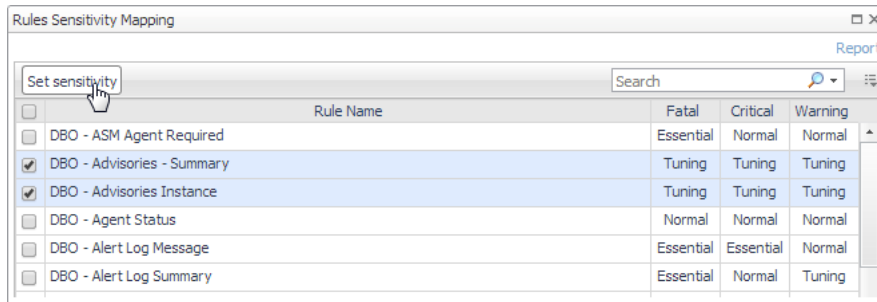
Viewing and modifying alarms assigned to severity levels

You can view a list of multiple severity rules to see which severities are assigned to which sensitivity level. If desired, you can change the assignments. Changes to sensitivity levels affect all agents.

-  **TIP:** If you want to see descriptions of the rules, follow the steps described in “Reviewing Rule Definitions” on page 187.

To view and edit alarms assigned sensitivity levels:

- 1 In the Alarms view, click the Sensitivity Level tab.
- 2 Click View rules sensitivity mapping.
Review the assignments. Recall that the Normal level includes all alarms assigned to the Essential level, and the Tuning level includes both Normal and Essential alarms.
- 3 If you want a record of the existing settings, click Report and export the settings to a file.
- 4 Select the check boxes for the rules you want to edit, and then click Set sensitivity.



- 5 In the Set Rule dialog box, select a sensitivity level for each severity.
- 6 Click Set.
All selected rules are updated with the new assignments.

Modifying Alarm Settings

You can customize how the alarms generated by the default Foglight *for Oracle* rules are triggered and displayed in the Alarm view's Settings tab. All changes to alarm settings apply to the selected agents, with the exception of thresholds, which can be customized by agent.

- TIP:** If you want to copy alarm settings to another agent that was not in the selected list— for example, you add a new agent — you can copy the alarm customizations to that agent. For instructions, see [Cloning Agent Settings](#) on page 98.
- IMPORTANT:** Avoid editing Foglight for Oracle rules in the Administration > Rules & Notifications > Rule Management dashboard. Default rules may be modified during regular software updates and your edits will be lost. Always use the Alarms view.

The Alarms list controls the contents displayed to the right and the tasks that are available.

- **All Alarms** - Displays all rules with configured alarms and indicates whether alarms are enabled. In this view, you can enable or disable alarms for all the rules at once. You can also set email notifications and define mail server settings.
- **Category of rules** - Displays a set of related rules with configured alarms. In this view, you can enable or disable alarms and also set email notifications for the category of rules.
- **Rule name** - Displays the alarm status for the selected rule. If the rule has multiple severity levels, displays the threshold configured for each severity level. In this view, you can enable or disable the alarm, edit the alarm text, and edit severity levels and their thresholds. You can also set email notifications for the alarm.

You can complete the following tasks:

- [Enabling or disabling alarms for selected agents](#)
- [Modifying alarm threshold values](#)
- [Editing the text of the alarm message](#)

Your changes are saved separately and applied over the default rules. This protects you from software upgrades that may change the underlying default rules.

Enabling or disabling alarms for selected agents

You can override the global [alarm sensitivity level](#) setting for the selected agents. You can enable or disable alarms for all rules, a category of rules, or an individual rule.

TIP: If you want to see descriptions of the rules, follow the steps described in [Reviewing Rule Definitions](#) on page 99.

To enable or disable alarms:

- 1 In the Alarms view, click the Settings tab.
- 2 Decide on the scope for the change: all alarms, a category of rules, or a selected rule.
- 3 Complete the steps for the selected scope:

Scope	Procedure
All alarms	Click All Alarms. In the Alarms Settings tab, click either Enable all or Disable all.
Category of rules	Click a category. Click either Enable all or Disable all.
Selected rule	Click the rule. In the Alarms Settings tab, click the link that displays the alarm status. Select Enabled or Disabled from the list and click Set.

- 4 Click Save changes.

Modifying alarm threshold values

You can and should modify the thresholds associated with alarms to better suit your environment. If you find that alarms are firing for conditions that you consider to be acceptable, you can change the threshold values that trigger the alarm. You can also enable or disable severity levels to better suit your environment.

When a rule has severity levels, a Threshold section appears in the Alarm Settings tab showing the severity levels and bounds by agent. Many rules, such as Baseline rules, do not have severity levels and thresholds.

When editing thresholds, ensure that the new values make sense in context with the other threshold values. For most metrics, threshold values are set so that Warning < Critical < Fatal. However, in metrics where normal performance has a higher value, such as *DBO - SGA Library Cache*, the threshold values are reversed: Warning > Critical > Fatal.

TIP: If you want to review the thresholds for all Foglight for Oracle rules in a single view, use the Rule Management dashboard. In the navigation panel, click Homes > Administration, then click Rules. Type DBO in the Search field to see the list of predefined rules for Oracle databases. For rules with severity levels, you can see the threshold values set for each level. If you want to edit threshold values, return to the Alarms view. Edits made directly to the default rules may be overwritten during software upgrades.

To change severity levels and thresholds:

- 1 In the Alarms view, click the Settings tab.
- 2 Click the multiple-severity rule that you want to edit.
- 3 Click the Alarms Settings tab.
- 4 In the Threshold section, review the defined severity levels and existing threshold bounds for all target agents. The colored bar across the bottom of the view indicates the percentage of collections that fall within each threshold.

Alarm Settings | Email Notification Settings

Alarm Status

Alarm status for the selected rule: ● Enabled Enhance alarm

Threshold

Agent Name	Normal Bound	Warning Bound	Critical Bound	Fatal Bound
<input checked="" type="radio"/> israix07-O112264 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> isrvmn808-o112364 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> isrvmn853-O1264CDB Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> isrvmpalw05-O112464 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> isrvmrh531-O121RAC3 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> isrvmrh532-O121RAC2 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> isrvmrh533-O121RAC1 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	
<input type="radio"/> RAC-isrvmrh531-O121RAC3 Edit	0.0 - 80.0	80.0 - 95.0	95.0 -	

5 From this view, you can perform the following tasks:

Task	Procedure
Edit severity levels and set threshold (lower bound) values for all agents.	Click Enhance alarm. Select the check boxes for the severity levels you want enabled and set the threshold values. Click Set.
Change the threshold (lower bound) values for one agent.	Click Edit beside the agent name. Set the new threshold values and click Set.
Copy the changes made to one agent's threshold values to all other agents.	Click Edit beside the agent name that has the values you want to copy. Select Set for all agents in table and click Set.

6 Click Save changes.

Editing the text of the alarm message

For individual rules, you can change the message displayed when an alarm fires. You cannot add or remove the variables used in the message. This is a global setting that affects all agents.

To change the alarm message:

- 1 In the Alarms view, click the Settings tab.
- 2 Select a rule.
- 3 Click the Alarm Settings tab.
- 4 Click Enhance alarm.
A Customize <rule> dialog box opens.
- 5 In the Message box, edit the message text.
If desired, you can restore the default message by clicking Reset message.
- 6 Click Set.
- 7 Click Save changes.

Configuring Email Notifications

We recommend that you set email notifications for the alarms you are most interested in tracking closely. For example, you may want to be notified by email of any Critical or Fatal situation. Or you may want to be informed whenever a key metric, such as CPU usage, is no longer operating within acceptable boundaries.

You can set up email notifications that are generated when an alarm fires and/or on a defined schedule, as described in the following topics:

- [Configuring an email server](#)
- [Defining Default Email settings](#)
- [Defining email notifications, recipients, and messages](#)
- [Defining variables to contain email recipients](#)
- [Defining scheduled email notifications](#)

Configuring an email server

You need to define the global mail server variables (connection details) to be used for sending email notifications. The setting of the email should be configured in **Foglight Administration > Email configuration**.

Defining Default Email settings

You can define a default email address to be used by every new agent created in the future, by selecting the Default email button when configuring email notification.

The Email addresses entered are applied to all monitored agents not only for the agents that were selected to enter the Alarm administration.

Enabling or disabling email notifications

You can enable or disable email notifications for all alarms, a category of alarms, or a selected rule. Email notifications are sent only if all the following conditions are met:

- The alarm email notification setting is enabled for the affected rule.
- The alarm is triggered by changes in the monitored environment.
- Alarm notification is enabled at the triggered severity level. See [Defining email notifications, recipients, and messages](#).

To enable or disable email notifications:

- 1 In the Alarms view, click the Settings tab.
- 2 Decide on the scope for the change: all alarms, a category of rules, or a selected rule.
- 3 Complete the steps for the selected scope:

Table 1. Enable or disable email notification settings.

Scope	Procedure
All alarms	Click All Alarms . In the Alarms Settings tab, click the Define Email Settings button. Select either Enabled or Disabled from the Alarms notification status list. Click Set .
Category of rules	Click a category. Click the Define Email Settings button. Select either Enabled or Disabled from the Alarms notification status list. Click Set .
Selected rule	Click a rule. In the Alarms Settings tab, click the Define Email Settings tab. Click the link that displays the alarm notification status. Select Enabled or Disabled and click Set .

- 4 Click **Save changes**.

Defining email notifications, recipients, and messages

You control who receives email messages, the subject line, and some text in the body of the email. The body of the email always contains information about the alarm. This information is not editable. You can also control whether an email is sent based on severity levels. You can set different distribution lists for different rules and different severity levels, or set the same notification policy for all rules.

To configure email notifications:

- 1 In the Alarms view, click the Settings tab.
- 2 Decide on the scope for the change: all alarms, a category of rules, or a selected rule.
- 3 Complete the steps for the selected scope:

Table 2. Configure email notification settings:

Scope	Procedure
All alarms	Click All Alarms . In the Alarms Settings tab, click the Define Email Settings button. Continue to Step 4 .
Category of rules	Click a category. Click the Define Email Settings button. Continue to Step 4 .
Selected rule	Click a rule. Click the Define Email Settings tab. <ul style="list-style-type: none">• To change the severity level that warrants an email notification, click the link that displays the severities. Select the desired level of severity and click Set.• To configure email recipients and the message, select the tab for a severity level, and click Edit. Skip to Step 5.

- 4 If you selected All Alarms or a category, in the Email Notification Settings dialog box, do one of the following:
 - To change the severity levels that warrant an email notification, from the Messages will be enabled for severities box, select the desired levels of severity.
 - To configure the same email recipients and message for all severity levels, click Configure settings for all severities and then click All severities.
 - To configure different email recipients and messages for each of the severity levels, click Configure severity specific settings and then click a severity level.
- 5 In the Message Settings dialog box, configure the email recipients and message.
 - To – Type the addresses of the people who need to take action when this alarm triggers.
 - CC – Type the addresses of the people who want to be notified when the alarm triggers.

NOTE: If a mail server is not found, you are prompted to configure a mail server. For instructions, see [Configuring an email server](#) on page 95.

You can use registry variables in place of email addresses. Type the variable name between two hash (#) symbols, for example: #EmailTeamName#. For more information, see [Defining variables to contain email recipients](#) on page 97.

- **Subject** – Optional. Edit the text of the subject line to better suit your environment. Avoid editing the variables, which are identified with the @ symbol.
- **Body Prefix** – Optional. Add text that should appear above the alarm information in the body of the email.

Figure 1. Message Settings.

Message Settings

Enter the recipient email address or multiple addresses separated by comma.
Note: the system supports entering names of registry variables, surrounded by pound keys (#registry_name#), in the To and CC field.

To:

CC:

Subject: ?

Body:

Body prefix:

Body: Foglight for <dbType> has generated a <Severity> alarm for instance <Instance name> on host <Host Name>
Alarm: <Alarm name>
Severity: <Severity>
Message: <Alarm message>
Created on: <Date of creation>


Set Cancel

- 6 Click Set to save the message configuration and close the dialog box.
- 7 If the Edit Notification Settings dialog box is open, click Set.
- 8 Click Save changes.

Defining variables to contain email recipients

You can create registry variables that contain one or more email addresses and (optionally) their scheduled notifications, and use these registry variables when defining email notifications. This procedure describes how to create a registry value. For schedules, see [Defining scheduled email notifications](#).

To create a registry variable:

- 1 On the navigation panel, under Dashboards, click Administration > Rules & Notifications > Manage Registry Variables.
- 2 Click Add.
The New Registry Variable Wizard opens.
- 3 Select the registry variable type String, and click Next.
- 4 In the Name field, enter a name, for example: EmailTeamName
Optional – Add a description.
- 5 Click Next.
- 6 Select Static Value.
- 7 In the Enter desired value box, enter one or more email addresses (separated by commas).
 **NOTE:** Email groups are not permitted.
- 8 Click Finish.



The Edit Registry Variable dashboard displays the newly created registry variable.

To use a registry variable in email notifications, type the variable name between two hash (#) symbols, for example: #EmailTeamName#. For more information, see [Defining email notifications, recipients, and messages](#) on page 96.

Defining scheduled email notifications

If someone wants to receive an email about an alarm on a regular basis, such as once a day, you use a registry variable schedule to set up the notification.

To schedule the sending of email notifications for a registry variable:

- 1 If you are continuing from [Defining variables to contain email recipients](#), the registry variable is already open for editing in the Edit Registry Variable dashboard.
-  **TIP:** To edit a different variable, navigate to the Administration > Rules & Notifications > Manage Registry Variables dashboard, click the variable name, and select View and Edit Details.
- 2 In the Performance Calendars List table, click Add.
The Performance Calendar Wizard opens.
- 3 Select a schedule, for example: End of Day
- 4 Click Next.
- 5 Select Static Value.
- 6 In the Enter desired value box, enter one or more email addresses (separated by commas) that should receive email notifications based on the schedule.
-  **TIP:** The addresses may be the same as or different from those assigned to the registry variable.
- 7 Click Finish.
The Edit Registry Variable dashboard displays the newly created schedule. If desired, repeat to add other schedules.

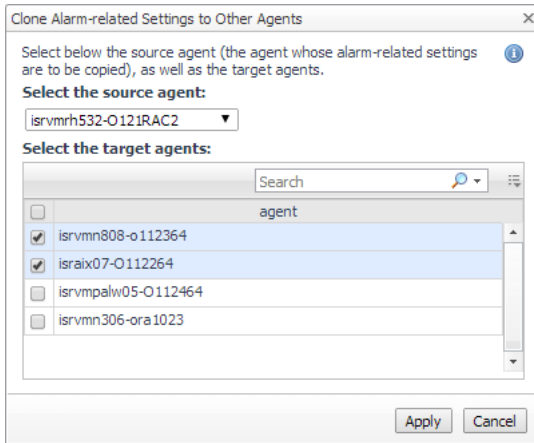
Cloning Agent Settings

You may want an agent to have the same settings as another agent. For example, if you add new agents, you may want them to use the same settings as an existing agent. In this case, you can clone the settings from one agent to other agents. This process does not link the agents; in the future if you update the source agent, you also need to update the target agents.

This procedure walks you through selecting the source agent from the Databases dashboard. However, you can also open the Administration dashboard with multiple agents selected. In this case, you select the source agent in Clone Alarm-related Settings to Other Agents dialog box.

To clone alarm-related settings:

- 1 On the Databases dashboard, select the check box for the agent with the settings you want to clone.
- 2 Click Agent settings and then Administration.
- 3 In the Administration dashboard, click Alarms.
- 4 Click Set configuration on selected agents.
The Clone Alarm-related Settings to Other Agents dialog box opens.
- 5 In the Select the source agent drop-down list, you should see the agent you selected.
- 6 In the Select the target agents table, select the check boxes for agents that should inherit settings from the source agent.



- 7 Click Apply.
- 8 When prompted for confirmation, click Yes.

Reviewing Rule Definitions

If you want to review the conditions of a rule, open the rule in the Rule Management dashboard.

IMPORTANT: Avoid editing Foglight for Oracle rules in the Rule Management dashboard. These rules may be modified during regular software updates and your edits will be lost. Always use the Alarms view.

You can create user-defined rules from the Rule Management dashboard. If you want to modify a rule, we recommend copying the rule and creating a user-defined rule. User-defined rules need to be managed from the Rule Management dashboard; these rules are not displayed in the Alarms view of the Databases Administration dashboard. For help creating rules, open the online help from the Rule Management dashboard.

To open the Rule Management dashboard:

- 1 On the navigation panel, under Homes, click Administration.
- 2 In the Administration dashboard, click Rules.
- 3 Type DBO in the Search field to see the list of predefined rules for Oracle databases.
The Foglight *for Oracle* rules are displayed. From here, you can review threshold values, alarm counts, and descriptions.
- 4 To see the full rule definition, click a rule and then click View and Edit.
- 5 In the Rule Detail dialog box, click Rule Editor.
- 6 When you are done your review, click Rule Management in the breadcrumbs to return to the dialog box.

Defining Retention Policies

Use the Retention Policies view in the Databases Administration dashboard to define the requested duration for which each of the metric collections, which are defined in the Collections view, are to be kept.

Data can be retained for brief, moderate, or long periods, by selecting one of the following options:

- Short – retains data up to one month.
- Medium – retains data up to three months.
- Long – retains data up to one year.

To modify the retention policy:

- 1 Click the text that displays the current retention policy.
The Edit the Retention Policy dialog box opens.
- 2 Select the requested retention policy scheme from the list.
- 3 Click Set to apply the selected setting or Cancel to reject the setting.
- 4 At the bottom of the view, click Save changes.

Defining the Collection Settings

Use the Collection Settings view to limit the number of displayed items for each of the collections below.

Locks

Use the Locks pane to define the following settings:

- The maximum number of blocked segments to be displayed in the Lock Overview pane of the Activity > Locks panel.
- The minimal duration, in seconds, which a lock should reach or exceed in order to be collected. All locks that meet this criterion are displayed in the lower section of the Locks Historical Tree pane.

To modify the lock collection and display settings:

- 1 Click Edit.
The Edit Lock Collection and Display Settings dialog box opens.
- 2 Use the designated field to type the maximum number of blocked segments.
- 3 Use the designated field to type the minimal duration for collecting and displaying a lock in the Locks Historical Tree
- 4 Click Set to approve the new setting or Cancel to undo the operation.

