

# Oracle® Cloud

## Using Oracle IT Analytics



E60701-42  
August 2020



Oracle Cloud Using Oracle IT Analytics,

E60701-42

Copyright © 2015, 2020, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

# Contents

## Preface

---

Audience	vii
Documentation Accessibility	vii
Related Resources	vii
Conventions	viii

## 1 Get Started with Oracle IT Analytics

---

About Oracle IT Analytics	1-1
About Oracle IT Analytics Roles and Users	1-3
Target Audiences	1-3
Before You Begin with Oracle IT Analytics	1-4
Configure Oracle IT Analytics	1-6
Pre-installation and Planning Considerations	1-7
Perform the Prerequisite Tasks	1-8
Perform ITA Specific Configuration Tasks	1-13
Disable and Re-Enable SQL Execution Plan Collections	1-13

## 2 Administer Oracle IT Analytics

---

Typical Tasks for Administering Oracle IT Analytics	2-1
Configure Capacity Thresholds for Your Environment	2-2
About Alert Thresholds and Notifications	2-4
Modify an Alert Rule	2-4
Set Up Notification Channels	2-5
Create a Notification Channel	2-6
About Viewing Entity Card Details	2-7
View Entity Card Details	2-7

## 3 Optimize IT Resources

---

Read This Before You Begin	3-1
Analyze Database Resources	3-2

Typical Tasks for Analyzing Database Resources	3-3
Analyze Database Resource Usage	3-3
Analyze Database CPU Usage	3-5
Analyze Available CPU Resources	3-5
Compare CPU Utilization Between Databases	3-5
Analyze Current CPU Trend and Forecast Utilization	3-6
Analyze Database Memory Usage	3-7
Analyze Available Memory Resources	3-7
Compare Memory Utilization Between Databases	3-8
Analyze Current Memory Usage Trend and Forecast Utilization	3-8
Analyze Database I/O Activity	3-9
Analyze I/O Trend Over Time	3-9
Compare I/O Activity Between Databases	3-9
Analyze Current I/O Trend and Forecast Activity	3-10
Analyze Database Storage Usage	3-11
Analyze Available Storage Resources	3-11
Compare Storage Utilization Between Databases	3-12
Analyze Current Storage Trend and Forecast Utilization	3-12
View Tablespace Breakdown by Database	3-13
Analyze Database Interconnect Activity	3-14
Analyze Interconnect Trend Over Time	3-14
Compare Interconnect Activity Between Databases	3-15
Analyze Current Interconnect Trend and Forecast Activity	3-15
Analyze App Server Resources	3-16
Typical Tasks for Analyzing App Server Resources	3-16
Analyze App Server Resource Usage	3-17
Analyze Resource Usage Across Application Server Instances	3-17
Analyze Resource Usage by Date Ranges	3-17
Analyze Resource Usage by Specific Time Periods	3-18
Analyze Heap Usage	3-18
Monitor Total Heap Usage Across Servers	3-19
Monitor Heap Usage By Server	3-19
Monitor the Current Trend and Forecast Heap Usage	3-20
Analyze Memory Usage	3-21
Monitor Total Memory Usage Across Hosts	3-21
Monitor Memory Usage by Host	3-22
Monitor the Current Trend and Forecast Memory Usage	3-22
Analyze CPU Usage	3-24
Monitor Average CPU Usage Across Servers	3-24
Monitor CPU Usage By Server	3-24
Monitor the Current Trend and Forecast CPU Usage	3-25

Analyze Host Resources	3-26
Typical Tasks for Analyzing Host and Virtual Machine Resources	3-27
Analyze Host CPU and Memory Usage	3-27
Analyze Host Storage Usage	3-28
Analyze Virtual Machine CPU and Memory Usage	3-28
Analyze Exadata Resources	3-29
Typical Tasks for Analyzing Exadata Resources	3-29
Analyze Exadata Systems with the Integrated View of Logs and Analytics	3-30
Analyze the Top Resource Consuming Exadata Systems	3-33
Analyze the Critical Exadata Systems	3-33
Analyze the Spare Capacity for New Workloads	3-34

## 4 Maximize Performance and Availability

---

Read This Before You Begin	4-2
Maximize Database Performance	4-2
Typical Tasks for Maximizing Database Performance	4-3
Analyze Database Performance Degradation	4-3
Analyze SQL Execution Performance	4-5
Analyze SQL Execution Performance According to Database	4-5
Analyze SQL Execution Performance Across Databases	4-5
Analyze Database Inefficiency	4-5
Analyze Inefficiency Across Databases	4-5
Analyze Inefficiency According to Database	4-6
Analyze SQL Performance	4-6
Typical Tasks for Analyzing SQL Performance	4-7
Monitor SQL Performance Across Databases	4-7
Findings and Performance Analysis	4-9
Analyze the Performance of SQL Statements	4-15
Maximize Application Performance	4-16
Typical Tasks for Maximizing Application Performance	4-17
Analyze Page Performance	4-17
Analyze Server Request Performance	4-19
Maximize Application Server Performance	4-21
Typical Tasks for Maximizing Application Server Performance	4-22
Analyze Application Server Performance	4-22
Analyze Application Server Workload Performance	4-23
Analyze IT Server and Application Availability	4-25
Typical Tasks for Analyzing IT Server and Application Availability	4-25
Assess the Health of Your IT Servers and Applications	4-26

## 5 Explore Data and Build Dashboards

---

Read This Before You Begin	5-2
Typical Tasks for Searching and Exploring Data	5-2
Access Data Explorer	5-3
Search, Explore, and Analyze Data Using Data Explorer	5-3
Save Your Analysis	5-5
Build Dashboards	5-6

## A Data Collection and Storage

---

Linux Host	A-2
Oracle Access Manager	A-5
Oracle Automatic Storage Management	A-6
Oracle Cluster Node	A-9
Oracle Clusterware	A-9
Oracle Database	A-10
Oracle Engineered System Cisco Switch	A-19
Oracle Exadata Storage Server	A-21
Oracle Hadoop HDFS	A-23
Oracle Hadoop NameNode	A-24
Oracle Hadoop Yarn	A-26
Oracle Home	A-26
Oracle as Host	A-27
Oracle HTTP Server	A-29
Oracle InfiniBand Switch	A-34
Oracle Listener	A-34
Oracle Siebel Component	A-35
Oracle Siebel Enterprise	A-36
Oracle Siebel Server	A-36
Oracle SOA Infrastructure	A-37
Oracle Traffic Director Configuration	A-38
Oracle WebLogic Server	A-41
Service Bus	A-41
Storage Manager	A-42
Windows Host	A-43

# Preface

Oracle IT Analytics gives you complete insight into the performance, availability, and capacity of your applications and IT resources. It describes how to find systemic issues, analyze resource usage across application tiers, and forecast future demand for IT services based on historical performance trends.

## Topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Resources](#)
- [Conventions](#)

*Using Oracle IT Analytics* describes how to find systemic issues, analyze resource usage across application tiers, and forecast future demand for IT services based on historical performance trends.

## Audience

*Using Oracle IT Analytics* is intended for IT analysts, IT administrators, IT capacity planners, and application owners who want to plan their IT infrastructure resources, analyze resource utilization, identify resource bottlenecks, and perform what-if analyses for contingency planning, to align their IT resource needs with current and forecasted demand.

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

### Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

## Related Resources

Apart from *Using Oracle IT Analytics*, see the following resources for information about Oracle Management Cloud.

- <http://cloud.oracle.com>

- Understand the Architecture of Oracle Management Cloud in *Installing and Managing Oracle Management Cloud Agents*
- About Oracle Application Performance Monitoring in *Using Oracle Application Performance Monitoring*
- About Oracle Log Analytics in *Using Oracle Log Analytics*
- Video: Maximizing Your Exadata Investments With Oracle Management Cloud  [Video](#)

## Conventions

The following text conventions are used in this document:

Convention	Meaning
<b>boldface</b>	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

# 1

## Get Started with Oracle IT Analytics

Oracle IT Analytics gives you a 360-degree insight into the performance, availability, and capacity of your applications and IT resources. It enables line of business executives, analysts, and administrators to make critical decisions about their IT operations, based on comprehensive system and data analysis.

By using Oracle IT Analytics, you can find systemic issues, analyze resource usage across application tiers, and forecast future demand for IT services based on historical performance trends.

### Note:

Oracle IT Analytics retains the raw data (the metric data that is collected and uploaded by agents on the target hosts) for eight days. Before purging the raw data, it rolls up the data into hourly and daily intervals. The rolled up data is then retained for 13 months to provide historical metric information. Data becomes inaccessible at the end of the 13-month period.

### Topics:

- [About Oracle IT Analytics](#)
- [Target Audiences](#)
- [Before You Begin with Oracle IT Analytics](#)
- [About Oracle IT Analytics Roles and Users](#)
- [Configure Oracle IT Analytics](#)

## About Oracle IT Analytics

Oracle IT Analytics enables you to view information about the availability, capability, and performance of your IT resources through rich, powerful analytics applications and data explorers.

Oracle IT Analytics includes the following applications and data explorers:

### Resource Analytics

- **Database Resource Analytics:** Analyze resource utilization trends of your database systems and compare them with available capacity. Perform a side-by-side comparison to identify which database consumes more resources and how the trend of resource consumption has changed. In addition, you can optimize your IT resources by identifying under-utilized and over-utilized resources, analyze database resources by different dimensions, such as cost center, and database version. You can also plan and forecast future resource usage based on either a single resource target or a group of targets.

- **App Server Resource Analytics:** Analyze the utilization of your app server (middleware) resources, and forecast utilization based on current trends. You can also analyze the resources by different usage dimensions, such as heap, memory, and CPU.
- **Host Resource Analytics:** View and analyze utilization trends for critical server resources such as CPU, Memory, and Storage. You can also analyze, compare, and contrast resource usage across servers.
- **Exadata Analytics:** Gain visibility into the performance of your workloads that run on a shared Exa infrastructure. You can analyze the top resource consuming Exadata systems, the critical Exadata systems, and the spare capacity for new workloads. This analysis helps you optimize the capacity based on proactive planning.

### Performance Analytics

- **Database Performance Analytics:** Perform a historical, comparative analysis to find any bottlenecks affecting enterprise-wide database performance. You can also determine the root cause of recurring systemic problems, identify the performance profile and activity types for CPU or I/O intensive database activities, and analyze SQL queries across multiple databases to identify which SQL statements consume the most database time and to help understand the trend of each query's performance.
- **Application Performance Analytics:** Monitor the performance of your applications through viewing the page performance, server request performance, page views and load time, page load time, tier response breakdown, error distribution, and a profile of the users.
- **SQL Analytics:** Analyze the SQL performance problems for enterprise-wide applications across a fleet of databases. In addition, view trends and key insights to SQL performance issues to proactively avoid future problems.
- **App Server Performance Analytics:** Draw analytical insights for heterogeneous application servers including garbage collection overhead, workload, resource usage, availability, configuration profiles, and so on.
- **Availability Analytics:** Gain visibility in to the overall health of your IT systems and applications. This application integrates with data collected from Oracle Application Performance Monitoring.

### Data Explorers

- **Search:** Search, browse, and compare systems running in your on-premises or cloud environment. With Data Explorer: Search, you can analyze your IT assets, including the distribution of Oracle WebLogic Servers by versions across enterprise-wide hardware platforms and configuration of database servers with associated core count, applied patches, and installed options. You can view real-time performance snapshots, such as the top 25 Oracle WebLogic Server workload by requests or heap utilization.

In addition, to take complete advantage of the capabilities of Oracle IT Analytics, use it with Oracle WebLogic Server Management Pack and Oracle Database Diagnostics Pack. See [Pre-installation and Planning Considerations](#).

You can access the Oracle IT Analytics applications from the home page of Oracle Management Cloud (OMC). Select the application of your choice from the **IT Analytics** drop-down list on the OMC home page. Alternatively, you can also access

these applications from the OMC Navigation menu. Click **IT Analytics** from the menu, and then select the application of your choice.

## About Oracle IT Analytics Roles and Users

If you're a new customer and have purchased an Oracle Management Cloud license edition such as Standard edition or Enterprise edition, then after the instance is created, the following roles are provisioned:

- Oracle Management Cloud <instance name> Administrator
- Oracle Management Cloud <instance name> User

For more information about the tasks that users assigned the above roles can perform, see Add Users and Assigning Roles in *Getting Started with Oracle Management Cloud*.

However, if you're an existing customer and you've purchased the standalone Oracle IT Analytics service, then the following roles are created:

Role	Tasks
Oracle IT Analytics Administrator	Set up, enable, and manage Oracle IT Analytics on entities. See Manage Entities Using the Administration Console in <i>Working with Oracle Management Cloud</i> .
Oracle IT Analytics User	<ul style="list-style-type: none"><li>• Analyze and plan for future growth by forecasting consumption of resources. See <a href="#">Optimize IT Resources</a>.</li><li>• Visualize and analyze performance findings and trends across databases and middleware resources. See <a href="#">Maximize Performance and Availability</a>.</li></ul>

### Note:

The Oracle IT Analytics Administrator can create users and assign one of the Oracle IT Analytics roles to the newly created users. See Managing User Accounts and Managing User Roles in *Managing and Monitoring Oracle Cloud*.