Top SQL Statements

Use the Top SQL Statements pane to configure the requested settings for top SQL statements, that is, statements that generated the maximal workload, as defined by the criterion by which the selected SQL statements are sorted. Use this pane to define the criterion of top SQL statements, as well as the maximum number of such statements to be retrieved and displayed. Top SQL statements are displayed on the Activity > Top SQL Statements panel.

To define the settings for collecting and displaying SQL statements:

- 1 Click Edit.
The Edit Top SQL Statements Settings dialog box opens.
- 2 Use the *Maximum number of displayed SQL statements* field to enter the maximum number of SQL statements that are to be displayed in the Top SQL Statements grid.
- 3 Use the *Maximum number of retrieved SQL statements* field to enter the number of SQL statements that generated the maximal workload and are to be saved to the Top SQL statements collection.
- 4 Use the *Sort the collected SQL statements by* list to select the field by which the Top SQL Statements list is to be sorted.

The possible values are:

- CPU Time – the total CPU time consumed for carrying out the SQL statement executions

- Elapsed Time – the total time consumed for carrying out the SQL statement executions
 - Executions – the number of times the SQL script executed for a unique SQL Handle
 - Total Waits – the total wait time spent while executing the SQL statements
- 5 Use the *Maximum size of short SQL statement* field to type the maximum number of characters for the short SQL text.
 - 6 Click Set to save these settings.
 - 7 At the bottom of the view, click Save changes.

Tablespace Settings

Use the Tablespace Settings pane to configure the default monitoring settings for the Tablespaces Storage collection, by excluding from monitoring one or more of the following types of tablespaces:

- Undo tablespaces
- Temporary tablespaces
- Read-only tablespaces
- System-Sysaux tablespaces

This pane also includes the Tablespace Exclude List section. Use this section to add one or more tablespaces to the Exclude list.


To add tablespaces to the Exclude list:

- 1 Click Add.
The Add a Tablespace to the Exclude List dialog box opens.
- 2 Use the Tablespace name field to enter a field name. To exclude multiple tablespaces, insert a regular expression (wild card) and select the Regex check box.
- 3 Click Add to complete the operation.

Invalid Objects Settings

Use the Exclude List pane to excluding specific schemas from being monitored by the Invalid Objects collection.

To exclude a schema from being monitored:

- 1 Select the agent that monitors the instance where the schema resides.
 **IMPORTANT:** If the selected agent is a RAC node, changes made will apply to RAC agent, thereby affecting the entire RAC.
- 2 Click Add a Schema to the Exclude List.
- 3 Use the dialog box that appears to enter the name of the requested schema.
- 4 Click Add.
- 5 Repeat this procedure for all of the schemas that are to be excluded.
To remove schemas from the table, select the schema's row and click Delete.
- 6 At the bottom of the view, click Save changes.

Configuring the Alert Log Panel Display

Use the Alert Log view to configure the retrieval and display settings for the Alerts > Alert Log panel. In addition, this view is used to define default settings for invoking alarms based on the collected alerts.

- ① **IMPORTANT:** While alerts whose severity type is Informational are collected and displayed on the Alert Log panel, alarms are invoked only for alerts whose severity type is Warning or higher.

Use this view to define the following settings:


- Minimal severity for invoking summary alarms – can be either turned off or set to one of the defined severity values, that is: Warning, Critical or Fatal.
 - ① **IMPORTANT:** By default, the value of this parameter is Critical.
 - Minimal severity for invoking alarms – can be either turned off (the default setting) or set to one of the defined severity values: Warning, Critical or Fatal.
 - Number of rows to collect from logs – the maximum number of rows that would be retrieved from the alert logs.
 - Display under a pre-defined name in the Alert Log panel – using the Match List pane. This pane contains a default list of expressions within the alert logs that are to be retrieved and displayed in the Alert Log panel under a pre-defined name, category and severity.
 - Exclusion from the Alert Log panel display – using the Ignore List pane, which contains a default list of alerts that are to be excluded when setting the alert logs display.
- ① **IMPORTANT:** Only messages that are explicitly defined in the Ignore List will not be displayed. Messages that were not added to either the Match or Ignore lists appear under name Other, type Oracle alert messages and severity Informational. Therefore, ensure that messages that need not be displayed are added to the Ignore List.

Both the Match List and the Ignore List panes can be customized by adding, editing or removing alert logs. Each filter can be enabled or disabled separately by clicking Edit and selecting or clearing the *Enabled* check box. Alternatively, to enable or disable all of the filters, click the Enable All or Disable All button.

To add an alert to the Match List:

- 1 Click Add.

The Add an Alert to the Match List dialog box opens.



Use this dialog box to configure the alert filter by means of the following fields:

- Filter Name – used for assigning a name to the filter.
- Severity – used for assigning a severity for the alert. Use this field to set the type of the alarm that is invoked as a result of the alert, or to define that the alert is informational and therefore invokes no alerts.

- Type – used for entering a filter type, either one of the predefined types or a new type.
- Expression – identifies the text that is to be extracted as a message from the alert log and displayed in the Message column of the Alert Log panel > Alert Log Messages table. This text usually contains the beginning of the message and a wildcard (for example: ORA-12012.*).

IMPORTANT: Alert filters are enabled automatically upon addition. To disable a filter, use the Edit button.

- 2 Click Add to save your settings.

Each newly added alert filter is enabled by default. To disable the filter, click Edit and then clear the *Enabled* check box.

To edit an alert:

- 1 Click the requested alert log.
- 2 Click Edit.
The Edit the Selected Alert dialog box opens.
- 3 Use this dialog box to configure the alert by means of the same fields used for adding it: Filter Name, Severity, Type, and Expression.
- 4 To disable the filter, clear the *Enabled* check box.
- 5 Click Set to save your settings.
- 6 At the bottom of the view, click Save changes.

Configuring the List of Listeners per Instance

Use the Listeners view in the Databases Administration dashboard to edit the list of monitored listeners per instance, by adding or removing listeners from the list.

To add listeners to an instance:

- 1 Open the drop-down list below the title *Please select an agent:*, to display the entire list of selected agents.

IMPORTANT: As RAC agents whose version is earlier than 11G2 have no listeners, they are not included in the list, even if such agents appear in the list of selected agents.

- 2 Select a single agent from the list.
- 3 Click Add Listener.

The contents of the *Add a Listener* dialog box that appears now vary between Windows and UNIX. For instructions for Windows-based agent, go to [Step 4](#).

In UNIX, enter a listener service name in the designated field.

- 4 For a Windows-based agent, type data in the following fields:
 - Listener name – the name of the listener
 - TNS Admin – the path of the \$TNS_ADMIN directory (usually *\$ORACLE_HOME/network/admin*)

- 5 Click Add.

The newly added listener appears in the Listeners table.

To edit the properties of a single listener:

- 1 Select the requested listener from the table.
- 2 Click Edit Listener.

- 3 In the Edit a Listener dialog box, edit the existing data in the various fields (see [Step 4](#) of the procedure for adding a listener).
- 4 Click Set to apply the changes or Cancel to revert to the previous settings.

To stop monitoring a listener:

- 1 Select the requested listener from the table.
- 2 Click Delete Listener.
The Stop Monitoring a Listener dialog box opens.
- 3 Click Yes to confirm the deletion or Cancel to keep monitoring the listener.
- 4 At the bottom of the view, click Save changes.

Configuring User-defined Collections

Use the User-defined Collections view in the Databases Administration dashboard to add user-defined collections to all of the currently selected agents, in order to provide for queries not included in Foglight *for Oracle*.

After collections are added, this view displays all user-defined collections for all of the agents; for example, if a collection was added to 12 agents during its addition, the view will display 12 rows, showing the collection for each agent.

This view is also used for configuring the sampling frequency for each collection.

The available sampling frequencies are:


- Real-Time - When a user is currently focusing on a screen, the sampling frequency for all of the collections associated with this screen switches to Real-Time.
- Online - When at least one user is connected, the sampling frequency for all of the collections that are not currently running at Real-Time frequency switches to Online.
- Offline - when no user is currently connected to the application

To add user-defined collections:

- 1 Select the requested agents.
- 2 Click Add.

If the user-defined collections are enabled for the selected agent, the Add a User-defined Collection dialog box opens. Proceed to [Step 3](#).

If user-defined collections are not enabled, the User-defined Collections Not Enabled dialog box opens.

 **IMPORTANT:** Agents can also be enabled for user-defined collections through the Connection Details view. For details, see [Enabling and setting credentials for user-defined collections](#) on page 87.

To enable and configure the creation of user-defined collections:

- a Click Edit credentials.
The Edit Credentials for User-defined Collections dialog box opens.
- b Select whether to perform the collection using the existing agent credentials or by creating dedicated credentials.
- c If the user-defined collections are to be enabled using dedicated credentials, type the requested user name and password.
- d Select whether to connect as SYSDBA.
- e Click OK.

- f The *Applying modified settings* progress bar is displayed. If the modified settings are successfully applied, the Add a User-defined Collection dialog box opens.

- 3 Click Set details.

The Set Collection Details dialog box opens.

- 4 Enter a name in the *Collection name* field.

- 5 Enter a brief description of the collection in the *Collection description* field (optional).

- 6 Paste the query's SQL text in the Query Text field.

IMPORTANT: This field can hold up to 4096 characters. Any PL-SQL Functions (that are not getting input) can be written by wrapping the PL-SQL "Select <pl_sql_func> From dual" statement.

- 7 Type a value, in seconds, in the Query Timeout field.

- 8 Use the Start Time field to define the time when the collection starts running (optional).

IMPORTANT: If this field is left empty, the collection starts working immediately. The start time can only be defined using the wizard. After the collection starts working, at the time defined using the Start Time field, it will continue to run according to its defined frequency and can only be disabled.

- 9 Click Verify.

The collection is verified by running the query on each of the currently selected agents.

After the verification process is complete, the *Verification Results* pop-up appears, indicating whether the collection was verified successfully. In case the collection verification failed, the relevant error message is displayed.

If the verification succeeded on at least one agent, the collection's details are displayed on the table described below.

Section	Field	Description
	ID	The collections' ID
Database	Column Name	The name of the column
	Column Type	The field type, as retrieved by the query (String, Integer and so on)
Data Storage	Display Name	The column's display name

Section	Field	Description
	Type	<p>The topology type for storage purposes. This type can be one of the following:</p> <ul style="list-style-type: none"> • String • Integer • Double • Boolean • Date <p>NOTE: When the Frequently modified check box is selected (the default state), the field's change history is kept, including use of optional functionalities such as use of the IntelliProfile mechanism and aggregation type selection. Clearing this check box is recommended only for fields whose values change infrequently, such as IP address of a specific host, as storing the change history of such fields is highly CPU-intensive and may degrade the FMS performance.</p>
	Unit of Measurement	<p>The metric's unit of measurement.</p> <p>The possible measurement unit values are: Percent, Count, Millisecond, Second, Minute, Hour, Day, Byte, Kilobyte, Megabyte and Gigabyte.</p>
	Aggregation	<p>Used for defining the value that is displayed in this metric, out of several values that were returned in the relevant time range.</p> <p>The available aggregation values are:</p> <ul style="list-style-type: none"> • As is – value of the last sample taken during the selected time range • Sum – summarized value of all samples taken during the selected time range • Average – average value of all samples taken during the selected time range
	Is Key	<p>Indicates whether the field is the query's key for retrieval.</p> <p>NOTE: Fields that are indicated as keys should correspond with the database result set unique values. Selecting a field which has frequently changing and repeating results as key is allowed, but may result in unexpected behavior from the defined collection.</p>
	Configured to use IntelliProfile	<p>Select True in this column to enable the collection of baseline statistics for this metric.</p>

Except for the *Column Name* and *Column Type* fields, whose values are retrieved by the query and cannot be changed, all other fields can be edited by clicking any of them.

The Edit Collection Properties dialog box appears, allowing you to edit the values of the following parameters:

- Display name
- Type
- Unit of measurement
- Aggregation

In addition, the collection's sampling frequencies are displayed on the table at the bottom of the screen, and can be edited by clicking any of them.

10 Ensure that all settings are appropriate, and click OK to finish the collection creation process.

The newly created collection now appears on the table.

- ① **IMPORTANT:** After adding the requested user-defined collections, they can only be deleted or cloned to other agents. If one or more queries need to be modified, delete them and create new ones.

Configuring the On-demand Data Port

Use the On-demand Data Port view in the Databases Administration dashboard to define a port for each Foglight Agent Manager to be used by for retrieving data and integrating with Performance Analysis.

To configure a port for the collection of on-demand data:

1 Select the check box near the names of the requested Foglight Agent Managers.

2 Click Set port.

The Set On-demand Data Port dialog box opens.

3 Type the name of the requested port.

4 Click Set.

5 Click Validate connectivity.

The Validate connectivity progress bar appears. At the end of the validation process, the connection status appears on the Status column.

If the connection failed, take the requested correction measures and try again.

- ① **IMPORTANT:** Changes to the On-demand Data Port view take effect immediately and do not need to be saved.

Changing the Monitoring Mode of the RAC One Node Agents

Use the RAC One Node view in the Databases Administration dashboard to change the monitoring mode of agents configured as RAC One Node, enabling the monitoring of such agents as regular RACs.

- ① **NOTE:** The RAC One Node view is only relevant for agents configured as RAC One Node. If no such agent was selected before entering the Databases Administration dashboard, this view will be empty.

To change the monitoring mode of a RAC One Node agent:

1 Click anywhere on the row of the RAC One Node agent (the root).

The *Convert to Regular RAC* dialog box appears, displaying one or more hosts.

2 Click the row of the RAC that is to be converted to regular RAC.

The Convert to Regular RAC dialog box opens.

3 From the list of hosts that appear in the dialog box, select the host that is to be converted to a physical RAC.

4 Click Validate to verify the selected host's connectivity.

5 After the connectivity validation completes successfully, click Set.

- ① **IMPORTANT:** Converting a RAC One Node to a regular RAC is an irreversible process. The Save changes button is disabled on the RAC One Node view, and after clicking Set, the selected host is immediately converted to a physical RAC. Therefore, ensure that the requested host is selected before clicking Set.

After the selected host was converted to a physical RAC, use the Oracle Monitoring Installer wizard to add the other hosts monitored through this agent as nodes to the newly created RAC. For further details, see [Using Foglight for Oracle](#) on page 26.

Administering SQL Performance Investigator


The SQL Performance Investigator view in the Administration dashboard allows you to enable and disable SQL PI monitoring for selected agents. In addition, you can start and stop the collection of change tracking data.

In the SQL PI view, select one or more agents to enable or disable or for which to modify the change tracking status.

Reporting

Generating Reports

Foglight *for Oracle* enables generating reports about various aspects of the selected instance or RAC performance. This chapter provides information on how to generate the various reports, as well as a brief description of each report.

 **NOTE:** For detailed information regarding the use and configuration of reports, see Foglight User Guide > Working with Reports.

Generating Reports for a Foglight *for Oracle* RAC or Instance

To generate reports for a selected RAC or instance:

- 1 Go to Dashboards > Reports.
- 2 Click Run a Report.
The Run Report dialog box opens.
- 3 On the first screen, *Select Template*, Go to All Oracle Templates.
- 4 Select the requested report.
- 5 Click Next.
- 6 Select the requested time range.
- 7 Select the instance or RAC for which the report will be generated.
- 8 Click Next.
- 9 Assign a name for the report.
- 10 Select the report format (PDF, Excel or XML).
- 11 Use the *Email Recipients* field to type the name of the report's recipients.
- 12 Click Finish.
The report generation process starts. Upon successful completion of this process, the Report Generated Confirmation dialog box opens.
- 13 Click Download Now to download the file.
- 14 Open the file using a relevant program for the selected format.

Studying the Various Reports

The reports that can be generated for Foglight *for Oracle* are described below. Except for the Cluster Balance report, all of the reports can be generated for either a RAC or an instance.

- Cluster Balance report – used for monitoring the load balance, during the selected time range, between the various nodes in the selected RAC.

- Databases Storage report – displays a breakdown of the database storage space usage at the host, instance and tablespace levels. If the monitored instance is part of a multitenant architecture, this report at the instance level includes a graph of pluggable database distribution by storage.
- Health Check report – used for monitoring various aspects of the instance’s or RAC’s health, that is: availability, listener status, response time and connection time.
- Host Disk Capacity Report – provides information about the file system disk of each of the monitored host, such as the file system type, the disk size, the amount (MB) of used and free space on the disk, and the percentage of free space.
- I/O Activity report – provides several I/O activity indicators, such as workload, physical and logical read and write operations, and wait events.
- Instance Availability Report – provides storage-related information about the instance, RAC node or cluster, such as availability percentage, listener status, connection time, and average response time.
- Oracle Executive Summary report – provides key decision-makers with a summary view of the Oracle activity. If the monitored instance is part of a multitenant architecture, this report displays a grid showing data for each pluggable database.
- Storage Summary report – provides various storage-related indicators, such as the status of datafiles and ASM, and the top 10 tablespaces sorted by space usage. If the monitored instance is part of a multitenant architecture, this report at the instance level includes a graph of pluggable database distribution by storage.
- Top DB Users report – displays the DB users that generated the highest amount of activity for a given time frame.
- Workload Summary Report – provides various workload indicators, such as the workload resource breakdown, CPU utilization and SQL executions.
- ASM Summary Report – Displays general ASM instance-related information including Disk Groups and their associated Disks.
- Database Backup Report – Provides an overview of RMAN configuration and database backups.
- Invalid Objects Report – Displays a summary of invalid objects.
- Locks Report – Displays lock waits chart (Average Active Sessions) and lock tree.
- Storage Planning Report – Displays a prediction of storage growth using a linear regression method, based on the available historical data (up to 3 months).
- Top Statements Report – Displays the top-consuming SQL statements that were executed during the specified time range. To exclude specific statements from the report, use the optional Exclude Text fields using a LIKE format. For example, specify %backup% to exclude all the statements that contain the word ‘backup’ as part of their text.

Reference

This chapter describes the collections that are used with monitored Oracle databases.

- [Overview of Collections](#)

Overview of Collections

Foglight *for Oracle* collects raw data for a set of collection types. This data is used by the Rules to trigger alarms and is used to populate dashboards. The collections data tables are not viewable from Foglight *for Oracle*, but the data is implemented in the various drilldowns of Foglight *for Oracle*.

The collection data tables are not viewable from Foglight *for Oracle*.

NOTE: Several collections contain metrics whose name ends with Rate (for example: DBO Wait Buffer Busy Rate). Such metrics, which are not documented in this guide, are used for plotting the original metric (usually called by the same name without the “Rate” suffix) on a graph for the selected time range.

Alert Log

Purpose

The Alert Log collection retrieves the content of Oracle alert log.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	900

Metric Descriptions

Metric	Description
Critical Errors	The number of critical error notifications that were received during the specified time range.
Dump File Name	Indicates the name of dump file (if this name appears in the alert log message).

Metric	Description
Fatal Errors	The number of fatal error notifications that were received during the specified time range.
Informational Errors	The number of informational messages that were received during the specified time range.
Message	The text of the alert log message.
Name	The name of the alert log message.
Severity	Indicates the alert log message severity (Fatal, Critical, Warning, Informational).
Timestamp	Alert log message timestamp.
Type	Alert log message type (Internal Error, Storage, Media Failure and so on).
Warning Errors	The number of warning error notifications that were received during the specified time range.

Archive Destination Miscellaneous

Purpose

The Archive Destination Miscellaneous collection provides Archive Destinations activity rates and time considerations.

Collection Type

Oracle and Operating System

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	3600
Online	3600
Offline	3600

Metric Descriptions

Metric	Description
Archive Destination Free Space	The free space in the file system where the most critical archive destination is located.
Archive Destination Size	The size of the file system where the most critical archive destination is located.
Archive Rate	The rate of the archiving operation.
Archiving Volume	Archiving volume
Average Archive File Size	The size of the average archive file

Metric	Description
Binding	Binding type: <ul style="list-style-type: none"> MANDATORY – successful archival is required. OPTIONAL – successful archival is not required.
Critical Archive Time To Failure	The number of hours left until the critical archive destination fills.
Database	The database name.
Destination	The destination ID of the most critical archive destination.
Destination directory	The archive's destination directory
Destination Free Space	The total amount (in megabytes) of free space available for the archive destination target.
Destination Total	The total size (in megabytes) of the archive destination target.
Time To Failure	The number of hours remaining until the archive destination with the shortest time to failure is filled.

Archive Logs Storage

Purpose

The Archive Logs Storage collection provides Archive Destinations properties and file systems space data.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Archive Destination	The name of the destination archive file.
Archiver Status	The archiver status; can have one of the following values: Stopped, Started, and Failed.
Binding	Binding type: <ul style="list-style-type: none"> MANDATORY – successful archival is required. OPTIONAL – successful archival is not required.
Database	Database name
Destination Directory	The directory used for storing the archive.

Metric	Description
Destination End	The value of destination directory is replaced by <i>Recovery Destination</i> if the directory is used for recovery.
Destination Free Space	The total amount (in megabytes) of free space available for the archive destination target.
Destination Total	The total size (in megabytes) of the archive destination target.
Destination Type	The type of the archived log destination: <ul style="list-style-type: none"> • PUBLIC • PRIVATE
Destination Usable	The amount (in megabytes) of storage space usable for the archive logs in the destination directory.
Error	Displays the error text.
If database in archive mode	Indicates whether the database is in archive mode
Is ASM	Indicates whether the instance is an Automatic Storage Management (ASM) instance
Is Recovery	Indicates whether the directory is used for recovery (True) or not (False).
Log Sequence	Identifies the sequence number of the most recently archived redo log.
Multiple destinations in filesystem	Indicates if more than one archive destinations reside on the same filesystem.
Re-open	Identifies the retry time (in seconds) to open a destination after error.
Space Reclaimable	The amount of space that can be used if files are removed from the Flash Recovery Area.
Status	The current status of the destination in V\$ARCHIVE_DEST: <ul style="list-style-type: none"> • VALID • INACTIVE • DEFERRED • ERROR • DISABLED • BAD PARAM • ALTERNATE • FULL
Target	Specifies whether the archive destination is local or remote to the primary database: <ul style="list-style-type: none"> • PRIMARY - local • STANDBY - remote
Transmit Mode	The mode in which data is transmitted across the network: <ul style="list-style-type: none"> • SYNCHRONOUS • ASYNCHRONOUS • PARALLELSYNC
Unarchived redo logs	The percentage of redo logs files that are waiting for archiving
Valid Type	The redo log type or types that are valid for the destination: <ul style="list-style-type: none"> • ONLINE_LOGFILE • STANDBY_LOGFILE • ALL_LOGFILES

ASM Diskgroups

Purpose

The ASM Diskgroups collection collects storage and activity data regarding ASM Diskgroups.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	900

Metric Descriptions

Metric	Description
Allocation Size Unit	The size of allocation units in the disk group.
ASM Metadata Block Size	The block size for ASM metadata in the disk group.
Average Service Time	The average service time for all disks in the group.
Group Name	The name of the disk group.
Group Number	The number assigned to an ASM disk group.
Group State	The state of the specified ASM disk group relative to the Oracle instance (CONNECTED, BROKEN, UNKNOWN, DISMOUNTED, or MOUNTED).
Number of Disks	The number of disks in the disk group.
Offline Disks	The number of disks in the disk group that are offline. NOTE: This information is unavailable for Oracle 10.1
Read Errors	The total number of failed I/O read requests for the disk group.
Read IO Requests	The total number of I/O read requests for the disk group.
Total Group Size	The total capacity of the specified ASM disk group (in megabytes).
Usable Group Space	The unused space on the disk group (in megabytes). NOTE: This information is unavailable for Oracle 10.1
Used Group Space	The percentage of used disk space on the entire disk group. This value is calculated as follows: $100 * (\text{Total Group Size} - \text{Usable Group Space}) / \text{Total Group Size}$
Write Errors	The total number of failed I/O write requests for the disk group.
Write IO Requests	The total number of I/O write requests for the disk group.

ASM Diskgroups Balance

Purpose

The ASM Diskgroups Balance collection estimates the rate of balance of IO Activity and Response time of ASM disks on ASM diskgroups.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	900

Metric Descriptions

Metric	Description
Balanced IO	The extent, in percents, to which the I/O activity is balanced across the ASM diskgroups. NOTE: This information is unavailable for Oracle 10.1
Balanced Time	The extent, in percents, to which the I/O time is balanced across the ASM diskgroups. NOTE: This information is unavailable for Oracle 10.1
Group Number	The number assigned to an ASM disk group.

ASM Disks

Purpose

The ASM Disks collection collects storage and activity data regarding ASM disks.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	900

Metric Descriptions

Metric	Description
Average Service Time	The service time for each disk. Service time is the response time of a disk, plus any wait time experienced during disk operations.
Create Date	The date and time when the disk was added to the disk group.
Disk Free Space	The unused capacity of the ASM disk.
Disk Name	The name of the disk.
Disk Number	The number assigned to an ASM disk within its disk group.
Disk Path	The disk's path, as indicated on the operating-system.
Disk Size	The total capacity of the ASM disk.
Disk State	The state of the disk, regarding its ASM disk group (UNKNOWN, NORMAL, ADDING, DROPPING, HUNG, FORCING, or DROPPED).
Fail Group	The name of the failure group that contains the disk.
Group Number	The number assigned to an ASM disk group.
Header Status	The status of the disk at the discovery stage (UNKNOWN, CANDIDATE, INCOMPATIBLE, PROVISIONED, MEMBER, FORMER, CONFLICT, or FOREIGN).
Label	The label portion of the disk name at the discovery stage.
Mode Status	The access type granted to the disk (UNKNOWN, ONLINE, or OFFLINE). Read/write access is only granted to ONLINE disks; OFFLINE disks still remain part of their disk group.
Mount Date	The date and time when the disk was first mounted.
Mount Status	The mount status of the disk (MISSING, CLOSED, OPENED, CACHED, IGNORED, or CLOSING).
Read Errors	The total number of failed I/O read requests for the disk. Such failures can result from a loss of access to the entire disk, or media corruptions on a generally healthy disk.
Read Time	The total time spent on I/O read requests.
Redundancy	The external redundancy of the disk, if the ASM mechanism does not provide redundancy. The possible states are: UNKNOWN, UNPROT, MIRROR, or PARITY.
Total IO Percent	The number of I/O operations performed by the selected disk as a percentage of all I/O performed by all disks in all ASM disk groups assigned to the Oracle instance.
Total Reads	The total number of I/O read requests for the disk.
Total Size Percent	The size of the disk as a percentage of the size of all disks in all ASM disk groups assigned to the Oracle instance.
Total Writes	The total number of I/O write requests for the disk.
Used Space Pct	The percent of used space out of the entire space allocation on the ASM disk.
Write Errors	The total number of failed I/O write requests for the disk.
Write Time	The total time spent on I/O write requests.

ASM Instance Total

Purpose

The ASM Instance Total collection provides ASM-related data, summarized for the database instance.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	900

Metric Descriptions

Metric	Description
Total Disk Space Allocated	The total capacity allocated to all ASM disks residing on ASM groups that are used by this database instance.
Total Disks Free Percent	The total free space percentage on all ASM disks residing on ASM groups that are used by this database instance.
Total IO	The total number of I/O operations (reads and writes) on ASM managed disks.
Total Reads	The total number of I/O read requests for ASM managed disks.
Total Writes	The total number of I/O writes to ASM managed disks.

ASM Operations

Purpose

The ASM Operations collection retrieves details regarding the currently running ASM operations.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	900

Metric Descriptions

Metric	Description
Actual Power	The power actually allocated to the operation.
Estimated Minutes	The estimated amount of time (in minutes) that the operation is expected to take.
Estimated Rate	The estimated rate (%) at which allocation units will be moved by the operation
Estimated Work	The estimated number of allocation units that the operation needs to move.
Group Number	The number of the disk group where the operation is running.
Operation	The type of the operation; REBAL, COD, DSCV, ACD, RFRSH, or EXPEL.
Operation State	The state of the operation; WAIT, RUN, REAP, HALT, or ERRORS.
Requested Power	The power requested for the operation as specified by the ASM_POWER_LIMIT parameter. Higher value of ASM_POWER_LIMIT indicate faster operation, while lower values indicate slower operation that uses fewer resources.
So Far Moved Units Number	The number of allocation units moved so far by the operation.

Backup Config

Purpose

The Backup Config collection provides a list of configuration parameters value for RMAN.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	21600
Online	21600
Offline	21600

Metric Descriptions

Metric	Description
Configuration #	The number of configuration parameters
Configuration Name	The names of the configuration parameters
Configuration Value	The values of the configuration parameters
Database	Database name

Backup Jobs

Purpose

The Backup Jobs collection collects RMAN jobs statuses.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	21600
Online	21600
Offline	21600

Metric Descriptions

Metric	Description
Backup Job ID	The ID of the backup job
Backup Job Status	The status of the backup job
Backup Job Type	The type of the backup job
Database	Database name
Job End Time	The end time of the backup job
Job Start Time	The start time of the backup job
Output	The output message of the backup job
Session Rec ID	The Oracle identifier of the RMAN job
Session Stamp	The time stamp of the session

Backup Type Status

Purpose

Backup Type Status collection collects the last time stamp of each RMAN job that ran successfully.

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	21600
Online	21600
Offline	21600

Metric Descriptions

Metric	Description
Backup Type	The type of the backup job.
Days Since Backup	The number of days that passed since the last database backup.
Last Successful Backup	The date of the most recent successful backup that has been carried out for the database.

Cluster Latency

Purpose

The Cluster Latency collection provides Cluster Latency and Activity data.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Average Global Cache Convert Time	Average amount of time required for the completion of global conversions on GCS resources.
Average Global Lock Convert Time	Average amount of time required for the completion of synchronous and asynchronous global lock converts.
Average Global Lock Get Time	Average amount of time required for the completion of synchronous and asynchronous global lock gets.
Balance	Interconnect Load balance between the various cluster nodes.
Block Received Latency	The average transfer time of a block from another instance.
Block Sent Latency	The average time needed to prepare a requested block.
Blocks Corrupt Percent	The percentage of global cache blocks that encountered a corruption or checksum failure within the entire number of global cache blocks sent over the Interconnect.
Blocks from Interconnect	The number of blocks received from Interconnect after the requests for blocks were not satisfied from the local instance, and therefore had to make a request for the block from the remote.
Blocks to Interconnect	The number of blocks sent via Interconnect after the requests for blocks were not satisfied on the requesting instance, and therefore had to make a request for the block from the local instance.

Metric	Description
Blocks/sec from Interconnect	The rate per second of blocks received from Interconnect after the requests for blocks were not satisfied from the local instance, and therefore had to make a request for the block from the remote.
Blocks/sec to Interconnect	The rate per second of blocks sent via Interconnect after the requests for blocks were not satisfied on the requesting instance, and therefore had to make a request for the block from the local instance.
Cluster Latency Rate	The average time it took to transfer a block through the Interconnect, that is, the time that has elapsed since the request for the block was submitted until the block was received by the requesting instance.
Consistent Gets	The number of times a consistent read was requested for a block by this session
CR Blocks Send Time	Total amount of time spent sending constructed CR Blocks through the Interconnect.
Current Block Send Time	Total amount of time spent sending Current Blocks through the Interconnect.
DB Block Gets	The number of data block reads in CURRENT mode (namely, current reads).
GC Blocks Corrupt	The total number of global cache blocks that encountered a corruption or checksum failure during the specified time range. As such failures occur during data transfer, high values of this metric can indicate network-related issues.
GC Blocks Lost	Total number of global cache blocks lost during the specified time range. As the loss of global cache blocks takes place during data transfer, high values of this metric can indicate network-related issues.
GC CR Block Build Time	The average time taken by an LMS process to create a consistent read block on a holding instance within the cluster.
GC CR Block Flush Time	The average time taken by the LGWR (log writer) process to flush any changes required to a buffer on the log flush queue.
GC CR Block Receive Time	Total amount of time foreground processes spent waiting for a CR block to be sent through the Interconnect.
GC CR Block Send Time	The average time taken by an LMS process to initiate the sending of a consistent read block. This is NOT the same as the time elapsed before the block arrives at the requesting instance.
GC CR Blocks Received	Total number of consistent read data blocks that Oracle processes did not find in their local cache and had to fetch from the global cache, using a consistent read request.
GC CR Blocks Served	Total number of blocks constructed by the Block Server Process (BSP).
GC Current Block Flush Time	The average time taken by the LGWR (log writer) process to flush the changes from a block to disk (a forced log flush) before the block is shipped to the requesting instance.
GC Current Block Pin Time	The average time taken by an LMS process to pin the current block before shipping it to the requesting instance.
GC Current Block Receive Time	Total amount of time foreground processes waited for a current block to be sent through the Interconnect.
GC Current Block Send Time	The average time taken by an LMS process to initiate the sending of a current block. NOTE: This metric is not identical to the time elapsed before the block arrives at the requesting instance.
GC Current Blocks Received	Total amount of time foreground processes waited for a current block to be sent through the Interconnect.

Metric	Description
GC Current Blocks Served	The number of current blocks shipped to the requesting instance over the hardware Interconnect.
Global Cache Convert Time	Total amount of time spent waiting for global cache converts on all sessions.
Global Cache Converts	Resource converts of buffer cache blocks. This statistic is incremented whenever GCS resources are converted from Null to Exclusive, Shared to Exclusive, or Null to Shared.
Global Lock Convert Time	Total time spent waiting for all synchronous and asynchronous global lock converts.
Global Lock Get Time	Total elapsed time of all synchronous and asynchronous global lock gets.
Global Lock Releases	Total number of synchronous global lock releases.
Global Lock Sync Converts	The total number of resources that Oracle converted from an incompatible mode during the selected time range.
Interconnect Block Latency	The average time it takes to prepare and transfer a block across the Interconnect network.
Lost Blocks Ratio	The ratio of global cache blocks losses during data transfer via the Interconnect.
Received from Interconnect Time	Total amount of time the foreground processes waited for blocks to be sent through the Interconnect
Sent to Interconnect Time	Total amount of time spent sending blocks through the Interconnect
Total CR Blocks	The total number of CR blocks that were served/received through the Interconnect
Total Current Blocks	The total number of current blocks that were served/received through the Interconnect.
Total I/C Load	Total number of blocks (both CR and current) that were served and received from the Interconnect.

Cluster Misses

Purpose

The Cluster Misses collection collects data regarding global cache misses.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
2-way Misses	The number of current and consistent blocks that were received by the selected instance from a specific instance, after a 2-way roundtrip.
3-way Misses	The number of current and consistent blocks that were received by the selected instance from a specific instance, after a 3-way roundtrip.
Consistent Blocks 2-way Misses	The number of consistent blocks that were received by the selected instance from a specific instance, after a 2-way roundtrip.
Consistent Blocks 3-way Misses	The number of consistent blocks that were received by the selected instance from a specific instance, after a 3-way roundtrip.
Consistent Blocks Misses	The number of consistent blocks that were received by the selected instance from a specific instance, after a 2-way or 3-way roundtrip.
Current Blocks 2-way Misses	The number of current blocks that were received by the selected instance from a specific instance, after a 2-way roundtrip.
Current Blocks 3-way Misses	.The number of current blocks that were received by the selected instance from a specific instance, after a 3-way roundtrip
Current Blocks Misses	The number of current blocks that were received by the selected instance from a specific instance, after a 2-way or 3-way roundtrip.

Datafile Fragmentation

Purpose

The Datafile Fragmentation collection collects information regarding the layout of the free space left in the datafile due to fragmentation.

Collection Type

Oracle

Collection Sampling Settings

None (on-demand collection).

Metric Descriptions

Metric	Description
Datafile name	The name of the datafile
Extent Size	Offset end of the extent from the beginning of the file.
Extent Start	Offset start of the extent from the beginning of the file.
Extent Type	Defines the type of the extent: Used, Fully Usable, Partly Usable or Unusable.
Sequence ID	The datafile's sequence ID.