## Target Audiences

Oracle IT Analytics offers analytics solutions targeted at the following audiences:

Audience	Tasks
 <b>IT Analyst</b>	<ul style="list-style-type: none"> <li>Analyze Database Resources</li> <li>Analyze App Server Resources</li> <li>Explore Data and Build Dashboards</li> <li>Analyze IT Server and Application Availability</li> <li>Analyze Host Resources</li> </ul>
 <b>IT Administrator/Database Administrator/Middleware Administrator</b>	<ul style="list-style-type: none"> <li>Analyze Database Resources</li> <li>Analyze App Server Resources</li> <li>Maximize Database Performance</li> <li>Maximize Application Server Performance</li> <li>Viewing the trend of resource utilization by all database and middleware targets and forecasting utilization.</li> <li>Identifying which databases have a specific patch or patch set applied.</li> <li>Viewing the top database and middleware targets, based on specific metrics. For example, the top 25 WebLogic Servers by heap.</li> <li>Analyze IT Server and Application Availability</li> <li>Analyze Host Resources</li> </ul>
 <b>IT Capacity Manager</b>	<ul style="list-style-type: none"> <li>Analyze Database Resources</li> <li>Analyze App Server Resources</li> <li>Viewing the trend of resource utilization by all database and middleware targets and forecasting utilization.</li> <li>Viewing the top database and middleware targets, based on specific metrics. For example, the top 25 databases by CPU utilization.</li> <li>Maximize Application Performance</li> <li>Analyze IT Server and Application Availability</li> <li>Analyze Host Resources</li> </ul>
 <b>Application Owner</b>	<ul style="list-style-type: none"> <li>Maximize Application Performance</li> </ul>

## Before You Begin with Oracle IT Analytics

Here are some basic terminologies about Oracle Management Cloud.

Term	Definition
Agents	Oracle Management Cloud agents collect configuration, performance, availability, and log data from monitored entities and make this information available in Oracle Management Cloud.

Term	Definition
Alerts	Information generated in response to an availability issue or when a metric crosses its thresholds. Conditions for generating alerts are defined in Alert Rules. Alerts sent to administrators by using various channels, such as, email and SMS are known as <b>notifications</b> .
Alert Rules	A set of rules that define the conditions under which alerts are generated and notifications sent when an alert is raised. Alert conditions consist of a metric, a comparison operator, and thresholds against which metric values are evaluated.
Associations	Associations (association instances) define a relationship between two managed entities. The association type that you define, either via the user interface or based on a written document Oracle provides, determines how data is correlated and visualized in Oracle Management Cloud. In many cases, associations are defined automatically by Oracle Management Cloud.
Entities	Entities are monitored resources such as databases, host servers, compute resources, or application servers.
Entity Types	Entity types are a type of monitored resource, such as a host or database, which define where that entity fits in the Oracle Management Cloud hierarchical structure. In Oracle Management Cloud, each entity is defined by a set of characteristics, it has a parent and may have other children. For example, a generic host is an operating system (OS) independent target and it has children entities that are specific OS hosts, such as Linux and Windows. The metrics collection functionality takes advantage of this inheritance model so each monitored entity has entity-specific metrics as well as metrics inherited from each level it descended from. For example, Oracle Management Cloud collects metrics at level three that are common to all generic hosts, independent of the vendor. A Linux host, since its parent is a generic host, inherits all the metrics collected for generic hosts and its ancestors, as well as Linux-specific ones, if any.

Term	Definition
JSON	JavaScript Object Notation (JSON) allows data to be concisely and precisely defined in a format that is both human and machine-readable. Oracle provides sample JSON files for defining entities. JSON files are then edited with your own custom parameters and are passed on to agents. This configuration step defines the entities with that agent and Oracle Management Cloud.
License Editions	License editions are pre-defined categories of Oracle Management Cloud offerings.
omcli	Oracle Management Cloud agent control command line interface utility (omcli) is used to interface with Cloud agents and define entities using customized JSON files.
Data Collector	A Data Collector agent collects data from your on-premises OMR and uploads it to Oracle Management Cloud.
Gateway	A Gateway agent acts as a channel between Oracle Management Cloud and other Cloud agents. Multiple Data Collector or Cloud agents can communicate with Oracle Management Cloud through a single Gateway.
Cloud Agent	A Cloud agent collects the host, entity, and log data from the host where you deploy the Cloud agent. It can connect to Oracle Management Cloud directly or through a Gateway.
Oracle Management Repository (OMR)	OMR is a schema in an Oracle Database where all the information collected by Oracle Enterprise Manager Cloud Control Management Agents is stored. It consists of objects such as database jobs, packages, procedures, tables, views, tablespaces, and so on.

## Configure Oracle IT Analytics

To get started with Oracle IT Analytics, you must install the cloud agent and enable Oracle IT Analytics on entities.

In addition, Oracle recommends a few pre-installation and planning considerations that help you leverage the complete capabilities of Oracle IT Analytics.

### Topics:

- [Pre-installation and Planning Considerations](#)
- [Perform the Prerequisite Tasks](#)
- [Perform ITA Specific Configuration Tasks](#)
- [Disable and Re-Enable SQL Execution Plan Collections](#)

## Pre-installation and Planning Considerations

Oracle recommends the following when planning your OMC setup, Oracle IT Analytics in particular:

- Oracle IT Analytics App Server Performance Analytics Application and Oracle WebLogic Server Management Pack:**  
 To take advantage of the WebLogic metrics collected by Oracle Enterprise Manager with the Oracle WebLogic Server Management Pack, plan to install a Data Collector. The Data Collector harvests data from the Oracle Enterprise Manager repository. If you do not have Oracle Enterprise Manager and the Oracle WebLogic Server Management Pack, you can deploy an OMC cloud agent with your WebLogic Server and benefit from the Oracle IT Analytics App Server Performance Analytics application capabilities.

- Oracle IT Analytics Database Applications and Oracle Database Diagnostics Pack:**  
 Leverage the full capabilities of the Oracle IT Analytics database applications with the Oracle Diagnostic Pack by installing a Data Collector.

The Oracle Diagnostic Pack provides automatic performance diagnostic and advanced system monitoring functionality. By configuring Oracle Management Cloud to collect data from the Oracle Diagnostics Pack, Oracle IT Analytics database applications can display and analyze historic AWR data.

- Oracle IT Analytics Data Collectors for ITA Exadata Systems**  
 Exadata Analytics presents a fully unified view of your inventory, capacity and performance data for all Exadata systems, on-premise Exadata Database Machine as well as Exadata CS and Exadata CC. Exadata Analytics gives you a clear insight into Exadata storage systems, allowing you to do in-depth resource analysis, capacity planning and forecasting for all your Exadata systems.

To leverage the capabilities of Oracle IT Analytics Exadata Analytics for all on-premise and cloud Exadata systems you will need to have data collectors deployed. For more information see: [Perform ITA Specific Configuration Tasks](#).

The following table lists the availability of the database applications for different editions of Oracle Database, with and without the Oracle Diagnostic Pack.