Datafile IO Activity

Purpose

The Datafile IO Activity collection collects metrics and performance ratios related to datafile I/O activity.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Avg IO Time	Average time spent on a single IO request for a specific datafile
Avg Read Time	Average time spent on a single physical read request for a specific datafile
Avg Write Time	Average time spent on a single physical write request for a specific datafile
Contents	The type of the tablespace contents (temporary, permanent, undo)
File Full Name	The name of the datafile (including full path)
File ID	The datafile ID
File Name	The name of the datafile (without the path)
File Path	The datafile path
IO Time	The total time spent on carrying out IO operations (reading and writing data) on a specific datafile
Last IO Time	The time spent while performing last I/O operation on a specific datafile
Max Single Read Time	The maximum time spent while reading a single block from a specific datafile
Max Single Write Time	The maximum time spent while writing a single block to a specific datafile
Min Single IO Time	The minimum time spent while reading a single block from a specific datafile
PDB	In multitenant architectures, the name of the pluggable database.
Physical IO	Total size, in blocks, of all disk reads and writes for a specific datafile
Physical IO Percent	The IO activity (block reads + block writes) for a specific datafile as a percentage of the total IO activity for all datafiles
Physical Reads	The number of data blocks read by Oracle, which are not present in the buffer cache or in the PGA, thereby requiring I/O read operations, that is: the total number of data blocks read from disk.

Metric	Description
Physical Reads Percent	The block reads carried out for a specific datafile as a percentage of the of total block reads carried out for all datafiles
Physical Writes	Total size in blocks of all disk writes for a specific datafile
Physical Writes Percent	The block writes carried out for a specific datafile as a percentage of the of total block writes carried out for all datafiles
Read Requests	The number of read requests for a specific datafile that read one or more database blocks per request.
Read Time	The total time spent on reading data from a specific datafile.
Single Block Read Time	The total time spent on read requests that read one database block per request from a specific datafile.
Single Block Reads	Number of read requests for a specific datafile, which read one database block per request.
Tablespace Name	The tablespace name.
Write Requests	The number of write requests for a specific datafile that wrote one or more database blocks per request.
Write Time	The total time spent on writing data to specific datafile.

Datafiles Storage

Purpose

The Datafiles Storage collection provides a list of all datafiles and their overall properties and storage details.

Collection Type

Oracle and Operating System

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	900
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Autoextend	Indicates whether the datafile is autoextensible.
Autoextensible	Indicates if at least one datafile in the database is not autoextensible
File ID	The ID of the selected datafile.
File Name	The name of the datafile.
Filesystem Name	The name of the file system where the datafile resides.
Free	The amount of free space on the selected datafile.

Metric	Description
Free Pct	The percentage of free space on the selected datafile.
Is ASM File	Indicates whether the datafile is an Automatic Storage Management (ASM) file.
Maximum Used Percent	The maximum percent of the used disk space of all the database datafiles.
Maxsize	The maximum size of this datafile; used for autoextensible datafiles.
PDB	In multitenant architectures, the name of the pluggable database.
Size	The size of the datafile (in megabytes).
Status	Indicates whether the datafile is available or invalid.
Tablespace Name	The name of the table space to which the datafile belongs.
Total Allocated Space	The total allocated space in the database.
Total Allocated Space Delta	
Total Free Space Percent	The percentage of free space in the database datafiles.
Used	The size (in megabytes) of used datafile storage.
Used Pct	The percentage of used datafile storage.

Dataguard Destination Status

Purpose

The Dataguard Destination Status collection retrieves status and statistics for standby destinations in Data Guard.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	300

Metric Descriptions

Metric	Description
Destination Directory	The dataguard destination directory
Destination Error	Displays the error message for destination status.
Destination ID	The destination ID

Metric	Description
Destination Status	Identifies the current status of the destination: <ul style="list-style-type: none"> • VALID – initialized and available • INACTIVE – no destination information • DEFERRED – manually disabled by the user • ERROR – error during open or copy • DISABLED – disabled after error • BAD PARAM – parameter has errors • ALTERNATE – destination is in an alternate state • FULL – exceeded quota size for the destination
Gap in Applied Logs	The gap between the last archive log on the primary database and last applied archive log on the standby database.
Gap in Received Logs	The gap between the last archive log on the primary database and the last received archive log on standby database.
Last Applied Sequence#	Number of the last applied sequence
Last Dataguard error	The most recent error reported by Oracle for this standby destination
Last Dataguard message	The most recent DataGuard message reported by Oracle for this standby destination.
Last Received Sequence#	Number of the last received sequence.

Dataguard Redo Apply

Purpose

The Dataguard Redo Apply Rates collection retrieves performance data for database configured with Data Guard.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	3600
Online	3600
Offline	3600

Metric Descriptions

Metric	Description
Destination ID	The destination ID
Rate of Redo Apply	The amount of MB of the redo logs applied to the standby database.
Rate of Redo Receive	The amount of MB of the redo logs received on the standby database.

Dataguard Status

Purpose

The Dataguard Status collection retrieves general definitions and status of database configured with Data Guard.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	300

Metric Descriptions

Metric	Description
Database Role	The database role
Number of Logs not Applied on Standby	Number of logs that were not applied on standby mode.
Number of Logs not Received on Standby	Number of logs that were not received on standby mode.
Number of Online Redo Logs	The number of Online redo logs defined for the standby database.
Number of Standby Redo Logs	The number of Standby redo logs defined for the standby database.
Protection Level	Aggregated protection mode currently in effect for the database: <ul style="list-style-type: none">• MAXIMUM PROTECTION – Database is running in maximized protection mode• MAXIMUM AVAILABILITY – Database is running in maximized availability mode• RESYNCHRONIZATION – Database is running in resynchronization mode• MAXIMUM PERFORMANCE – Database is running in maximized protection mode• UNPROTECTED – Database is unprotected (this normally occurs when the primary database is mounted and not open)

Metric	Description
Protection Mode	Protection mode currently in effect for the database: <ul style="list-style-type: none"> • MAXIMUM PROTECTION – Database is running in maximized protection mode • MAXIMUM AVAILABILITY – Database is running in maximized availability mode • RESYNCHRONIZATION – Database is running in resynchronization mode • MAXIMUM PERFORMANCE – Database is running in maximized protection mode • UNPROTECTED – Database is unprotected (this normally occurs when the primary database is mounted and not open)
Redo Logs Configuration Mismatch	The number of online redo logs and the number of standby redo logs on the primary database do not match.
Total Number of Dataguard Errors	The number of DataGuard errors reported by Oracle for the standby database.

Dump File

Purpose

The Dump File collection retrieves the contents of the Oracle dump file.

Collection Type

Oracle

Collection Sampling Settings

None (on-demand collection).

Metric Descriptions

Metric	Description
Dump File Name	The name of the dump file
Message	The dump file message

Event Wait

Purpose

The Event Wait collection retrieves the instance wait events (workload) that took place within the specified time range.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Event ID	The wait event's ID
Event Name	The wait event's name
PDB	In multitenant architectures, the name of the pluggable database.
Wait Time	The elapsed time of the selected wait event, plotted over time

Exadata Cell Flashcache

Purpose

The Exadata Cell Flashcache collection retrieves performance, utilization and usage statistics about the Exadata cell flashcache.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Cell ID	The Cell ID, provided by the hardware vendor.
Flashcache hit ratio	The ratio, within the total physical disk requests, of physical disks requests that were resolved from the Flashcache, thereby saving the need to carry out physical disk reads.
Flashcache I/O Errors	The number of IO errors that were encountered on the Flashcache
Flashcache Keep read request misses (Keep Objects)	The number of read IO requests for 'keep' objects that did not find all data in Flashcache and therefore had to perform physical read operations.

Metric	Description
Flashcache Keep read requests (Total)	The number of read IO requests for 'keep' objects that were resolved from the Flashcache, thereby saving the need to carry out physical disk reads.
Flashcache read request misses (Total)	The total number of read IO requests that did not find all the requested data in the Flashcache, and therefore had to perform read from disks.
Flashcache read requests (Keep objects)	The number of read IO requests that were resolved from the Flashcache, thereby saving the need to carry out physical disk reads.
Flashcache usage (Keep objects)	The amount of space (number of megabytes) used for 'keep' objects on the Flashcache.
Flashcache usage (total)	The number of megabytes used on the Flashcache.
MB of Flashcache misses	The total amount of data (number of megabytes) read from disks because not all requested data was found in the Flashcache.
MB Pushed out of the Flashcache for Keep objects	The amount (megabytes) of data pushed out of the Flashcache because the space limit for storing 'keep' objects has been exceeded.
MB Pushed out of the Flashcache for Keep objects rate	The rate (megabytes/second) at which data was pushed out of the Flashcache because the space limit for storing 'keep' objects has been exceeded.
MB read from Flashcache	The total amount of data (number of megabytes) from the Flashcache during the specified time range.
MB written to Flashcache	The total amount of data (number of megabytes) written to the Flashcache during the specified time range.

Exadata Cell General

Purpose

The Exadata Cell General collection retrieves general information and status about the Exadata cell.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
BMC Type	Name and version of the Baseboard Management Controller (BMC).
Cell ID	The Cell ID, provided by the hardware vendor.
Cell Model	The brand and model of the cell hardware, as provided by the vendor.

Metric	Description
Cell Name	The Exadata cell name.
Cell Offload Efficiency	The cell's offload efficiency. Cell Offload Processing is a procedure used by the oracle database server to push query predicates (that limit data retrieved by the query using filters and join conditions) to the Exadata cell.
Cell Status	The Exadata cell status.
CPUs	The total number of CPU units on the cell.
Interconnects	The total number of network interconnection cards (NICs).
Kernel Version	The host kernel's software version.
OS Version	The Exadata cell's OS version.
Power Supplies (Working/ Total)	The total number of power supplies on the Exadata cell, displayed as working/total.
Temperature	The server's temperature (Celsius degrees), as provided by the Baseboard Management Controller (BMC).
Temperature Status	The server's temperature status, as provided by the Baseboard Management Controller (BMC).
Total Fans	The total number of fans on the Exadata cell.
Up Since	Calculates the last cell restart or start time, based on the cell_uptime
Uptime	The time period that has elapsed since the cell was last restarted.
Working Fans	The total number of working fans on the Exadata cell.

Exadata Cell InfiniBand

Purpose

The Exadata Cell InfiniBand Nodes collection retrieves information and status of the cell InfiniBand nodes.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
HCA	The InfiniBand node's host channel adapter (HCA)
Node IP Address	The InfiniBand node's IP address.

Metric	Description
Node Name	The name of the InfiniBand node.
Node Status	The InfiniBand node's status.

Exadata Cell InfiniBand Ports

Purpose

The Exadata Cell InfiniBand Ports collection retrieves the cell InfiniBand ports' information and status.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Link Speed Active	The actual link speed of the Exadata cell InfiniBand port, as retrieved in the last sample.
Link Speed Enabled	The link speed enabled by the user, as retrieved in the last sample.
Link Speed Supported	The supported link speed of the Exadata cell InfiniBand port, as retrieved in the last sample.
Link State	The state of the InfiniBand port link of the Exadata cell.
Link Width Active	The actual link width of the Exadata cell Infiniband port, as retrieved in the last sample.
Link Width Enabled	The link width enabled by the user, as retrieved in the last sample.
Link Width Supported	The supported link width of the Exadata cell InfiniBand port, as retrieved in the last sample.
Node Name	The name of the InfiniBand node.
Physical Link State	The physical state of the Exadata cell InfiniBand port's link.
Port Number	The InfiniBand port number.
Switch LID	The InfiniBand Switch's local identifier (LID), as assigned by the Subnet Manager.

Exadata Cell InfiniBand Switches

Purpose

The Exadata Cell InfiniBand Switches collection retrieves the cell InfiniBand ports' information and status.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Health Details	Retrieves the Infiniband health check detailed results of the parameters checked in the environment test.
Health Status	Retrieves the Infiniband Switch's current health status, as indicated by the success or failure of the environment test.
Link Status	The Infiniband Switch's current link status, as indicated during the last sample.
Subnet Manager State	The Subnet Manager's current state on the Infiniband switch.
Switch LID	The InfiniBand Switch's local identifier (LID), as assigned by the Subnet Manager.
Switch Name	The Infiniband switch's name.

Exadata Cell Statistics

Purpose

The Exadata Cell Statistics collection retrieves load, throughput and utilization statistics about the Exadata cell.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60

Frequency Mode	Collection Interval (Seconds)
Online	300
Offline	900

Metric Descriptions

Metric	Description
Cell ID	The Cell ID, provided by the hardware vendor.
CPU Busy Percent	The percentage of time within the previous minute during which the system CPU units were not idle.
KB Received by Ethernet Interfaces Rate	The rate (kilobytes/second) at which data was received by the Ethernet interfaces during the last sample.
KB Transmitted by Ethernet Interfaces Rate	The rate (kilobytes/second) at which data was transmitted by the Ethernet interfaces during the last sample.
MB Received by InfiniBand Interfaces Rate	The rate (megabytes/second) at which data was received by the InfiniBand interfaces during the last sample.
MB Transmitted by InfiniBand Interfaces Rate	The rate (megabytes/second) at which data was transmitted by the InfiniBand interfaces during the last sample.
Physical Memory Utilization	The average percentage of physical memory used on the cell during the last sample.
Run Queue Length	The average number (over the preceding minute) of processes in the Linux run queue that were marked as running or uninterruptable. This reading is taken from <code>/proc/loadavg</code> .
Swap Utilization	The percentage of swap used.

Exadata Cell General

Purpose

The Exadata Cell General collection retrieves summarized cell disks performance and throughput statistics.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Average Small Read Latency Per Request	Average latency in milliseconds to read small blocks per request on all the cell disks.
Average Small Write Latency Per Request	Average latency in milliseconds to write small blocks per request on all the cell disks.
Cell ID	The Cell ID, provided by the hardware vendor.
MB Read Rate	The rate (megabytes/second) at which data was read from all of the cell disks during the specified time range.
MB Written Rate	The rate (megabytes/second) at which data was written on all of the cell disks during the specified time range.
Read Requests Rate	Number of requests to read large and small blocks per second on all the cell disks.
Write Requests Rate	Number of requests to write large and small blocks per second on all the cell disks.

Filesystem Disk

Purpose

The Filesystem Disk collection provides storage definitions regarding the filesystem disk where datafiles, dump files or archive files are located.

Collection Type

Operating System

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Filesystem size	The size of the filesystem
Filesystem free %	The percentage of the free disk space on the filesystem.
Filesystem Free Space	The amount (in megabytes) of free disk space on the filesystem
Filesystem Name	The name of the filesystem
Filesystem Type	The type of filesystem (datafiles, dump files or mixed files)
Long Term Estimated Till Full	Estimated number of hours until the filesystem is full, based on the long-term growth rate
Long Term Growth Rate	The estimated long-term growth rate of the filesystem

Metric	Description
Mounted On	Mount point name.
Operating System	The name of the operating system where this instance runs
Short Term Estimated Till Full	Estimated number of hours until the filesystem is full, based on the short-term growth rate
Short Term Growth Rate	The filesystem's estimated short-term growth rate
Used	The amount of used disk space on the filesystem
Used %	The percentage of the used disk space on the filesystem
Used Growth	The growth (in megabytes) of used disk space on the filesystem

Flashback Area

Purpose

If the flash recovery area feature was used for creating a location on disk that enables the database to create various backup and recovery-related files, the Flashback Area collection provides information about this location.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	300
Offline	900

Metric Descriptions

.....

Metric	Description
File Type	Indicates the types of the files stored on the flashback area: <ul style="list-style-type: none"> ONLINELOG FLASHBACKLOG CONTROLFILE ARCHIVELOG IMAGECOPY BACKUPPIECE
Non Reclaimable Space Percent	The percentage, within the flashback area, of non-reclaimable space (space currently in use, which cannot be reclaimed).
Number of Files	Indicates the number of files currently stored on the flashback area.
Used Space Percent	The percentage of used space within the entire disk space allocated to the flashback area.

History Locks Tree Instance

Purpose

The History Locks Tree Instance collection collects snapshots of instance-level locks that reached or exceeded the minimal duration, in seconds, required for collecting the lock. All locks that meet this criterion are displayed in the lower section of the Locks Historical Tree pane.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	60

Metric Descriptions

Metric	Description
Blocked SQL	The SQL statement that waits for lock release
Blocked SQL Hash	The hash value of the statement waiting for lock
Blocking Serial	Serial number of the locking session
Blocking SID	A SID that holds locks
Client Machine	The name of the computer where the client is running
Cross Instance	Indicates whether the global lock and the locking and blocked sessions are running on different RAC nodes (True) or not (False)
DB User Name	The Oracle database account of the started session
Host	The name of the host where the instance is installed
Instance	Instance name of the session in the lock tree
Lock Duration	The number of seconds the lock is running
Lock Event	The wait event on which the blocked statement is waiting
Lock Mode	The lock mode in which the session holds the lock. The lock mode can have one of the following values: <ul style="list-style-type: none">• 0 - none• 1 - null (NULL)• 2 - row-S (SS)• 3 - row-X (SX)• 4 - share (S)• 5 - S/Row-X (SSX)• 6 - exclusive (X)
Lock Request Mode	The lock mode in which the process requests the lock. The lock request mode has the same possible values as the lock mode.

Metric	Description
Lock Type	The type of user or system lock. The locks on the user types are obtained by user applications. Any process that is blocking others is likely to be holding one of these locks. The user type locks are as follows: <ul style="list-style-type: none"> • TM - DML enqueue • TX - Transaction enqueue • UL - User supplied
Object Wanted	The object on which the statement is blocked
OS User Name	The operating system user for the Client
PDB	In multitenant architectures, the name of the pluggable database.
PID	The operating system process ID for the Oracle server process
Program	The server program name
Serial	The session's serial number. The combination of SID and serial number provides a unique identification of the session.
SID	The session identification number

History Locks Tree RAC

Purpose

The History Locks Tree RAC collection collects snapshot of all locks within the RAC that reached or exceeded the minimal duration, in seconds, required for collecting the lock. All locks that meet this criterion are displayed in the lower section of the Locks Historical Tree pane

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	60

Metric Descriptions

.....

Metric	Description
Blocked SQL	The SQL statement waiting for lock release
Blocked SQL Hash	The hash value of the statement waiting for lock
Blocking Serial	Serial number of the locking session
Blocking SID	A SID that holds locks
Client Machine	The name of the computer where the client is running

Metric	Description
Cross Instance	Indicates whether the global lock and the locking and blocked sessions are running on different RAC nodes (True) or not (False)
DB User Name	The Oracle database account of the started session
Host	The name of the host where the instance is installed
Instance	Instance name of the session in the lock tree
Lock Duration	The number of seconds the lock is running
Lock Event	The wait event on which the blocked statement is waiting
Lock Mode	The lock mode in which the session holds the lock. The lock mode can have one of the following values: <ul style="list-style-type: none"> • 0 - none • 1 - null (NULL) • 2 - row-S (SS) • 3 - row-X (SX) • 4 - share (S) • 5 - S/Row-X (SSX) • 6 - exclusive (X)
Lock Request Mode	The lock mode in which the process requests the lock. The lock request mode has the same possible values as the lock mode.
Lock Type	The type of user or system lock. The locks on the user types are obtained by user applications. Any process that is blocking others is likely to be holding one of these locks. The user type locks are as follows: <ul style="list-style-type: none"> • TM - DML enqueue • TX - Transaction enqueue • UL - User supplied
Object Wanted	The object on which the statement is blocked
OS User Name	The operating system user for the Client
PDB	In multitenant architectures, the name of the pluggable database.
PID	The operating system process ID for the Oracle server process
Program	The server program name
Serial	The session's serial number, which, in combination with SID, provides a unique identification of the session.
SID	The session identification number

Host Properties

Purpose

The Host Properties collection provides operating system version, CPU, memory and network host properties.

Collection Type

Operating System

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	86400

Metric Descriptions

Metric	Description
Machine Model	The model of the host machine (for example: Sun Fire V240)
Number of CPU	The number of CPU units installed on the host
Number of IO Devices	The number of IO devices configured on the host
Number of Network Interfaces	Number of network interfaces installed on the host
Operating System	The name of the operating system installed on the host
Operating System Version	The version of the operating system installed on the host
Physical Hostname	The host name, defined per network interface
RAM	The amount of RAM available on the physical host

Initialization Parameters

Purpose

The Initialization Parameters collection collects definitions and change tracking history of Oracle's initialization parameters.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	900
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Current Value	Parameter value for the session (if the value was modified within the session); otherwise, the instance-wide parameter value.
Description	The parameter description

Metric	Description
Instance Level	Indicates whether the parameter can have a different value for each instance (TRUE) or must retain the same value for all RAC nodes (FALSE).
Is Default?	Indicates whether the parameter's value is set to the default value (TRUE) or was modified (FALSE)
Is Deprecated?	Indicates whether the parameter has been deprecated (TRUE) or not (FALSE)
Is Dynamic?	Indicates whether the parameter can be changed with ALTER SESSION (TRUE) or not (FALSE)
Is Modified?	Indicates whether the parameter has been modified after instance startup: <ul style="list-style-type: none"> MODIFIED – the parameter has been modified with ALTER SESSION SYSTEM_MOD – the parameter has been modified with ALTER SYSTEM, thereby modifying all of the currently logged-in session values FALSE – the parameter has not been modified after instance startup
Is System Modifiable?	Indicates whether the parameter can be changed with ALTER SYSTEM and when the change takes effect. This metric can have one of the following values: <ul style="list-style-type: none"> IMMEDIATE – the parameter can be changed with ALTER SYSTEM regardless of the type of parameter file used for starting the instance. The change takes effect immediately. DEFERRED – the parameter can be changed with ALTER SYSTEM regardless of the type of parameter file used for starting the instance. The change takes effect in subsequent sessions. FALSE – the parameter cannot be changed with ALTER SYSTEM unless a server parameter file was used for starting the instance. The change takes effect in subsequent instances.
Parameter	Initialization parameter name
PDB	In multitenant architectures, the name of the pluggable database.
Show Hidden Parameters	Indicates whether the parameter is hidden (TRUE) or not (FALSE)
Type	Indicates the parameter type. The possible values are as follows: <ul style="list-style-type: none"> Boolean String Integer Parameter file Reserved Big integer

Instance General

Purpose

The Instance General collection collects general instance data such as version and archive logging.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	3600

Metric Descriptions

Metric	Description
DB Version	The database version.
Is Archived	Indicates whether the instance is configured in archive mode.
Is Connected to ASM	Indicates whether the instance is an Automatic Storage Management (ASM) instance.
Is Data Guard Configured	Indicates whether DataGuard is configured on this instance.
Is Exadata	Indicates if this instance is part of Exadata Architecture.
Is Flashback Configured	Indicates whether the flashback recovery is configured.
Is RMAN Configured	Indicates whether RMAN is configured on the database for backup and recovery.

Invalid Objects

Purpose

The Invalid Objects collection collects data about invalid objects in each schema.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Invalid Functions	Number of invalid functions in the schema
Invalid Package Bodies	Number of invalid package bodies in the schema.
Invalid Packages	Number of invalid packages in the schema.
Invalid Procedures	Number of invalid procedures in the schema.

Metric	Description
Invalid Types	Number of invalid types in the schema.
Invalid Views	Number of invalid views in the schema.
New Invalid Functions	The number of invalid functions that are found to be present in the schema, and were not detected during the last run.
New Invalid Package Bodies	Number of invalid package bodies that were added to the schema during the last sampling period.
New Invalid Packages	Number of invalid packages that were added to the schema during the last sampling period.
New Invalid Procedures	Number of invalid procedures that were added to the schema during the last sampling period.
New Invalid Types	Number of invalid types that were added to the schema during the last sampling period.
New Invalid Views	Number of invalid views that were added to the schema during the last sampling period.
Others	Number of invalid objects of other object types in the schema.
PDB	In multitenant architectures, the name of the pluggable database.
Schema Name	The name of the schema that contains the invalid objects
Total Invalid Objects	Total number of invalid objects in the schema

Invalid Objects List

Purpose

The Invalid Objects List collection lists the invalid Object of the specific type currently detected in the database.

Collection Type

Oracle

Collection Sampling Settings

None (on-demand collection).

Metric Descriptions

Metric	Description
Object Name	The name of the invalid object
Object Type	The type of the invalid object
PDB	In multitenant architectures, the name of the pluggable database.
Schema Name	The name of the schema that contains the invalid objects

IO Activity

Purpose

The IO Activity collection provides metrics, wait events, and performance ratios related to input/output activity.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Avg Read Time	The average time (ms) spent reading a data block from the disk into the database buffer cache.
Avg Write Time	The average time spent (ms) writing a data block from the database buffer to the disk.
Balance Average Read Time	The extent, in percents, to which average read time is balanced across RAC nodes
Balance Average Write Time	The extent, in percents, to which average write time is balanced across RAC nodes
Balance Cache Hit Ratio	The extent, in percents, to which cache hit ratio is balanced across RAC nodes
Balance Logical Reads	The extent, in percents, to which logical reads are balanced across RAC nodes
Balance Physical Reads	The extent, in percents, to which physical reads are balanced across RAC nodes
Balance Physical Writes	The extent, in percents, to which physical writes are balanced across RAC nodes
Block Changes per Call	This ratio measures the amount of data manipulation within existing database tables, which each transaction performs via DML commands, by computing the average number of blocks changed per single user call. Because creating or dropping indexes modifies index blocks, such operations increment the value of this metric.
Block Reads	This rate determines the block get rate, which is a basic measure of the rate at which the application system references the database. Change in metric size may indicate that the database workload has materially increased or that the referenced segment sizes increased.

Metric	Description
Cache Hit Ratio	The effectiveness of the buffer cache, computed by comparing the number of block requests, which found the block in the buffer cache, with the entire block requests (including requests that required disk reads).
Chained Row Ratio	This ratio measures the percentage of chained or migrated rows encountered during a fetch as a percentage of the entire fetched rows.
Consistent Changes	The number of times a user process has applied rollback entries to perform a consistent read on the block.
Consistent Gets	The number of data block reads in CONSISTENT mode (that is, consistent reads). A CONSISTENT read retrieves the block version as of the time the query started.
Consistent Gets Direct	The number of times a consistent read was carried out directly for a block, bypassing the buffer cache.
Consistent Gets From Cache	The number of times a consistent read was carried out for a block from buffer cache.
DB Block Changes	The number of database blocks in memory (SGA) that were modified by the session as a result of DML activity. Such changes generate redo log entries, thereby becoming permanent changes to the database if the transaction is committed.
DB Block Gets	The number of data block reads in CURRENT mode (that is, current reads). A CURRENT read retrieves the latest version of the block as of the time of reading. Current mode reads are usually associated with DML activity.
DB Block Gets Direct	Number of times a CURRENT block was retrieved directly, bypassing the buffer cache.
DB Block Gets From Cache	Number of times a CURRENT block was retrieved from the buffer cache.
Disk Sorts Ratio	The number of sort operations that require creation of temporary segments on disk to store the intermediate sort results. This occurs if the data being sorted cannot be fit into memory block specified by the SORT_AREA_SIZE parameter in the <i>init.ora</i> file.
Full Scan Ratio	This ratio measures the percentage of the rows, which were retrieved as a result of full table scans, within the total number of retrieved rows.
Oracle Data I/O Operations	Total number of I/O operations (both physical and logical) performed by the Oracle session.
Oracle Physical Data Processed	the total size (in megabytes) of data processed via physical read and write operations during the specified time range.
Oracle Physical I/O Operations	Total number of physical I/O operations performed by the Oracle session.
Physical Read Bytes	The number of data blocks read by Oracle, which are not present in the buffer cache or in the PGA, thereby requiring I/O read operations, that is: the total number of data blocks read from disk.
Physical Read IO Requests	Number of read requests for application activity (mainly buffer cache and direct load operation) that read one or more database blocks per request.

Metric	Description
Physical Read Total Bytes	Total size in bytes of disk reads by all database instance activity, including application reads, backup and recovery, and other utilities. The difference between this value and the value of the metric <i>physical read bytes</i> provides the total read size in bytes by non-application workload.
Physical Read Total IO Requests	Number of read requests that read one or more database blocks for all instance activity including application, backup and recovery, and other utilities. The difference between this value and the value of the metric <i>physical read total multi block requests</i> provides the total number of single block read requests.
Physical Read Total Multi Block Requests	Total number of read requests which read in two or more database blocks per request for all instance activity including application, backup and recovery, and other utilities.
Physical Reads	The number of data blocks read by Oracle, which are not present in the buffer cache or in the PGA, thereby requiring I/O read operations, that is: the total number of data blocks read from disk.
Physical Reads Cache	Total number of data blocks read from disk into the buffer cache.
Physical Reads Direct	Number of reads that were carried out directly from disk into the PGA, bypassing the buffer cache.
Physical Reads Refetch Warmup	Number of data blocks that were read from the disk during the automatic pre-warming of the buffer cache.
Physical Write Bytes	Total size, in bytes, of all disk writes from the database application activity (this metric does not include any other kind of instance activity).
Physical Write IO Requests	Number of write requests for application activity (mainly buffer cache and direct load operation) that wrote one or more database blocks per request.
Physical Write Total Bytes	Total size, in bytes, of all disk writes for the database instance, including application activity, backup and recovery, and other utilities. The difference between this value and the value of the metric <i>physical write bytes</i> provides the total write size, in bytes, resulting from non-application workload.
Physical Write Total IO Requests	Number of write requests which wrote one or more database blocks from all instance activity including application activity, backup and recovery, and other utilities. The difference between this metric value and the value of the metric <i>physical write total multi block requests</i> provides the number of single block write requests.
Physical Write Total Multi Block Requests	Total number of Oracle instance write requests that wrote two or more blocks per request to the disk for all instance activity, including application activity, recovery and backup, and other utilities.
Physical Writes	Total number of data blocks written to disk. This metric represents the rate at which Oracle generates system I/O calls.
Physical Writes Direct	Number of writes directly to disk, bypassing the buffer cache
Physical Writes From Cache	Total number of data blocks written to disk from the buffer cache
Physical Writes Non Checkpoint	Number of times a buffer is written for reasons other than checkpoint. This statistic measures the number of writes that would have occurred had no checkpoint operation taken place. Subtracting this value from the value of the metric <i>physical writes</i> provides the number of I/O operations carried out as a result of checkpoint operations.

Metric	Description
Read Consistency Overhead	This ratio measures the extent to which applications must make use of the read consistency mechanism. The queries that processing parts of the UPDATE and/or DELETE operations are subject to read consistency.
Read Time	The total time spent on reading data from the data file.
Reads Per Second	Number of physical reads carried out per second
Redo Blocks Written	Total number of redo blocks written. Dividing this metric by the value of the metric <i>redo writes</i> provides the average number of blocks per write.
Redo Buffer Allocation Retries	Total number of retries necessary to allocate space in the redo buffer. Retries are needed either because the redo writer has fallen behind or because an event such as a log switch is occurring.
Redo Entries	The number of redo records created (copied into the redo log buffer).
Redo Log Space Requests	Number of times the active log file is full and the session must wait for a log switch to be performed. After the completion of the log switch, the redo log entries can be allocated, and the session can be resumed.
Redo Log Space Wait Ratio	
Redo Size	Total size, in bytes, of the redo volume.
Redo Writes	Total number of write operations issued by LGWR to the redo log files.
Session Logical Reads	The number of data blocks read by Oracle, which are already present in the buffer cache, thereby requiring no I/O read operations. The value of this metric is a summary of the <i>db block gets</i> and <i>consistent gets</i> metrics.
Sorts (Disk)	Sorts (disk) are sorts that require creation of temporary segments on disk to store the intermediate sort results. Such sorts are carried out if the data being sorted cannot be fit into memory block specified by the SORT_AREA_SIZE parameter in the init.ora file.
Sorts (Memory)	Number of sort operations that were carried out entirely in memory, without requiring any disk writes. The sort operation can take place in memory if the sorted data can be fit into the memory block specified by the SORT_AREA_SIZE parameter in the init.ora file.
Sorts (Rows)	The total number of rows sorted.
Table Scan Blocks Gotten	The number of blocks encountered during the full table scans
Table Scan Short Tables	The number of times Oracle carried out a complete scan of a short table
Table Scans Cache Partitions	Number of range scans performed on tables that have the CACHE option enabled
Table Scans Direct Read	Number of table scans performed using direct read (bypassing the buffer cache)
Table Scans Fetch by ROWID	Number of rows that were fetched using a ROWID (usually recovered from an index)
Table Scans Long Tables	The number of times Oracle carried out a complete scan of a long table
Table Scans ROWID Ranges	The number of table scans conducted during parallel query by the query slaves, with the specified ROWID ranges

Metric	Description
Table Scans Rows Gotten	The number of rows scanned during full table scans. High values may indicate use of non-optimal queries, absence of proper indexing, or both.
Tables Fetch Continued Row	Number of times a chained or migrated row is encountered during a fetch.
Total Physical Data I/O Operations	The total number of read and write operations carried out for the data file.
Total Sorts	Total number of sort operations, including both in-memory and disk sorts.
User Calls	Number of user calls such as login, parse, fetch, or execute. The value of this metric represents the rate at which the user application addresses Oracle.
User Commits	Number of transactions committed during the current interval. After a transaction is committed, the redo generated by that transaction reflects the changes made to database blocks, and the database blocks must be written to disk.
User Rollbacks	Number of transactions that were rolled back, either because of an error or as a result of a manually executed ROLLBACK command
Write Time	The total time spent on writing data to the datafile
Writes Per Second	The number of times DBWR was required to write per second.

Jobs Queue

Purpose

The Jobs Queue collection provides jobs queue information and status.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	300

Metric Descriptions

Metric	Description
Active Jobs	The number of currently running jobs
Jobs Idle	The number of scheduled jobs that are currently not running
Jobs Waiting	The number of jobs that are scheduled to be running now but waiting in queue for the resources
Total Jobs	The total number of scheduled jobs

Latches Activity

Purpose

The Latches Activity collection collects latches activity-related data.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Cache Buffer Chain Latch	This latch is acquired whenever a block in the buffer cache is accessed (pinned)
Cache Buffer LRU Chain Latch	This latch is acquired in order to introduce a new block into the buffer cache and when writing a buffer back to disk, specifically when trying to scan the LRU chain that contains all of the dirty blocks in the buffer cache
Latch Free Waits	Number of waits for a latch that is currently busy (held by another process)
Library Cache Latch	This latch is acquired in order to add a new statement to the library cache. During a parse, Oracle searches the library cache for a matching statement. If one is not found, then Oracle will parse the SQL statement, obtain the library cache latch and insert the new SQL.
Library Cache Pin Latch	This latch is acquired when a statement in the library cache is re-executed
Redo Allocation Latch	This latch is acquired in order to allocate space within the redo log buffer. There is one redo allocation latch per instance.
Redo Copy Latch	This latch is acquired in order to write redo records into the redo log buffer. There may be up to LOG_SIMULTANEOUS_COPIES allocation latches.
Shared Pool Latch	This Latch is acquired when allocating and freeing memory in the shared pool
Total Active Time	Total amount of time waited for the event (in seconds)

Locks Activity

Purpose

The Locks Activity collection provides locks activity-related data.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	300

Metric Descriptions

Metric	Description
Average Lock Duration	The average time (in milliseconds) that a resource was locked (held in contention). High values of lock duration may indicate aggressive lock incidence which may be improved through tuning. The metric is calculated as follows: $100 * (\text{Lock Wait} / \text{Enqueue Waits})$
Enqueue Waits	The total number of times an enqueue request or conversion resulted in a wait.
Enqueue Waits Ratio	The percentage of enqueue wait events (waits resulting from enqueue requests or conversions) within the total time consumed by enqueue operations.
Gets	The total number of enqueue requests or enqueue conversions for this type of enqueue.
Gets Pct	The percentage of lock requests of this type within the total number of lock requests
Lock Efficiency	The percentage of lock requests that did not escalate to a blocking lock.
Lock Name	The name of the lock.
Lock Wait	The total amount of time spent waiting for the enqueue or enqueue conversion.
Waits Pct	The percentage of waits for this type of lock, of all the locks.