ITA Apps	Oracle DB Standard Edition	Oracle DB Enterprise Edition without the Oracle Diagnostic Pack	Oracle DB Enterprise Edition with the Oracle Diagnostic Pack
Database Performance Analytics	Yes	Yes	Yes
Database Resource Analytics*	Yes	Yes	Yes
SQL Analytics	No	No	Yes
Data Explorer for AWR metrics*	No	No	Yes
Exadata Analytics	NA	NA	Yes

\* Interconnect metrics in Database Resource Analytics requires Oracle Diagnostic Pack. Without the Oracle Diagnostic Pack, the interconnect portion will be blank.

\*Additional metrics related to AWR is available in Data Explorer only with the Oracle Diagnostics Pack when database entities are enabled for Oracle Management Cloud Enterprise Edition (which includes Oracle IT Analytics).

The SQL Analytics application requires the Oracle Diagnostics Pack for all its data.

Oracle IT Analytics database applications do not require the Oracle Database Tuning Pack.

Note that, before you set up Oracle Management Cloud, to enable the Oracle Diagnostics Pack, you must purchase the necessary licenses on the target databases and set the initialization parameter `CONTROL_MANAGEMENT_PACK_ACCESS` to the desired value. See Oracle Management Packs and Their Permitted Features in *Oracle Database Licensing Information* for information on Oracle Diagnostics Pack.

## Perform the Prerequisite Tasks

To get started with Oracle IT Analytics, you must enable the license editions and install the cloud agents. Cloud agents collect performance and metric data from entities that you want to monitor and upload the data to Oracle Management Cloud. To complete these tasks, you must have the *Oracle Management Cloud Administrator* role.

If the *Oracle Management Cloud Administrator* role isn't assigned to you or you're not sure, then ask your system administrator to ensure that the correct role is assigned to you in Oracle Management Cloud. See [Modifying User Roles](#) in *Managing and Monitoring Oracle Cloud* and [About Oracle IT Analytics Roles and Users](#).

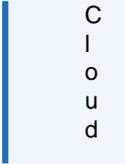
You can use Oracle IT Analytics to monitor Autonomous Databases, which include Autonomous Data Warehouse – Shared and Autonomous Transaction Processing – Shared. The prerequisite tasks to set up Oracle Management Cloud to work with Oracle Cloud Infrastructure, discover Autonomous Databases, and assign the required privileges to users are different from the tasks listed in the following table. See [Set Up Oracle Database Management for Autonomous Databases](#) in *Using Oracle Database Management for Autonomous Databases*.

Task	Required / Optional	Description
<i>Task 1:</i> Install a gateway on a host in your data center (the host should have internet access to Oracle Management Cloud)	Optional	A gateway acts as a channel between Oracle Management Cloud and all other Oracle Management Cloud agents. See <i>Install a Gateway</i> in <i>Installing and Managing Oracle Management Cloud Agents</i> .

Task	Required / Optional	Description
<p>Task 2: Install a data collector on the target host (a host on which the Oracle Management Repository is present or another host with SSH access to the remote host)</p>	<p>Required when using Oracle Enterprise Manager to harvest entities into Oracle Management Cloud</p>	<p>A data collector is required <i>only</i> if you want to collect data from entities that are monitored by an existing on-premises Oracle Enterprise Manager Cloud Control. See <i>Install a Data Collector</i> in <i>Installing and Managing Oracle Management Cloud Agents</i>.</p>

 **Note**: The Data Collector is optional if you are not harvesting

Task	Required / Optional	Description
		ti n g e n t i t i e s f r o m O r a c l e E n t e r p r i s e M a n a g e r i n t o O r a c l e M a n a g e m e n t

Task	Required / Optional	Description
		
Task 3: Install Cloud Agents.	Required, Cloud Agents <b>must</b> be installed as Data Collectors by themselves are not sufficient.	<p>If you are using Oracle Enterprise Manager and installing a Data Collector, you are required to install Cloud Agents on the same hosts where the EM agents are installed. The Cloud Agents collect performance data from Oracle Database targets for Database Performance Analytics and SQL Analytics applications.</p> <p>If you are <b>not</b> using Oracle Enterprise Manager and installing a Data Collector, install Cloud Agents on hosts where your entities to be discovered reside.</p> <p>See Install Cloud Agents in <i>Installing and Managing Oracle Management Cloud Agents</i>.</p>
Task 4: Define monitoring credentials	Required only if <b>not</b> using a Data Collector	Monitoring credentials are required to collect performance and availability metrics by using cloud agents. See Prerequisites and Monitoring Credentials in <i>Using Oracle Infrastructure Monitoring</i> .
Task 5: Enable secure database communication on the cloud agent	Required for database entities if you are using TCP settings for the listener.	Enable the TCPS connections to facilitate communication between the cloud agent and the secure database. See <b>Enabling TCPS Connections</b> under Oracle Database in <i>Using Oracle Infrastructure Monitoring</i> .

Task	Required / Optional	Description
Task 6: Add entities.	Required only if you are <b>not</b> using a Data Collector and are adding entities via a Cloud Agent.	<p>Add the entities that your cloud agent will monitor. Oracle IT Analytics makes use of the entities discovered from Oracle Enterprise Manager (via Data Collector, if you've installed one).</p> <p>If you want to monitor the performance and availability of entities by using cloud agents, add entities to Oracle Infrastructure Monitoring. See Adding Entities for Infrastructure Monitoring in <i>Using Oracle Infrastructure Monitoring</i>.</p>
Task 7: Assign licenses. <ul style="list-style-type: none"> <li>• Option 1: Auto-enable entities</li> <li>• Option 2: Assign license per entity</li> </ul>	Required	<p>Oracle IT Analytics is available to you as part of the Oracle Management Cloud – Enterprise Edition.</p> <p><b>Option 1:</b> Enable license editions from the Oracle Management Cloud console. Assign the Enterprise Edition automatically to all entities. See Auto-Assign License Editions to All Entities in <i>Getting Started with Oracle Management Cloud</i>.</p> <p><b>Option 2:</b> If you do not enable auto-assignment of license editions, then after you've added entities, you can assign Enterprise Edition license to selected entities. See Assign License Editions to Selected Entities in <i>Getting Started with Oracle Management Cloud</i>.</p>

 **Note:**

If you are using a previous version of Oracle Management Cloud, which allows you to enable a standalone service on entities, then after you've deployed your cloud agent, you can enable Oracle IT Analytics on an entity or a set of entities. You can perform this task from the Administration console of Oracle Management Cloud. See Enabling or Disabling Services on Entities in *Working with Oracle Management Cloud*.

## Perform ITA Specific Configuration Tasks

The following is a list of requirements that ITA Exadata Analytics requires in order to support Exadata Cloud Service (Exadata CS) and Cloud at Customer (Exadata CC) system information, resource utilization and performance metrics.

- IT Analytics comes included with Oracle Management Cloud Enterprise Edition License. For more information see: [Oracle Management Cloud Offerings](#).
- ITA Exadata Analytics can monitor your Exadata on premises, Exadata CS and Exadata CC systems. You will need Enterprise Manager 13.4 setup, with all types Exadata targets already discovered. For more information see: [Discover the Exadata Cloud Service Target](#) in *Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud* and [Exadata Database Machine Discovery](#) in *Oracle Enterprise Manager Oracle Exadata Database Machine Getting Started Guide*.
- A data collector must be deployed to monitor any Exadata system. Oracle Management Cloud data collector version 1.50 or higher is required for Exadata CC and Exadata CS systems. For more information see: [Install a Data Collector](#) in *Oracle Cloud Installing and Managing Oracle Management Cloud Agents*.
- For Exadata CS and Exadata CC systems Oracle Management Cloud Agents must be deployed on the user domain `domU`. For more information see: [Install Cloud Agents](#) in *Oracle Cloud Installing and Managing Oracle Management Cloud Agents*, and [Oracle VM and Oracle Exadata Database Machine](#) in *Oracle Exadata Database Machine Maintenance Guide*.
- If monitoring Exadata CS, a Service Gateway must be configured. For more information see: [Access to Oracle Services: Service Gateway](#) and [Service Gateway: Supported Cloud Services in Oracle Services Network](#).

## Disable and Re-Enable SQL Execution Plan Collections

Oracle Management Cloud automatically collects SQL Execution Plan data. By upgrading your agents to the latest version, OMC can begin collecting SQL Execution Plan data across all your databases. You can also manually disable and re-enable SQL Execution Plan data collection should your environment require this.

Collecting SQL text and plans data, across all your databases, allows you to analyze problematic SQL and then use your experience, or tools like SQL Tuning Advisor, to tune your SQL Statements. Some examples of this can be:

- Fast running queries may turn slow if the new execution plan is bad
- SQL query changing plans frequently can cause variations in performance
- Improperly parallelized or not parallelized execution plans
- Improper Optimizer configuration

 **Note:**

If you have used a custom defined user, and you've granted privileges with a previous script, you will need to grant additional privileges to avoid the following error: Error evaluating SQL Plan Details:30MinASHITACollection - FetchletException: SQLException encountered while extracting the ASH Data: java.sql.SQLException: ORA-00942: table or view does not exist. To grant these additional privileges run the following SQL grant commands:

```
grant select ON V_$SQLAREA_PLAN_HASH to <YOUR MONITOR ROLE>
grant select ON GV_$SQLAREA_PLAN_HASH to <YOUR MONITOR ROLE>
```

For CDBs you will need to add the *C##* prefix to <YOUR MONITOR ROLE>. If future users are created, use the 1.49 or later *grantPrivileges.sql* script, located in: \$ <AGENT BASE DIR>/plugins/oracle.em.sgfm.zip/<CURRENT OMC VERSION>/scripts/grantPrivileges.sql

**Disable SQL Text/Plan Collection**

If you need to disable SQL data collection, update the Oracle Management Cloud Agent that is monitoring the database with the following steps:

1. Set the `_collectASHSQLPlan` property to false:

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -
allow_new -name _collectASHSQLPlan -value FALSE
Oracle Management Cloud Agent
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.
```

```
EMD setproperty succeeded
```

Ensure that the output shows `EMD setproperty succeeded`.

2. Set the `_collectASHSQLStatement` property to false:

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -
allow_new -name _collectASHSQLStatement -value FALSE
Oracle Management Cloud Agent
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.
```

```
EMD setproperty succeeded
```

Ensure that the output shows `EMD setproperty succeeded`.

3. Set the `_maskLiteralsInSQLStatement` property to false:

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -
allow_new -name _maskLiteralsInSQLStatement -value FALSE
Oracle Management Cloud Agent
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.
```

```
EMD setproperty succeeded
```

Ensure that the output shows `EMD setproperty succeeded`.

4. Set the `_collectSQLStatement` property to false:

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -
allow_new -name _collectSQLStatement -value FALSE
Oracle Management Cloud Agent
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.
```

```
EMD setproperty succeeded
```

Ensure that the output shows `EMD setproperty succeeded`.

5. Verify that the property settings from the configuration file `emd.properties` are set to false:

```
$ cat <AGENT_BASE_DIR>/agent_inst/sysman/config/emd.properties |
grep SQL
```

The command should return the following prompt with all four values set to *false*.

```
_collectSQLStatement=false
_maskLiteralsInSQLStatement=false
_collectASHSQLPlan=false
_collectASHSQLStatement=false
```

### Re-Enable SQL Text/Plan Collection

If you need to re-enable SQL data collection, update the OMC Cloud Agent that is monitoring the database with the following steps:

#### Note:

To re-enable SQL Execution Plan data collection for Oracle databases, you must have a special database user with appropriate privileges to access this data. If you use the database user `DBSNMP`, these privileges are **not** required. The `DBSNMP` user already has all the privileges required to collect Oracle IT Analytics data.

1. Obtain the property settings from the configuration file: `emd.properties`

```
$ cat <AGENT_BASE_DIR>/agent_inst/sysman/config/emd.properties |
grep SQL
```

Here, the four properties can be seen with a *true*, *false*, or *empty* value:

```
_collectSQLStatement= <true/false/empty>
_maskLiteralsInSQLStatement= <true/false/empty>
```

```
_collectASHSQLPlan= <true/false/empty>  
_collectASHSQLStatement= <true/false/empty>
```

 **Note:**

When running this command some agents might see `_collectSQLStatement` and `_maskLiteralsInSQLStatement` already set as *true*. In this case you can skip directly to step 3.

2. Set the property values to true for `_collectSQLStatement` and `_maskLiteralsInSQLStatement` with these commands:

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -  
allow_new -name _collectSQLStatement -value true
```

```
Oracle Management Cloud Agent  
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.  
EMD setproperty succeeded
```

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -  
allow_new -name _maskLiteralsInSQLStatement -value true
```

```
Oracle Management Cloud Agent  
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.  
EMD setproperty succeeded
```

3. Set the property values to true for `_collectASHSQLStatement` and `_collectASHSQLPlan` with these commands:

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -  
allow_new -name _collectASHSQLStatement -value true
```

```
Oracle Management Cloud Agent  
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.  
EMD setproperty succeeded
```

```
$ <AGENT_BASE_DIR>/agent_inst/bin/omcli setproperty agent -  
allow_new -name _collectASHSQLPlan -value true
```

```
Oracle Management Cloud Agent  
Copyright (c) 1996, 2019 Oracle Corporation. All rights reserved.  
EMD setproperty succeeded
```

4. Verify the property settings from the configuration file `emd.properties` are set to *true*.

```
$ cat <AGENT_BASE_DIR>/agent_inst/sysman/config/emd.properties |  
grep SQL
```

The command should return the following prompt with all four values set to *true*.