Locks Breakdown

Purpose

The Locks Breakdown collection provides lock requests and wait information, including the percentage of time spent on locks of the following types: system, TX, TM, User-defined.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	300

Metric Descriptions

Metric	Description
PDB	In multitenant architectures, the name of the pluggable database.
System Locks Percentage	The percentage of time the sessions spent waiting on System locks, out of the entire time spent waiting on locks.
System Locks Wait Time	The percentage of time the sessions spent waiting on System locks.
TM Locks Wait Time	The amount of time the sessions spent waiting on TM locks.
TM Percentage	The percentage of time the sessions spent waiting on TM locks, out of the entire time spent waiting on locks.
Total Lock Wait	The total time the sessions spent waiting on locks.
TX Locks Wait Time	The amount of time the sessions spent waiting on TX locks.
TX Percentage	The percentage of time the sessions spent waiting on TX locks, out of the entire time spent waiting on locks.
UL Locks Wait Time	The amount of time the sessions spent waiting on user-defined lock types.
UL Percentage	The percentage of time the sessions spent waiting on user-defined locks, out of the entire time spent waiting on locks.

Locks Instance

Purpose

The Locks Instance collection provides summarized lock requests and application efficiency information.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	300

Metric Descriptions

Metric	Description
Balance Enqueue Waits	The extent, in percents, to which enqueue wait events are balanced across RAC nodes
Balance Lock Efficiency	The extent, in percents, to which lock efficiency is balanced across RAC nodes
Balance Lock Wait	The extent, in percents, to which lock wait events are balanced across RAC nodes
Enqueue Waits	Total number of waits on enqueues
Lock Efficiency	The percentage of total locks waits within the total lock requests

Locks Tree

Purpose

The Locks Tree Instance collection provides lock tree information at the instance level.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20

Metric Descriptions

Metric	Description
Blocked SQL	The SQL statement waiting for lock release.
Blocking Serial	The session serial number is either: <ul style="list-style-type: none">Waiting for lock (blocked).Holding the object (locking).
Blocking SID	The session identifier for the session is holding the object (locking).
Client Machine	The name of the Client machine
Cross Instance	The lock is between different cluster nodes for global locks.
DB User Name	The Oracle user, which is either: <ul style="list-style-type: none">Waiting for lock (blocked)Holding the object (locking)
Host	The host where the Oracle instance is installed, which is either: <ul style="list-style-type: none">Waiting for lock (blocked)Holding the object (locking)

Metric	Description
Instance	The Oracle instance session is either: <ul style="list-style-type: none"> • Waiting for lock (blocked). • Holding the object (locking).
Lock Duration	The wait length of time, in seconds.
Lock Event	The event on which the blocked session is waiting
Lock Mode	The lock mode in which the session holds the lock
Lock Request Mode	The lock mode in which the process requests the lock
Lock Type	The type of user or system lock.
Object Wanted	The object requested for lock.
OS User Name	The OS user name with which the session logged in is either: <ul style="list-style-type: none"> • Waiting for lock (blocked). • Holding the object (locking).
PDB	In multitenant architectures, the name of the pluggable database.
PID	The operating system process ID for the Oracle server process. The process ID is either: <ul style="list-style-type: none"> • Waiting for lock (blocked) • Holding the object (locking)
Program	The operating system program is either: <ul style="list-style-type: none"> • Waiting for lock (blocked) • Holding the object (locking)
Serial	The session serial number is either: <ul style="list-style-type: none"> • Waiting for lock (blocked). • Holding the object (locking).
SID	The blocked session's system ID (SID)

Locks Tree RAC

Purpose

The Locks Tree RAC collection provides lock tree information at the global level.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20 seconds

Metric Descriptions

Metric	Description
Blocked SQL	The SQL statement waiting for lock release
Blocking Serial	The session serial number is either: <ul style="list-style-type: none">• Waiting for lock (blocked).• Holding the object (locking).
Blocking SID	The session identifier for the session is holding the object (locking).
Client Machine	The name of the host where the instance is running
Cross Instance	The lock is between different cluster nodes for global locks.
DB User Name	The Oracle user (locking/blocked) that is either: <ul style="list-style-type: none">• Waiting for lock (blocked).• Holding the object (locking).
Host	The host where that the Oracle instance is installed is either: <ul style="list-style-type: none">• Waiting for lock (blocked).• Holding the object (locking).
Instance	The instance name of the session in the lock tree
Lock Duration	The wait length of time, in seconds.
Lock Event	The wait event on which the blocked session is waiting
Lock Mode	The lock mode in which the session holds the lock.
Lock Request Mode	The lock mode in which the process requests the lock.
Lock Type	The type of user or system lock.
Object Wanted	The object requested for lock.
OS User Name	The OS user name that the session logged in with is either: <ul style="list-style-type: none">• Waiting for lock (blocked).• Holding the object (locking).
PDB	In multitenant architectures, the name of the pluggable database.
PID	The operating system process ID for the Oracle server process. The process ID is either: <ul style="list-style-type: none">• Waiting for lock (blocked).• Holding the object (locking).
Program	The operating system program name is either: <ul style="list-style-type: none">• Waiting for lock (blocked).• Holding the object (locking)
SID	The blocked session's system ID (SID)

OS and Instance Statistics

Purpose

The OS and Instance Statistics collection provides the following data:

- CPU and memory information about the operating system
- CPU and memory-related information about the operating system and the monitored instance.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20 seconds
Online	60 seconds
Offline	300 seconds

Metric Descriptions

Metric	Description
ASM Buffer Cache Size	The amount of memory allocated to the ASM buffer cache.
Automatic SGA	Indicates whether the database is configured in manual or automatic memory management mode.
Average Connected Sessions	Total number of current logons.
Balance System CPU Utilization	The extent (displayed as a percentage) to which system CPU utilization is balanced across RAC nodes.
Buffer Cache Size	The amount of memory currently allocated to the buffer cache. The buffer cache, which is used for caching frequently accessed data blocks, is normally the largest memory cache used by Oracle.
Bytes received via SQL*Net from dblink	Total number of bytes received from a database link over Oracle Net Services.
Bytes sent via SQL*Net to client	Total number of bytes sent to the client from the foreground processes.
Bytes sent via SQL*Net to dblink	Total number of bytes sent over a database link.
Calls	Measures the workload being placed on the instance from all work sources.
Current SGA Size	The actual size currently allocated to the System Global Area (SGA) memory.
Cursor Authentications	Number of privilege checks conducted during execution of an operation.
Data Dictionary Cache Hit Ratio	The data dictionary cache hit ratio is used for measuring the ratio of dictionary hits to misses. The data dictionary cache is used for holding rows from the internal Oracle metadata tables, including SQL stored in packages.
Default 16K Buffer Cache Size	The default space allocated within the buffer cache to 16k-size data blocks.
Default 2K Buffer Cache Size	The default space allocated within the buffer cache to 2k-size data blocks.
Default 32K Buffer Cache Size	The default space allocated within the buffer cache to 32k-size data blocks.
Default 4K Buffer Cache Size	The default space allocated within the buffer cache to 4k-size data blocks.
Default 8K Buffer Cache Size	The default space allocated within the buffer cache to 8k-size data blocks.

Metric	Description
Default Buffer Cache Size	Size in MB of default cache area of SGA that cache blocks of data and indexes that are not assigned to specific buffer pools.
Execute Count	Number of statements executed during the specified time range.
Fixed Size	The fixed part of the SGA is an internal area set by Oracle, which cannot be manually modified. This area contains the following data: <ul style="list-style-type: none"> General information about the state of the database and the instance that needs to be accessed by background processes Information communicated between processes
Flash Cache Insert Skip: DBWR Overloaded	The number of writes to the flash cache that were skipped (insert skip) as a result of an overloaded DBWR, which has exceeded a predefined threshold.
Flash Cache Read Hit Ratio	The percentage of read requests that were resolved from Flash Cache, out of the total of physical read total IO requests.
Flash Cache Read Hits	The number of read requests, resolved from Flash Cache.
Free Memory	Amount of free RAM in the system.
Free SGA Pct	The percentage of free memory out of the total amount of memory allocated to SGA.
Free SGA Size	The size (in megabytes) of free memory out of the total amount of memory allocated to SGA.
Hard Parse Ratio	The ratio of hard parse calls (real parses) to total parses.
Instance CPU	The amount of CPU time used by a session from the time a user call starts until it ends.
Instance CPU Load	The amount of time the instance was using CPU within the total elapsed time (as a percentage).
Instance RAM Percentage	Total memory, in percentage, allocated to the instance out of the total RAM.
Java Pool Free Pct	The amount of free space in the Java pool, expressed as a percentage of the Java pool's total size.
Java Pool Free Size	The amount (in megabytes) of free space in the Java pool.
Java pool Size	The size of the Java pool, which caches class definition, Java methods, and Java objects. An improperly sized Java pool can affect the parse rate.
Java Pool Used Size	The amount (in megabytes) of used space in the Java pool.
Keep Pool Cache Size	The size of the Keep buffer pool. This pool contains frequently accessed segments in the buffer cache.
Large Pool Free Pct	The amount of free space in the Java pool, expressed as a percentage of the large pool's total size.
Large Pool Free Size	The amount (in megabytes) of free space in the large pool.
Large Pool Size	The size of the Large pool. This pool is used in multi-threaded servers (MTS) for session memory, parallel execution, and backup processes.
Large Pool Used Percent	The percentage of used space in the large pool.
Large Pool Used Size	The amount (in megabytes) of used space in the large pool.

Metric	Description
Library Cache Hit Ratio	The effectiveness of the library cache, computed by comparing the number of queries submitted to Oracle server for execution, which found the parsed execution plan in the library cache (soft parsing), with the entire number of queries submitted during the specified time range (including queries that required the parsing of a new statement, also known as "hard parsing").
Logons Cumulative	The total number of current logons.
Logons per Second	The average number of logons per second for the specified time range.
Maximum Memory Size	The maximum size (in megabytes) of memory that can be allocated to an Oracle instance.
Maximum SGA Size	The maximum defined size of the System Global Area (SGA) memory.
Memory Usage	The amount of memory consumed by all OS processes (including Oracle).
Memory Utilization	The relative share, in percents, of memory consumed by all OS processes (including Oracle) out of the total memory.
Messages Received	Number of messages sent and received between background processes.
Opened Cursors Cumulative	Total number of cursors opened since the instance started.
Opened Cursors Current	Total number of currently open cursors.
Oracle/System CPU Usage	The amount of CPU consumed by the entire system, broken down into monitored instance activity and all other activities
Parse count (hard)	The number of times an SQL statement has undergone a hard parsing process, which requires reloading the SQL statement to the shared pool, along with the associated RAM memory allocation.
Parse count (total)	Total number of parse calls (hard and soft).
PDB	In multitenant architectures, the name of the pluggable database
PGA Currently Allocated	The amount of Program Global Area (PGA) memory allocated for the current session.
Physical Read Total IO Requests	The total number of I/O requests due to physical reads.
Potential SGA Growth	The amount (in megabytes) by which the SGA memory has increased during the specified time range.
Recursive Calls	Number of recursive calls generated at both the user and system level. For more information, see on page 183 .
Recycle Pool Cache Size	The amount of memory currently allocated to the Recycle buffer pool. This pool contains blocks belonging to segments that are not frequently accessed.
Redo Buffers Cache Size	The size of the redo buffer, which contains redo entries that must be written to the redo log. A properly sized redo buffer improves performance by minimizing the time waiting for writes to the log file to complete.
Result Cache Find/Create Ratio	<p>The result cache find/create ratio is used for measuring the ratio of result cache hits to misses.</p> <p>The Result Cache is an area within the shared pool, which contains the end results of a query execution. Failure to find these end results, that is: a result cache miss, leads to creation operation.</p> <p>Low values of this metric indicate that the contents of the result cache are being inefficiently used.</p>

Metric	Description
Result Cache Max Size	The maximum size (in megabytes) of memory allocated to the result cache.
Session Cursor Cache Count	Total number of cursors cached. This statistic is incremented only if <code>SESSION_CACHED_CURSORS > 0</code> .
Session Cursor Cache Hits	Number of hits in the session cursor cache. A hit means that the SQL statement did not have to be re-parsed.
Shared Pool Free Pct	The amount of free space in the shared pool, expressed as a percentage of the shared pool's total size.
Shared Pool Free Size	The amount (in megabytes) of free space in the shared pool.
Shared Pool Size	The size of the shared pool. The shared pool caches SQL statements, PL/SQL programs, object definitions and session memory for MTS sessions. Performance can improve when the shared pool is sized properly, thereby avoiding reparsing and reloading.
Shared Pool Used Percent	The amount of used space in the shared pool, expressed as a percentage of the shared pool's total size.
Shared Pool Used Size	The amount (in megabytes) of used space in the shared pool.
SQL*Net roundtrips to/from client	Total number of Net8 messages sent to and received from the client
SQL*Net roundtrips to/from dblink	Total number of Net8 messages sent and received over a database link.
Stream Pool Size	The size (in megabytes) of the stream pool area (the area that caches Oracle Streams objects) within the SGA.
System CPU Utilization	The overall operating system CPU usage, including Oracle.
System Kernel CPU Utilization	Percentage of CPU time consumed by the operating system's processes (kernel mode activities).
System User CPU Utilization	Percentage of CPU time consumed by the operating system's processes (user mode activities).
Target Memory Size	The size (in megabytes) required for optimal memory utilization by the monitored Oracle instance. When this parameter, which was introduced in Oracle 11g, is set, Oracle automatically tunes memory allocation by resizing the SGA and PGA as needed.
Target SGA Size	Specifies the total size (in megabytes) of all SGA components. Specifying the <code>SGA_TARGET</code> results in automatic sizing of the following memory pools: <ul style="list-style-type: none"> • Buffer cache • Shared pool • Large pool • Java pool • Streams pool
Transactions	The number of transactions running.
Used RAM Pct Rate	The overall utilization of RAM by all processes, including Oracle.
Used SGA Pct	The amount of used space in the SGA memory, expressed as a percentage of the SGA memory's total size.
Used SGA Size	Total size of the currently used SGA memory.
User Calls	Number of user calls such as login, parse, fetch, or execute. This metric value represents the rate at which the user application addresses Oracle.

Metric	Description
User Calls Per Parse	The average number of calls in relation to statement parses. Changes in this rate most likely result from application change, but may also indicate that usage patterns are changing and users are moving from one module to another, either more or less frequently.
User Commits	Number of transactions committed during the current interval. After a transaction is committed, the redo generated by that transaction reflects the changes made to database blocks and they must be written to disk.
User Rollback Ratio	The percentage of rolled back transactions, that is: transactions for which the ROLLBACK statement was issued, either manually or as a result of an error that occurred during the transaction.
User Rollbacks	Number of transactions that were rolled back. Rollbacks occur either because of a manually issuing a ROLLBACK command or in the event of an error.
Variable Others Size	The size of variable portion of SGA excluding shared pool, large pool and java pool.
Variable Size	The variable portion of SGA, which contains shared pool, large pool and java pool.

OS File System IO Summary

Purpose

The OS File System IO Summary collection provides IO information about the operating system.

Collection Type

Operating System

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20 seconds
Online	60 seconds
Offline	300 seconds

Metric Descriptions

Metric	Description
Operating System	The name of the operating system where the selected instance is running.
System Physical Reads	The number of physical read operations carried out from this operating system.

Metric	Description
System Physical Writes	The number of physical write operations carried out from this operating system.
System RW Operations	Total number of disk operations (both read and write) carried out by the operating system. This metric does not include operations that were satisfied using the file system.

OS General Network Summary

Purpose

The OS General Network Summary collection provides network information about the operating system.

Collection Type

Operating System

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Net Collisions	The number of times two computers sent packets at the same time on the network, and the packets “collided”, resulting in the need to retransmit both packets.
Net Errors	Number of errors reported on the various network interfaces.
Network Name	The network interface name.
Operating System	The name of the operating system where the selected instance is running.
Packets	The total sum of packets received and sent on this network interface (the sum of the packets_in and packets_out values).
Packets In	Number of packets received by this network interface.
Packets Out	Number of packets sent from this network interface.

Parallel

Purpose

The Parallel collection provides information about the instance’s parallel executions.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Active Parallel Sessions (11i)	Number of active parallel sessions
Active Serial Sessions (11i)	Number of active sessions that were serialized
DDL Statements Parallelized	Number of DDL statements that were executed in parallel
DML Statements Parallelized	Number of DML statements that were executed in parallel
Parallel operations Downgraded	Total number of times parallel execution was requested but execution was downgraded
Parallel Operations Downgraded 1 to 25 Pct	Number of times parallel execution was requested but execution was downgraded up to 25% due to insufficient parallel execution servers
Parallel Operations Downgraded 25 to 50 Pct	Number of times parallel execution was requested but execution was downgraded 25% or more as a result of insufficient parallel execution servers
Parallel Operations Downgraded 50 to 75 Pct	Number of times parallel execution was requested but execution was downgraded 50% or more as a result of insufficient parallel execution servers
Parallel Operations Downgraded 75 to 99 Pct	Number of times parallel execution was requested but execution was downgraded 75% or more as a result of insufficient parallel execution servers
Parallel Operations Downgraded to Serial	Number of times parallel execution was requested but execution was serial as a result of insufficient parallel execution servers
Parallel Operations not Downgraded	Number of times parallel execution was executed at the requested degree of parallelism
Parallel QC (11i)	The number of parallel query coordinators
Parallel Slaves (11i)	The number of parallel slaves
PQ Servers Busy	The number of servers that were simultaneously executing SQL queries and DML statements.
PQ Servers Busy Percent	The percent of busy parallel servers within the total number of parallel servers
PQ Servers Idle	The number of currently idle parallel servers
PQ Servers Limit	The maximum allowed number of parallel servers
Queries Parallelized	Number of SELECT statements executed in parallel

PGA Statistics

Purpose

The PGA Statistics collection collects information about the use of the PGA memory.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
PGA Aggregate Target	<p>The value of the PGA Aggregate Target parameter in bytes.</p> <p>The <code>pga_aggregate_target</code> parameter, which was introduced in Oracle 11g, defines the total amount of memory available in Oracle for allocation to the PGA memory structure.</p> <p>If the value of this parameter is set correctly, all work areas executed by the system since start-up are processed exclusively in the PGA memory, requiring no extra passes. In such a case, the PGA hit percent's value is 100%, and Oracle does not need to allocate extra memory to the PGA. However, when this value is set too small, the hit percent's value decreases, and Oracle may need to perform one or more extra memory allocations.</p>
PGA Hit Percent	<p>The PGA cache hit ratio (as a percentage).</p> <p>When the value of this metric is 100%, it indicates that all work areas executed by the system processed the entire input data exclusively via the PGA memory, thereby making optimal use of this memory. Any lower value indicates that at least one work area required one or more extra passes over the input data.</p>
PGA Overallocation Count	<p>The number of extra PGA memory allocations that were carried out during the specified time range.</p> <p>Over-allocating PGA memory can happen if the value of the dynamic initialization parameter, <code>PGA_AGGREGATE_TARGET</code>, which is set to all server processes, is too small to allow executing all work area workload (operations) exclusively in the PGA memory. Such a scenario requires Oracle to allocate extra PGA memory.</p> <p>Ideally, the PGA cache hit ratio should be close to 100%, resulting in an extremely small number (close to zero) of PGA overallocations. High values of the PGA Overallocation Count metric indicate the need to increase the value of <code>PGA_AGGREGATE_TARGET</code>.</p>

Pluggable Databases Information

Purpose

This collection contains general information about pluggable databases.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	3600
Online	3600
Offline	3600

Metric Descriptions

Metric	Description
PDB ID	The identifier used for the pluggable database.
PDB	The name of the pluggable database.
Cloned From PDB Name	Indicates from which pluggable database the selected pluggable database is cloned.
Created Time	The time when the pluggable database was created.
Share	The portion of the system's CPU resources that is allocated to the pluggable database.