```
_collectSQLStatement=true  
_maskLiteralsInSQLStatement=true  
_collectASHSQLPlan=true  
_collectASHSQLStatement=true
```

Once all four property settings are set to *true*, SQL Text/Plan collection is now enabled and will begin at the next execution cycle. Oracle Management Cloud will collect these plans and begin automatically populating the repository over time.

For further information on how SQL Text and SQL Plan Collections work, how to view and analyze the collected data and specific use cases see: [Findings and Performance Analysis](#) and [Monitor SQL Performance Across Databases](#).

# 2

## Administer Oracle IT Analytics

This section explains the tasks that help you set up Oracle IT Analytics and administer your Oracle IT Analytics warehouse. This section also explains the procedure to configure capacity thresholds for resource utilization, set alert rules and notification preferences, and access entity-related information by using the Entity Card. You must have the Oracle Management Cloud Administrator role or Oracle IT Analytics administrator role to administer Oracle IT Analytics.

### Topics:

- [Typical Tasks for Administering Oracle IT Analytics](#)
- [Configure Capacity Thresholds for Your Environment](#)
- [About Alert Thresholds and Notifications](#)
- [View Entity Card Details](#)

## Typical Tasks for Administering Oracle IT Analytics

You can perform Oracle IT Analytics specific configuration tasks such as setting capacity thresholds, alert thresholds, and notification channels for your IT environment. You must have the *Oracle Management Cloud Administrator* role or *Oracle IT Analytics administrator* role to be able to perform these tasks.

If one of these roles isn't assigned to you or you're not sure, then ask your system administrator to ensure that the correct role is assigned to you in Oracle Management Cloud. See [Modifying User Roles](#) in *Managing and Monitoring Oracle Cloud* and [About Oracle IT Analytics Roles and Users](#).

Task	Description
Configure capacity thresholds for your IT environment.	Configure the high and low resource utilization thresholds (in percentage) for critical resources such as CPU, memory, and storage across your environment.  See <a href="#">Configure Capacity Thresholds for Your Environment</a> .
Set up alert thresholds and notification channels.	Set alert rule conditions and specify your preference for receiving notifications when alerts are triggered.  See <a href="#">About Alert Thresholds and Notifications</a> .
View Entity Card details.	Use the Entity Card to view information related to specific targets, in the form of histograms.  See <a href="#">View Entity Card Details</a> .

 **Note:**

Oracle IT Analytics retains the raw data (the metric data that is collected and uploaded by agents on the target hosts) for eight days. Before purging the raw data, it rolls up the data into hourly and daily intervals. The rolled up data is then retained for 13 months to provide historical metric information. Data becomes inaccessible at the end of the 13-month period.

## Configure Capacity Thresholds for Your Environment

As an Oracle Management Cloud administrator, you can configure the high and low resource utilization threshold values (in percentage) for critical resources across your environment.

You can configure the threshold values for the following resources:

- Host CPU
- Host Memory
- Host Storage
- Exadata Storage
- Exadata IOPS

By default, the high and low values are set at 75% and 25%, respectively. You can modify these values, if required. When you modify, you can view the updated data in the scatter chart for CPU, memory, and storage; in your time-series plots; and in your Trend and Forecast charts.

Based on the actual usage of these critical resources (CPU, memory, storage, and IOPS), over a defined period of time, ITA forecasts when these resources will reach their utilization threshold values (high or low) in the next 30, 90, 180, and 180+ days. This information will help you derive insights and take appropriate actions to maximize resource performance.

To modify the capacity thresholds:

1. From Oracle IT Analytics navigation menu, click **IT Analytics Admin**, and then click **IT Analytics Settings**.

**Figure 2-1 IT Analytics Settings Page**

The screenshot shows the 'IT Analytics Settings' page with the following sections:

- Utilization Thresholds:** A table of settings for various resources. Each resource has 'High' and 'Low' threshold values with a 'Default' value below them.
 

Resource	High	Low	Default
Host CPU	75	25	75
Host Memory	75	25	75
Host Storage	75	22	75
Exadata Storage	82	22	75
Exadata IOPS	79	22	75
- Time Ranges for Reaching Utilization Thresholds:** A section with a radio button for 'Add 365+ days time range'.
- Trend and Forecast:** Includes 'Forecast Period (Day)' set to 90 and 'Confidence (%)' set to 95.
- Time Percentile:** Includes 'Time Series Percentile' with radio buttons for 90, 95 (Default), and Maximum. An 'Advanced' section below it includes 'Time Range Percentile' with radio buttons for 90 (Default), 95, and Maximum.

Alternatively, from the Oracle Management Cloud Console navigation pane, click **Administration**, click **IT Analytics Admin**, and then click **IT Analytics Settings**.

2. In the **Utilization Thresholds** section of the **IT Analytics Settings** page, modify the high (between 75 and 100) and low (between 0 and 25) values for the applicable resources.
3. **Time Ranges for Reaching Utilization Thresholds:** Enable this option if you want to view the number of database systems that are forecasted to reach high resource utilization threshold in the next 365+ days. You can view the details in the Overview page of the Database Resource Analytics application. See [Analyze Database Resource Usage](#).
4. In the **Trend and Forecast** section, specify the following details:
  - **Forecast Period:** Specify the period in days. You can use a value between 1 and 366. The default value is 30.
  - **Confidence (%):** Specify the confidence level, which indicates the percentage of confidence that future values of resource utilization are expected to fall within the estimated high and low capacity thresholds. For example, a 95% of confidence interval indicates that 95 out of 100 forecasted points are expected to fall within the high-low range.

You can set these parameters to use the default values for interactive trend and forecast in the different ITA applications.

5. In the **Time Percentile** section, specify the following details:

- **Time Series Percentile:** Specify the value to indicate the resource utilization for a time period. If you are computing the daily resource usage, the time series percentile represents the value of the metric data points of all the resources, collected in a day. A time series percentile value of 95 indicates the value below which 95% of the data points are found.
- **Time Range Percentile:** Click **Advanced**, and then specify the percentile value of resource utilization on a daily or hourly basis. This value is used to compute the aggregate resource usage. For example, if you want to arrive at a single data point for current CPU utilization for the last 30 days with the time range percentile value set at 90, then this value represents the 90th percentile value of the CPU usage metric data collected in the last 30 days.

6. Click **Save**.

You can use these threshold values to generate alerts based on the alert rules you define.

## About Alert Thresholds and Notifications

You can set alert rule conditions and specify your preference for receiving notifications when alerts are triggered. To set up alert rules for your enterprise, you must log in to the services as an *Oracle Management Cloud Administrator* user.

By default, the following alert rules are available for your use:

- Exadata IOPS Capacity Headroom Rule
- Exadata Storage Capacity Headroom Rule
- Host CPU Capacity Headroom Rule
- Host Memory Capacity Headroom Rule

You receive these alerts as early warnings that inform you about capacity utilization of your host resource types such as Exadata IOPS, Exadata storage, host CPU and host memory. You can then take remedial measures to manage capacity, such as move the workload to another host, add more capacity, or analyze the reason for increased resource usage and resolve the issue.

The alerts appear both on applications such as Host Resource Analytics and on the Oracle Management Cloud console.

## Modify an Alert Rule

*Alert Rules* enable you to define how alerts should be triggered (if they are not automatic) and how you want to get notified. You can modify the default alert rules that are predefined for you.

To modify an alert rule condition:

1. From Oracle IT Analytics navigation menu, click **IT Analytics Admin**, and then click **Alert Rules**.

Alternatively, from the Oracle Management Cloud Console navigation pane, click **Administration**, click **IT Analytics Admin**, and then click **Alert Rules**.

You can view the application-specific available alerts by clicking **Alerts** on the top left corner of the page.

- From the predefined list of alert rules in the **Alert Rules** page, click the rule that you want to modify. The details of the selected rule appear.

**Figure 2-2 Alert Rule Page in the Edit Mode**

The screenshot shows the 'Edit Alert Rule' interface. At the top right, it displays the date 'Aug 28, 2018' and a 'Settings' icon. Below the title, there are 'Save' and 'Cancel' buttons. The form contains the following fields:

- Rule Type:** Capacity Headroom (dropdown menu)
- Name:** Host CPU Capacity Headroom Rule
- Severity:** Warning
- Description:** Alert rule for hosts projected to reach high CPU utilization in the specified number of days or less
- Enable Rule:**  (with an information icon)
- Entities:** Host (Generic) (dropdown menu)
- Alert Condition:**
  - Resource Type: CPU
  - Utilization Threshold (%): 75
  - Days to Reach High Utilization: 30 (input field)

- The **Enable Rule** option (checked by default) ensures that you receive alert notifications for the rule. You can uncheck this option if you do not want to receive these notifications.
- Modify the **Days to Reach High Utilization** value per your requirement. You should specify a value greater than zero. The default value is 30. This value indicates the number of days that remain to reach the defined threshold value for resource utilization. An alert is triggered when resource utilization reaches or exceeds the threshold value before 30 days (or the number of days you specify).

 **Note:**

The **Utilization Threshold** value that you define for the resource type in the **IT Analytics Settings** page, is displayed here. See [Configure Capacity Thresholds for Your Environment](#).

- Click **Save**.

## Set Up Notification Channels

When an alert is raised, worsens in severity, or clears, you may want to be actively notified through email, by push notifications (mobile devices), or have a third-party application take action. You can define these types of notifications by using the *notification channels*.

### Types of Notification Channels

Classes of notification destinations are called *notification channels*. Notification channels allow you to set up and reuse functional groups of notification recipients, such as regional administrators, IT managers, or other Web servers without having to specify a large number of individual destinations repeatedly. After you set up a notification channel, you can reuse the channels across different alert rules.

The following notification channels are available for use:

- **Email:** Allows you to manage/logically organize the email recipients of alert notifications. For example, you can create an email channel for the *On-Call Team* that contains the email addresses of the IT operators that staff the on-call team, or a *DBA Team* email channel that contains all email addresses of database administrators in your department.
- **Mobile:** Allows you to send push notifications to one or more mobile phones that are associated with Oracle Management Cloud users.
- **Integration:** Allows you to send HTTP POST (WebHook) messages to a destination URL by integrating the alerts with your organization's existing channels.

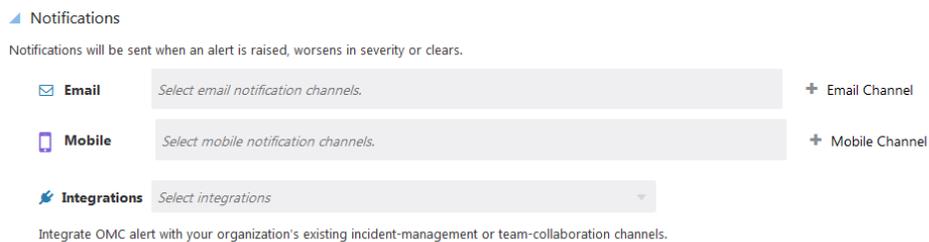
## Create a Notification Channel

You can define notification channels while creating/editing an alert rule from the Edit Alert Rule page.

To create a notification channel:

1. Expand the **Notifications** section of the Edit Alert Rule page. You can view the list of available notification channels.

**Figure 2-3 Notifications Section of the Edit Alert Rule Page**



2. Click the desired notification channel type and specify the requisite information:

**For creating an email notification channel:**

- a. Specify an intuitive name for the channel.
- b. Specify a comma-separated list of e-mail addresses.
- c. Click **Create**.

**For creating a mobile notification channel:**

- a. Specify an intuitive name for the channel.
- b. Specify a comma-separated list of recipient Oracle Management Cloud users.
- c. Click **Create**.

To receive push notification on your mobile device, you must install and remain signed in to the Oracle Management Cloud Mobile application on your device.

**For creating integrations channel:** In addition to notifying people, Oracle Management Cloud can also send relevant information to third-party Web applications (such as Slack or Hipchat) if an alert is raised, thus allowing you to extend Oracle Management Cloud functionality by having third-party applications carry out actions in response to an Oracle Management Cloud alert notification.

- a. Select the desired third-party application from the drop-down list.
3. Click **Save**.

## About Viewing Entity Card Details

You can use Entity Card to view information related to specific targets, in the form of histograms. You can access the entity-related information easily instead of going to other views or performing a separate search.

The Entity Card visualization displays an entity's status, alerts for that entity, and the associated log records (in the form of a histogram). You can also use the Entity Card to get information from other applications such as Oracle Log Analytics, Oracle Infrastructure Monitoring, and so on.

The Entity Card is available from the following applications:

- Host Resource Analytics (HRA)
- Database Resource Analytics (DBRA)
- Exadata Analytics

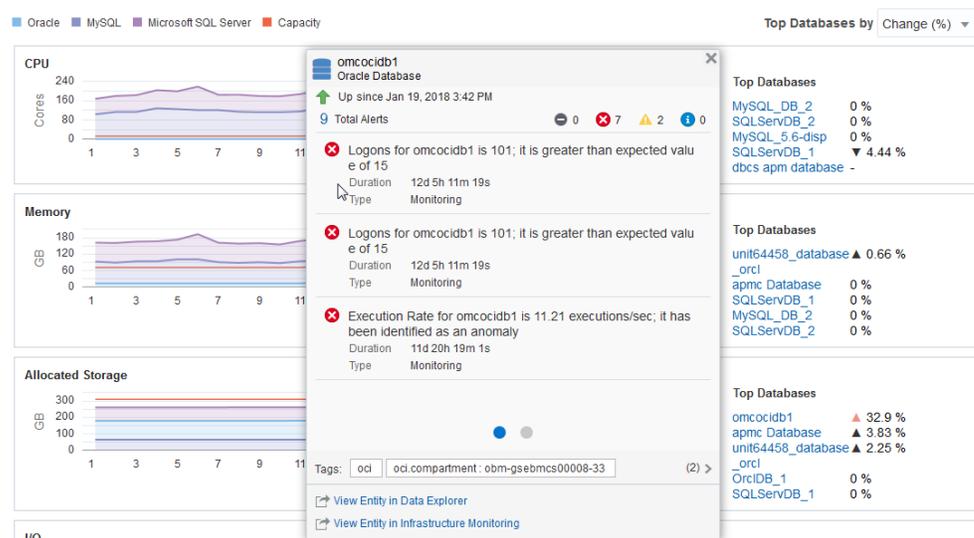
The Entity Card display provides a link to the Data Explorer application of Oracle IT Analytics from where you can view and analyze the entity. Based on the services that are enabled on the selected entity, the entity card provides a link to view the entity in that service.

## View Entity Card Details

You can view the Entity Card details from Host Resource Analytics, Database Resource Analytics, and Exadata Analytics:

1. From the required application, hover over an entity name from the list.