Pluggable Databases State

Purpose

This collection contains the status, state, and startup time for pluggable databases.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	900

Metric Descriptions

Metric	Description
PDB ID	The identifier used for the pluggable database.
PDB	The name of the pluggable database.
PDB Availability	Status of the pluggable database, expressed as a percentage.
Open Mode	State of the pluggable database.
Open Time	Startup time of the pluggable database.

Redo Archive Performance

Purpose

The Redo Archive Performance collection provides redo and archive activity information.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Average Redo Write Time	The average time (ms) spent writing a redo log entry to the log files. User sessions experience log write waits when, upon commits, an individual session waits for the LGWR to flush its redo information to the redo log file.
Balance Archive Waits	The extent to which archive wait events are balanced across RAC nodes.
Balance Average Redo Write Time	The extent to which average redo write time is balanced across RAC nodes.
Balance Redo Waits	The extent to which average redo wait events are balanced across RAC nodes.
Balance Redo Writes	The extent to which redo writes are balanced across RAC nodes.
Commits	Number of transactions committed during the current interval. After a transaction is committed, the redo generated by that transaction reflects the changes made to database blocks, which must be written to disk.
Log File Sync Event	Average time spent on writing redo log entries during a commit.

Metric	Description
Log Switch Wait	The time the session spent waiting for the completion of a redo log switch.
Log Write Wait	Time spent waiting for the completion of a write (physical I/O) to a logfile. LGWR experience log write waits when it periodically flushes the contents of the log buffer into the active log file.
Redo Waits	Elapsed redo wait time, including the time spent waiting on contention resulting from: <ul style="list-style-type: none"> • User sessions writing redo entries to the redo log buffer • LGWR writing the log buffer contents into the active redo log file • ARCH archiving redo logs to the destination location
Redo Writes	Total number of write operations issued by the log writer (LGWR) to the redo log files.
Redo Writes Volume	Total amount (in megabytes) of redo operations generated and written to redo logs
Time to Sync a Redo Log Entry	Average time spent on writing single redo log entry to logfile during a commit.

Redo Instances Miscellaneous

Purpose

The Redo Instances Miscellaneous collection provides Redo Log Switches information.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	3600
Online	3600
Offline	3600

Metric Descriptions

Metric	Description
Log Switches	The number of redo Logs switches during the specified time range

Redo Logs Member Storage

Purpose

The Redo Logs Member Storage is a list of Redo log members and their properties.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Group Number	The redo log group number
Member Name	The redo log group member name
Member Status	The redo log group member status is one of the following: <ul style="list-style-type: none">• invalid• stale• deleted• in use
Redo Log Group	Redo Log Group name

Redo Log Storage

Purpose

The Redo Log Storage collection provides a list of all Redo log groups and their overall properties.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Archived	The archive status of the log file: <ul style="list-style-type: none">• yes• no
First change#	A system-allocated number (SCN), identifying the first time a change was recorded in this log.
First Time	The time at which the first SCN in the log was recorded.
Group Status	The status of the group is one of the following: <ul style="list-style-type: none">• unused• current• clearing• clearing_current• inactive This information is derived from the v\$log table.
Group#	The log group number.
Sequence#	The sequence number.
Size	The size (in megabytes) of the log .

Sessions Summary

Purpose

The Sessions Summary collection retrieves information regarding session types.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Active Sessions	The average number of active sessions during the specified time range
Active Sessions PCT	The average percentage of active sessions within the total sessions during the specified time range
Background Sessions	The average number of background sessions that ran during the specified time range.

Metric	Description
Blocked Sessions	The average number of blocked processes that were collected during the specified last sample.
Cached Sessions	The average number of cached sessions during the specified time range
Foreground Sessions	The average number of foreground sessions that ran during the specified time range.
Inactive Sessions	The average number of inactive sessions during the specified time range
Killed Sessions	The average number of killed sessions during the specified time range
PDB	In multitenant architectures, the name of the pluggable database
Sniped Sessions	The average number of sniped sessions during the specified time range
Total Sessions	The average number of sessions during the specified time range
Total Sessions Limit	The maximum allowable number of sessions

Servers

Purpose

The Servers collection provides instance processes.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20
Online	60
Offline	300

Metric Descriptions

Metric	Description
Dedicated Servers	The number of dedicated server processes. Each dedicated server process service a single client process.
Dispatchers	The number of multi-threaded server (MTS) dispatchers. MTS dispatchers coordinate the allocation of shared servers to client tasks.
Shared Servers	The number of servers that service more than one client process. This value includes both shared servers and multi-threaded servers (MTS).

Tablespaces Free Space Fragmentation

Purpose

The Tablespaces Free Space Fragmentation collection retrieves information regarding the fragmentation of tablespace free space.

Collection Type

Oracle

Collection Sampling Settings

None (on-demand collection).

Metric Descriptions

Metric	Description
Average Extent Size	The average size of free extents in this tablespace
Extents	The number of free extents in this tablespace
Max Extent Size	The size of the largest free extent in the selected tablespace
Min Extent Size	The size of the smallest free extent in this tablespace
Partly Usable	The total sum of fragments that are smaller than the maximum allocated extent in the tablespace, but bigger than, or equal to, the minimum allocated extent in the tablespace
PDB	In multitenant architectures, the name of the pluggable database
Tablespace Name	The name of the currently selected tablespace
Unusable	Total sum of fragments whose size is less than the size of the minimum extent in the tablespace
Usable	Total sum of fragments that are greater than or equal to the minimum extent in the tablespace.

Tablespaces Storage

Purpose

The Tablespaces Storage collection provides a list of all tablespaces and their overall properties and storage characteristics.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	3600

Frequency Mode	Collection Interval (Seconds)
Online	3600
Offline	14400

Metric Descriptions

Metric	Description
Allocated Space Used Pct	The percent of used allocated space within this tablespace
Allocation Type	The type of extent allocation to be used for this tablespace. The possible values for this metric are as follows: <ul style="list-style-type: none"> • USER • SYSTEM • UNIFORM
Auto DF ASM	The number of autoextensible ASM datafiles
Auto DF FS	The number of autoextensible OS datafiles
Average Fragment Size	The average size of fragments in this tablespace
Big File	Indicates whether the tablespace is a BIGFILE tablespace. BIGFILE tablespaces are supported only for locally managed tablespaces with automatic segment-space management.
Block Size	The tablespace blocksize
Contents	The type of tablespace contents is one of the following: <ul style="list-style-type: none"> • PERMANENT • TEMPORARY • UNDO
Database	The database name
Default Table Compression	Indicates whether table compression is enabled by default for this tablespace
Encryption	Indicates whether the tablespace data is encrypted.
Estimated Date Until Full	The estimated date when the tablespace becomes full.
Estimated Until Full	The estimated number of days until the tablespace is full.
Extent Management	Indicates the extent management type: <ul style="list-style-type: none"> • DICTIONARY – where the management of extents is within the Oracle data dictionary. • LOCAL – where the management of extents is within the tablespace itself.
Force Logging	Indicates whether the tablespace is in FORCE LOGGING mode: <ul style="list-style-type: none"> • Yes – logs all changes to all objects in the tablespace except changes to temporary segments. • No – logs are not created.
Fragments	The number of fragments in this tablespace
Free	The total amount of free space in this tablespace
Free Pct	The percent of free space in this tablespace
Logging	Indicates whether a redo log is to be generated for the tablespace
Long Term Growth Rate	The long term growth rate.
Manual DF ASM	The number of manual ASM datafiles

Metric	Description
Manual DF FS	The number of manual operating system datafiles
Max Fragment Size	The size of the largest fragment in this tablespace
Min Fragment Size	The size of the smallest fragment in this tablespace
Number of Datafiles	The number of datafiles in the tablespace
PDB	In multitenant architectures, the name of the pluggable database
Recycle Free Pct	The percentage of recycle bin space which occupies the allocated free space of the tablespace.
Recycle Free	The total amount of space used by the recycle bin within the allocated free space of the tablespace
Retention	Indicates whether unexpired UNDO operations are retained in the tablespace. This applies only to UNDO tablespaces.
Segment Space Management	Indicates the method used for managing free and used segments: <ul style="list-style-type: none"> manually – by free lists automatically – by bitmaps
Short Term Growth Rate	The short term growth rate.
Status	The current status of the tablespace. <ul style="list-style-type: none"> ONLINE OFFLINE
Tablespace Name	The name of the currently selected tablespace.
Total	The total amount of space in the currently selected tablespace.
Used	The amount (in megabytes) of space used in this tablespace
Used Growth	The amount (in megabytes) by which the size of the tablespace has increased during the specified time range

Top Blocked Objects

Purpose

The Top Blocked Objects collection provides a list of top blocked objects for TX locks.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	300

Metric Descriptions

Metric	Description
Segment Name	The name of the segment
Segment Waits	The number of waits for a specific segment during the selected time range

Top Sessions

Purpose

The Top Sessions collection provides instance current sessions details.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	20 seconds

Metric Descriptions

Metric	Description
Action	The currently executing action name set by DBMS_APPLICATION_INFO.SET_MODULE
BG CPU Time	The amount of CPU time used by a session from the time a user call starts until it ends
BG Time	The amount of time used by a session from the time a user call starts until it ends
Blocking Session Status	The status of the blocking session; indicates whether the session is currently executing SQL statements (ACTIVE) or not (INACTIVE).
Client Info	The client information; contains the information set by DBMS_APPLICATION_INFO.SET_CLIENT_INFO
Client Machine Name	The name of the computer where the Client is running
Client PID	The operating system process ID for the Client program
Consistent Changes	The number of times a user process has applied rollback entries to carry out a consistent read on the block
Consistent Gets	The number of times a consistent read was requested for a block by this session
Current SQL	The last SQL statement executed by the user
DB Block Changes	The total number of changes that took part in an Update or Delete operation made to blocks by this session
DB Block Gets	The number of times a current block was requested by this session

Metric	Description
DB CPU Time	Time spent waiting for the CPU to process SQL statements and other types of calls made to Oracle
DB Time	The amount of time used by the database for this session
Event	The name of the event that resulted in the wait event
Failed Over	In a RAC environment, a user can connect to any one of the multiple instances in a cluster. If the connected instance fails, the session can be reconnected automatically to other available instance. This metric indicates whether the session is running in failover mode and failover has occurred (YES) or not (NO).
Failover Method	The failover method used for determining how failover occurs from the primary node to the backup node
Failover Type	Failover type settings: SESSION, SELECT, NONE
Fixed Table Sequence	This metric's value is incremented each time the session completes a call to the database. Lack of changes in this metric's value indicates that the selected session was idle during the selected time range.
Hard Parse Count	The number of hard parse requests performed by this session
Logical Reads	The number of logical reads performed by session since login. This value is the total summary of all requests for database blocks, irrespective of whether they were found in the database cache.
Logical Writes	The number of logical writes performed by the session since login.
Logon Time	The date and time when the user logged on to the database. The format is year/month/day and hours:minutes:seconds (on a 24 hour clock).
Module Name	The currently executing module name, set by DBMS_APPLICATION_INFO.SET_MODULE.
Oracle User Name	The Oracle database account of the started session
OS User Name	The operating system user used by the client
Parallel Status	The Parallel Query status; can have one of the following values: ENABLED – the session is in a Parallel Query enabled mode DISABLED – the Parallel Query enabled mode is not supported for the session FORCED – the session has been altered to force Parallel Query)
PDB	In multitenant architectures, the name of the pluggable database
Physical Reads	The number of data blocks read by Oracle, which are not present in the buffer cache or in the PGA, thereby requiring I/O read operations, that is: the total number of data blocks read from disk.
Physical Writes	The number of physical writes carried out by the session since logon
Redo Entries	The amount of redo entries issues by the session since logon
Resource Consumer Group Name	The name of the session's current resource consumer group
Schema Name	The name of the database schema
Seconds in Wait	The wait event's length of time, in seconds, since logon
Serial Number	The session's serial number. The combination of SID and serial number provides a unique identification of the session.
Server	The server's type: DEDICATED, SHARED, PSEUDO. or NONE
Server PID	The operating system process ID for the Oracle shadow process.
Server Program Name	The server program name
Service Name	The service name of the session

Metric	Description
Session Action	The currently executing action name, set by DBMS_APPLICATION_INFO.SET_MODULE.
Session Cursor Cache Hits	The number of session cursor cache hits.
Session ID	The session identification number.
Session PGA Memory	The amount of PGA memory allocated for the current session
Session SQL ID	The session identification number for the SQL statement
Session Status	The session status; indicates whether the session is currently executing SQL statements (ACTIVE) or not (INACTIVE)
Session Type	The session type (USER or BACKGROUND)
Session UGA Memory	The amount of UGA memory allocated to the current session
Sort Disk	The number of sort operations that require creation of temporary segments on disk to store the intermediate sort results. This occurs if the data being sorted cannot be fit into memory block specified by the SORT_AREA_SIZE parameter in the <i>init.ora</i> file.
Wait Class	The classification of wait events.

Top SQLs

Purpose

The Top SQLs collection retrieves top n SQL statements (namely: statements that experienced the longest time of CPU consumption or total wait events during the specified time range) and their details.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Array Wastage	<p>The gap (by a percentage) between the average number of rows per fetch that the array can hold, as defined in the <i>arraysize</i> parameter of the various fetch operations of the specified SQL query, and the number of rows actually retrieved.</p> <p>For example: if the average setting of the <i>arraysize</i> parameter was 5, and the average number of rows actually retrieved per operation was 10, the value indicated will be 50.</p> <p>Ideally, the value of this parameter is 0; any other value indicates the need for additional sql.net roundtrips per each SQL query execution.</p>
Buffer Gets	Number of buffer gets for this child cursor
Child Number	Hash value of the parent statement in the library cache
CPU Time	CPU time used by this cursor for parsing, executing, and fetching
Direct Writes	Number of direct writes for this child cursor
Disk Reads	Number of disk reads for this child cursor
Elapsed Time	Elapsed time used by this cursor for parsing, executing, and fetching
Executions	Number of executions that took place on this object since it was brought into the library cache
Fetches	Number of fetches associated with the SQL statement
Fetches per Execution	The average number of fetches per a single execution associated with the SQL statement
First Load Time	The first time the object was loaded to the shared pool
Hash Value	Hash value of the parent statement in the library cache
Java Time	Java execution time
Loads	Number of times the object was either loaded or reloaded
Loads per Execution	The average number of loads per a single execution associated with the SQL statement
Other Buffer Gets	Total time spent waiting for buffer gets on all SQL statements other than the selected statement.
Other CPU Time	Total CPU time consumed for executing all SQL statements other than the selected statement.
Other Disk Reads	The total number of disk reads carried out by all SQL statements except the selected statement.
Other Executions	Number of executions that took place on this object since it was brought into the library cache.
Other Total Waits	Total time spent on wait events due to executing all SQL statements except the selected statement.
Parallel Executions	Total number of executions performed by Parallel execution servers.
Parse Calls	Number of parse calls for this child cursor.
PLSQL Time	PL/SQL execution time.
Rows Processed	Total number of rows the parsed SQL statement returns
Sorts	Number of sorts that were done for the child cursor
Sorts per Execution	The average number of sorts per a single execution associated with the SQL statement
Total Waits	The total time spent on wait events.

Top SQLs Full Text

Purpose

This collection retrieves the SQL statement's full text.

Collection Type

Oracle

Collection Sampling Settings

The Top SQLs Full Text collection is an on-demand collection, which retrieves data only upon clicking View full text.

Metric Descriptions

Metric	Description
Hash Value	The SQL statement's hash value
SQL Full Text	Full text for the SQL statement exposed as a Character Large Objects (CLOB) column

Top SQLs Text

Purpose

The Top SQLs Text collection retrieves the first n characters of the SQL statement text (the exact number of the characters is defined in the Top SQL Statements pane of the Collection Settings screen).

Collection Type

Oracle

Collection Sampling Settings

This collection retrieves SQL text only for new hash values.

Metric Descriptions

Metric	Description
Hash Value	The SQL statement's hash value.
Is Full Text	A boolean value that indicates whether the first n characters are a full or partial SQL text for the current cursor.
SQL Text	First n characters of the SQL text for the current cursor.

Total Archive Storage

Purpose

The Total Archive Storage collection collects disk space data for critical archive destinations.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Archive Destination Free Space	The free space size (in megabytes) in the file system where the most critical archive destination resides.
Archive Destination Size	The total size (in megabytes) of the file system where the most critical archive destination resides.
Critical Archive Time To Failure	The number of hours remaining until the archive destination with the shortest time to failure is filled.
Destination	The destination ID of the most critical archive destination.

Total Database Storage

Purpose

The Total Database Storage collection retrieves the database's total storage, that is: the summarized value of total free space and total allocated space.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	300
Online	900
Offline	3600

Metric Descriptions

Metric	Description
Maximum Used Percent	The maximum percent of the used disk space of all the database datafiles.
Total Allocated Space	The total size (in megabytes) of allocated space in the database.
Total Allocated Space Delta	The incremental change in the total size of the database's allocated space during the specified time range.
Total Allocated Space Rate	The rate (in megabytes/seconds) at which disk space was allocated on the database during the last sample.
Total Free Space Percent	The percent of the total free disk space within the database.

Total Instance Waits

Purpose

The Total Instance Waits collection collects the total amount of time (in seconds) that the instance spent on wait events since the last sample.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	300
Offline	900

Metric Descriptions

Metric	Description
Total Waits	The total amount of time (in seconds) that the instance spent on wait events since the last sample.

Undo Activity

Purpose

The Undo Activity collection collects UNDO operations-related information.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	600
Online	600
Offline	600

Metric Descriptions

Metric	Description
Active UNDO Tablespace	The name of the last active undo tablespace of the last period when undo statistics were collected (by default: every 10 minutes)
Begin Time	The beginning of the last period when undo statistics were collected
End Time	The end of the last period when undo statistics were collected
Expired Blocks Removed	Number of expired undo blocks stolen from other undo segments
Expired Blocks Reused	Number of expired undo blocks reused within the same undo segments
Expired Blocks Steals	Number of attempts to steal expired undo blocks from other undo segments
Highest Concurrently TNXs	The highest number of transactions executed concurrently within the specified time range
Longest Query	The length, in seconds, of the longest query executed in the instance during the specified time range
Longest Query SQL ID	SQL identifier of the longest running SQL statement within the specified time range
No Space Error Count	The number of times requests for space in the undo tablespace failed to find any free space available
SnapShot Too Old Error Count	The number of times the error ORA-01555 (snapshot too old) occurred
Transactions	The total number of transactions executed within the specified time range
Unexpired Blocks Removed	Number of unexpired blocks removed from certain undo segments so they can be used by other transactions
Unexpired Blocks Reused	Number of unexpired undo blocks reused by transactions
Unexpired Blocks Steals	Number of attempts to obtain undo space by stealing unexpired extents from other transactions
Used Blocks	The total number of undo blocks consumed

Total Summary

Purpose

The Total Summary collection collects UNDO definitions, storage and retention.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	600
Online	600
Offline	600

Metric Descriptions

Metric	Description
Active Undo Tablespace	Name of the last active undo tablespace
Actual Undo Size	The size (in megabytes) of the undo tablespace
Autoextensible	Indicates whether the datafile is autoextensible
Current Retention	Specifies (in seconds) the low threshold value of undo retention
Minimum Required Undo Size	The minimum size (in megabytes) of undo space required. This value takes into account transactions rate, undo generation rate and current undo retention.
Optimal Retention	The optimal undo retention, taking into account undo generation rate and actual undo size
Retention Guarantee	Indicates the state of the Undo tablespace retention (GUARANTEE/NOGUARANTEE) Retention value of GUARANTEE indicates that unexpired undo, in all undo segments in the undo tablespace, should be retained even at the cost of failure of forward-going operations that need to generate undo in those segments.
Transactions Count	The rate of transaction executed within the specified time range
Undo Generated Volume	The rate per second (in megabytes) of undo generated within the specified time range
Undo Management	Indicates whether undo management is automatic or manual

Usability

Purpose

The Usability collection provides retrieved instance usability based on response time and instance availability.

Collection Type

Oracle

Collection Sampling Settings

Frequency Mode	Collection Interval (Seconds)
Realtime Collection	60
Online	60
Offline	300

Metric Descriptions

Metric	Description
Availability	The percentage of time, within the selected time range, during which the database was available: <ul style="list-style-type: none">• For a RAC node – the percentage of nodes that are up and running• For a single instance – the percentage of the time during this period the instance was up and running
Connect Availability	The percentage of time, within the selected time range, during which the connection to the database instance was successful.
Connection Time	The amount of time it takes to establish a connection.
Error Message	Last error message that was received upon trying to connect to the database (if an unsuccessful connection trial took place).
Listener Error Message	The last error message (if any) for each listener
Listener Name	The listener name for each listener
Listener Status	For each listener: the percentage of time, within the selected time range, during which the listener was available. IMPORTANT: Retrieving the Listener Status information requires having an OS user with the permission to run the lsnrctl command.
OS Connect Availability	The percentage of time, within the selected time range, during which the connection to the OS was successful.
OS Error Message	The last error message (if any) received upon an unsuccessful connection trial.
Response Time	The amount of time (in microseconds) used by the active cursor for carrying out parsing/executing/fetching operations
Up Since	The date and time when the instance started

Glossary

Alarm

The mechanism by which Foglight *for Oracle* alerts users to a condition that might be a problem in the Oracle instance.

Authentication

The process of identifying and verifying a user who is attempting to establish an Oracle session. With Oracle Authentication, users must have an Oracle login and password which is validated against an Oracle system table.

Buffer Wait

Buffer wait events result from a contention in accessing database block buffers; for example, if the requested data block is in an incompatible mode (*buffer busy*), or if the buffer does not contain sufficient disk space for reading the data block (*free buffer*).

Calibration

The process by which Foglight *for Oracle* determines the maximum and minimum values for every dataflow on the home page, by observing data moving through the database system. This information helps Foglight *for Oracle* display the dataflows correctly.

Cartridge

Cartridges extend the functionality of Foglight and are installed on the Foglight Management Server. A cartridge contains one or more components, such as agents for deployment, communication capabilities, and modifications to the way that data is transformed or handled, as well as rules, reports, and views. When a cartridge is installed and enabled, its components become part of the Management Server. Adding cartridges allows users to monitor additional parts of their environment. For further details about managing cartridges, see the *Foglight Administration and Configuration Guide*.

Cluster Wait

Cluster wait events take place when the session waits for all types of cluster resources, such as Interconnect hardware, shared disks, and Oracle instances.

CPU Usage

When SQL statements and other types of calls are made to Oracle, processing the calls requires spending a certain amount of CPU time. Whereas processing average calls requires a small amount of CPU time, a SQL statement involving a large amount of data or a runaway query can consume a much larger amount, thereby reducing the CPU time available for other processing.

CPU utilization is a key operating system statistic in the tuning process. Excessive CPU usage can result from an inadequately-sized system, untuned SQL statements, or inefficient application programs.

CPU Wait

CPU wait events take place when the session is waiting in the system's run queue to be granted for CPU cycles. The length of these wait events (the amount of time spent) depends upon the number of concurrent processes and threads requesting CPU time.

This metric's value should be inspected in conjunction with the value of the Run Queue Length.

Cursors

Extensions to result sets that provide the mechanism for working with individual rows, or a small block of rows, in a table.

Because a cursor points to a currently selected set of records, they can be used by only one connection at a time. However, the compiled plan to which the cursor is linked can be used simultaneously by multiple connections.

DataGuard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.

Using Data Guard's set of services for creating, maintaining, managing, and monitoring one or more standby databases, enables production Oracle databases to survive disasters and data corruptions. These standby databases are then maintained as transactionally-consistent copies of the production database. If the production database is unavailable due to outage (either planned or unplanned), these copies enable Data Guard to switch any standby database to the production role, thereby minimizing the downtime associated with the outage.

DataFlow

A Dataflow displays the current level of activity. As the rate of data transfer increases, so too does the speed of the flow. If the statistic that the flow represents moves to another threshold, the flow may change color. The combination of movement and color makes it easy to spot congested areas. A graph above the flow shows how the load has varied over time.

Deadlock

A deadlock occurs when two or more session are waiting for data locked by each other, resulting in a situation where all sessions are being blocked.

Oracle automatically detects and resolves deadlocks by rolling back the statement associated with the transaction that detects the deadlock.

Disk Transfer Time

Reading or writing data requires a disk to access the disk sector where the requested data resides. After this sector is accessed, the amount of time required for a disk to read or write data from or to storage media is referred to as disk transfer time.

Transfer time, usually expressed in milliseconds, is part of the disk access time, that is, the total time required for the computer to process the data request from the processor and then retrieve the needed data from a storage device.

Disk Utilization

The percentage of elapsed time during which a disk is busy servicing I/O requests.

DiskPerf

A Windows command line utility that enables or disables the collection of I/O statistics.

Dispatcher

Dispatcher is a background process, which is responsible for routing requests from connected user processes to available shared server processes, and for returning the responses back to the appropriate user processes.

The dispatcher process, which is only utilized with shared server configuration, handles and directs multiple incoming network session requests to shared server processes. At least one dispatcher process has to be created for every communication protocol available on the server.

Drilldown

A Foglight *for Oracle* view that provides more detailed information about a particular element (for example, a monitored object or an alarm) than the current view or dashboard. Foglight *for Oracle* drilldowns often contains charts or tables showing Oracle or Windows statistics or objects.

External Procedures

An external procedure, also sometimes referred to as an external routine, is a procedure stored in a dynamic link library (DLL) on Windows or shared library under UNIX. The external procedure is registered with the base language and then invoked to perform special-purpose processing.

Because external procedures run in a process separated from the database instance, using these procedures ensures that any problems on the client side do not adversely affect the database.

Foglight Agent Manager

The Foglight Agent Manager is a client application that manages Foglight agents installed on monitored hosts. It provides a centralized communications link between the Foglight Management Server and the agents. The Foglight Agent Manager also provides a number of support services such as the ability to deploy, upgrade, and configure agents.

For further details, see the *Foglight Agent Manager Installation Guide*.

Foglight Management Server

The Foglight Management Server (FMS) is the central component of Foglight. The Management Server receives information from agents and makes it available in the browser interface.

The Foglight database stores all system, application, and performance data. Over time, it becomes an invaluable source of historical information for planning future system capacity requirements and for doing point-in-time analysis.

Host Name

The name of the client computer that established the Oracle connection. This name appears in the Foglight *for Oracle* Sessions, Locks, and Blocking views.

Host Process

A string that uniquely identifies the instance of the client application that established the Oracle connection. This string appears in the Foglight *for Oracle* Sessions, Locks, and Blocking views.

I/O Wait

I/O wait events take place when the session is waiting for the completion of input/output operations. The length of these wait events (the amount of time spent) depends upon the number of concurrent processes and threads requesting CPU time.

This metric's value should be inspected in conjunction with the value of the Run Queue Length.

Kernel Memory

The physical memory allocated to Windows kernel.

Kernel Mode

See [Privileged Mode](#) on page 189.

Kill

A Transact-SQL statement that terminates an Oracle connection. Any outstanding transactions for the selected session are rolled back and all locks are released.

LRU

Least Recently Used. An algorithm used by Oracle for making room for new information in the memory space allocated, by replacing the oldest (least recently used) data.

LSN

Log Sequence Number. This number, which uniquely identifies a point in a database's Log, is used for determining when a page was last modified.

Latch Wait

Latch wait events occur when a session needs to acquire a latch that is held by another session. Latches are simple, low-level serialization mechanisms to protect shared SGA data structures and shared code segments from simultaneous session access. They are similar in purpose to locks: latches protect internal memory structures while locks protect data structures. Latch wait events are designed to be very quickly acquired and freed.

Licensing

The legal agreement that determines how many users can access each Oracle installation. See glossary definitions of [Per-Seat Licensing](#) on page 189 and [Per Server Licensing](#) on page 189.

Lock Wait

Lock wait events occur when a session attempts to acquire a lock that is already held by another session. A session will be blocked until the blocking session releases the lock. Locks are designed to ensure data integrity by limiting simultaneous data access.

Multi-user database locking generally consists of two levels: exclusive locks and share locks. You want to watch out for exclusive locks (that is, TX) as they prohibit resource sharing. For example, the first transaction that exclusively locks a resource is the only one that can alter the resource (except for the DBA) until the exclusive lock is released. Share locks, unlike exclusive locks, allow a resource to be shared.

Most locking issues are application-specific and can be addressed by tuning the concurrency logic in the application.

Logical Reads

The number of data blocks read by Oracle, which were already present in the buffer cache, thereby saving input/output read operations.

This value is also referred to as session logical reads and is the total of Current Reads plus Consistent Reads.

Metric

A unit of measurement that can be applied to a database. Metrics can help gauge the performance of a system.

A metric is an individual piece of information that Foglight *for Oracle* collects about the performance of a system. The information may be a numeric value (a number or percentage), a string of text, or some other piece of data.

Every time the Foglight *for Oracle* dashboard is refreshed, the cartridge retrieves the latest value of the metric, which can then be displayed in a drilldown or on the home page.

Network Wait

Network wait events occur when a session spends time waiting for messages to be sent or received over the network interface.

Network performance, which is measured in number (per second) of packets sent and received, can be used just like disk statistics to detect overload or non-optimal performance within a network or a network interface. For more information, see [SQL *net](#) on page 191.

Excessive network wait can result from either:

- Excessive network usage, originating in the application
- Physical issues, identifiable by network errors and network collisions

OLTP

Online Transaction Processing. OLTP allows real-time processing of SQL transactions, in order to support Customer Relationship Management (CRM), ERP, and other time-critical applications. OLTP is characterized by high rates of index lookups, single row modifications, and frequent commits.

Because real-time transaction processing is being increasingly carried out on a network and may include more than one company, OLTP databases use client/server processing and allow transactions to run on different platforms in a network.

Other Wait

Other wait events refer to time spent waiting for miscellaneous operations to complete. None of these operations fits into the separately identified wait event.

Paging

Disk I/O activity performed by the operating system to manage its virtual memory. High paging rates can adversely affect performance.

Panel

A group of related components on the Foglight *for Oracle* screens.

Per-Seat Licensing

The Oracle licensing mode that allows multiple instances of Oracle to be deployed, and allows specific users to access all of them. See [Licensing](#) on page 187 and [Per Server Licensing](#) on page 189.

Per Server Licensing

The Oracle licensing mode that allows a stated maximum number of simultaneous users to be supported for a specific installation of Oracle. See [Licensing](#) on page 187, and [Per-Seat Licensing](#) on page 189.

Physical Reads

The number of data blocks read by Oracle, which are not present in the buffer cache or in the PGA (Program Global Area) memory, resulting in an I/O read operations.

While physical read operations are inevitable in certain situations, they should generally be avoided.

PID

Process ID. A unique number that identifies a Windows process at any given point in time.

Privileged Mode

An operational state of hardware or software that has the highest priority. This mode is primarily used by operating system services such as I/O and paging. Also known as Supervisor Mode and Kernel Mode. See [User Mode](#) on page 191.

Process

An instance of an application executing in Windows.

Random I/O

I/O in which a specific disk block is directly accessed. This is typically the I/O that results from index lookups.

Recompile

The process of compiling a stored procedure part way through that procedure's execution.

Recursive Calls

When Oracle needs to make a change to the tables maintained for internal processing, it internally generates an internal SQL statement, which in turn generates a recursive call.

Recursive calls occur because of cache misses and dynamic storage extension. If the dictionary data is found in cache, a recursive call is not made and the data is read from cache directly.

Referential Integrity

Referential Integrity ensures that foreign keys correctly map to primary keys. A referential constraint prevents the insertion or update of foreign keys for which there are no matching primary keys. It either prevents the deletion of primary keys if foreign keys exist, or deletes these foreign rows (DELETE CASCADE).

Relational Data Engine

A major functional part of Oracle. Responsible for the parsing and optimization of SQL requests, controls query plan execution, and processes row sets from the storage engine.

SQL*Net

SQL *net is the Oracle network interface that allows Oracle tools running on network workstations and servers to access, modify, share, and store data on other servers. SQL *net is considered part of the program interface in network communications.

Scattered Read

I/O operation in which multiple disk blocks are read and scattered into different discontinuous buffers in the SGA. This is typical of the I/O that results from performing operations that require processing large amounts of data, such as full table scans, sorting, or Oracle Parallel Queries (OPQs).

Schema Locks

A lock mode that is used when compiling a query, to prevent changes to the underlying tables structures while the query is compiled.

Sequential Read

I/O operation in which disk blocks (usually single blocks) are read in sequence into contiguous memory space. This is typical of the I/O that results, for example, from using indexes.

Severity

Represents how critical an alarm is. A severity determines how Foglight *for Oracle* behaves when the values for a metric fall within a user-defined range of values. A severity specifies whether the information returned in the metric represents normal or abnormal behavior for the system under diagnosis. For example, unusually high values might mean that a metric has crossed a threshold into a high severity state. This, in turn, could change the color of a component on the home page; for example, from orange to red when moving from a critical to a fatal state.

The severity determines what action Foglight *for Oracle* takes when a metric value falls into the range defined by a threshold.

Spinner

A Spinner displays the current level of activity for a statistic that is not directional. As the load increases, so too does the speed of the spin. If the statistic represented by the flow crosses another threshold, the spinner may change color. The combination of movement and color makes it easy to spot congested areas.

SQL *net

SQL *net is the Oracle network interface that allows Oracle tools running on network workstations and servers to access, modify, share, and store data on other servers. SQL *net is considered part of the program interface in network communications.

Threshold

A range of values for a metric. If the metric falls within this range, Foglight *for Oracle* checks the threshold's severity to determine how to behave. For example, the component that represents this metric might change color.

Transaction

A group of one or more database modification statements that are combined into a logical unit of work that is either wholly committed or rolled back.

Unused Space

Disk space within a Database or File Group that is allocated to a table or index, but currently does not have any information stored in it. It is free space that can only be used by the table or index to which it is allocated.

User Mode

A processor mode that is primarily used for executing application code. Unlike the Privileged (Kernel) Mode, processes running in the user mode do not interact directly with the operating system's kernel, memory, or hardware directly. As a result, whereas a crash of a kernel mode process in most cases leads to the entire system's crash, a crash of a user mode process usually only terminates the application that was run using this mode. See [Privileged Mode](#) on page 189.

Dell listens to customers and delivers worldwide innovative technology, business solutions and services they trust and value. For more information, visit www.software.dell.com.

Contacting Dell

Technical support:

[Online support](#)

Product questions and sales:

(800) 306-9329

Email:

info@software.dell.com

Technical support resources

Technical support is available to customers who have purchased Dell software with a valid maintenance contract and to customers who have trial versions. To access the Support Portal, go to <https://support.software.dell.com/>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. In addition, the portal provides direct access to product support engineers through an online Service Request system.

The site enables you to:

- Create, update, and manage Service Requests (cases)
- View Knowledge Base articles
- Obtain product notifications
- Download software. For trial software, go to [Trial Downloads](#).
- View how-to videos
- Engage in community discussions
- Chat with a support engineer