2. In the pop-up window, click **View More** to display the Entity Card.

**Figure 2-4 Entity Card Details Page**



# 3

## Optimize IT Resources

Use Oracle IT Analytics to view the current trend of IT resource utilization, and use this information to forecast your future resource requirements.

As a business analyst or IT manager, you can use Oracle IT Analytics to:

- Analyze resource consumption by various dimensions, such as platform and lifecycle usage.
- Plan for future growth by forecasting consumption of resources by databases and middleware systems, and identifying hosts that are likely to run out of capacity.
- Compare resource usage across prior periods to identify unexpected growth.

If you see a break in the graphs in the applications, then it could be because the entities weren't using any resources during that time period, and so no data was collected. This behavior is applicable to all entities and graphs.

### Topics:

- [Read This Before You Begin](#)
- [Analyze Database Resources](#)
- [Analyze App Server Resources](#)
- [Analyze Host Resources](#)
- [Analyze Exadata Resources](#)

### Note:

Oracle IT Analytics retains the raw data (the metric data that is collected and uploaded by agents on the target hosts) for eight days. Before purging the raw data, it rolls up the data into hourly and daily intervals. The rolled up data is then retained for 13 months to provide historical metric information. Data becomes inaccessible at the end of the 13-month period.

## Read This Before You Begin

Here are some common terms and basic concepts about the Resource Analytics applications.

**Regression** is a data mining function that predicts a number. It's commonly used in prediction and forecasting.

A regression task begins with a data set in which the target values are known. For example, you can develop a regression model that predicts resource usage based on observed data for your middleware and database resources over a period of time. In the model build, a regression algorithm estimates the value of the target

as a function of the predictors for each case in the build data. These relationships between predictors and target are summarized in a model, which can then be applied to a different data set in which the target values are unknown. **Regression Models** are tested by computing various statistics that measure the difference between the predicted values and the expected values.

You can use a **linear regression** technique if the relationship between the predictors and the target can be approximated with a straight line.

You can use a **nonlinear regression** technique when the relationship between  $x$  and  $y$  can't be approximated with a straight line.

A **channel**, in the context of a forecasting report based on a given regression model, is simply additional information that may be part of the estimated values. For example, some regression algorithms may also yield confidence levels, which indicate the percentage of confidence that the future values will be within the range of estimated values.

The database **storage** subsystem controls the physical data files, logs, and other files used by a database. The database **storage utilization** shows the amount of disk space used by a database or various logical units within the database.

The **database interconnect** refers to the high-speed, low latency dedicated communication link between the nodes of a database cluster. For example, between the nodes of an Oracle Real Application Clusters (Oracle RAC) environment.

## Analyze Database Resources

Understand and analyze trends in the utilization of critical database resources, such as CPU, memory, I/O, storage, and interconnect activity. You can also analyze, compare, and contrast resource usage across databases.

As an IT administrator or a capacity planner, you can use Database Resource Analytics to understand how critical resources, including CPU, memory, interconnect activity, and storage, are used. Capacity planning ensures that your enterprise databases have sufficient resources to meet future business needs.

Database Resource Analytics can be also be used for both types of Autonomous Databases - Shared: Autonomous Data Warehouse – Shared and Autonomous Transaction Processing – Shared. For more information on Autonomous Databases - Shared, see: [Discover Autonomous Databases](#) in *Using Oracle Database Management for Autonomous Databases*.

### Topics:

- [Typical Tasks for Analyzing Database Resources](#)
- [Analyze Database Resource Usage](#)
- [Analyze Database CPU Usage](#)
- [Analyze Database Memory Usage](#)
- [Analyze Database I/O Activity](#)
- [Analyze Database Storage Usage](#)
- [Analyze Database Interconnect Activity](#)

## Typical Tasks for Analyzing Database Resources

As an administrator or capacity planner, perform these tasks using Database Resource Analytics to analyze how your database resources are being used.

Task	Description	More Information
Analyze database CPU usage.	View the current aggregate CPU utilization trend, compare CPU utilization across databases, and forecast future CPU needs to plan your resources for a later time period.	<a href="#">Analyze Database CPU Usage</a>
Analyze database memory usage.	View the current memory utilization trend, compare the memory utilization across databases, and forecast future memory resource needs to plan your resources for a later time period.	<a href="#">Analyze Database Memory Usage</a>
Analyze database I/O activity.	View the current I/O activity, compare I/O activity across databases, and forecast future I/O needs to plan your resources for a later time period.	<a href="#">Analyze Database I/O Activity</a>
Analyze database storage usage.	View the current space usage trend, compare storage utilization across databases, and forecast future storage resources for a later time period.	<a href="#">Analyze Database Storage Usage</a>
Analyze database interconnect activity.	View the current interconnect activity trend, compare interconnect activity across databases, and forecast future interconnect resources for a later time period.	<a href="#">Analyze Database Interconnect Activity</a>

## Analyze Database Resource Usage

Using Database Resource Analytics, you can view your current database resource utilization by CPU, memory, I/O, and storage. You can also view the interconnect activity details for a selected database. You can also compare and contrast resource usage between databases, and identify which databases use the most resources, and are expanding rapidly.

When you log in to the application, the Overview page is displayed. The **Inventory** section displays the total number of databases along with the database types. The pie chart also depicts this information. When you check the **Show Oracle Autonomous Databases** option, the pie chart displays the breakup of the Oracle databases into Autonomous Data Warehouse (ADW) databases and Autonomous Transaction Processing (ATP) databases. In addition, the CPU and Memory usage charts will display a separate line to show the resource consumption by these database targets.

The **Resource Utilization** section shows the current utilization of databases and the databases that are reaching the server capacity headroom.

- **Current Utilization of Database Systems**

This section shows the number of database systems running with 0–25%, 25–50%, 50–75%, and 75–100% utilization of resources. This categorization is based on CPU and memory utilization. The current utilization of some database systems cannot be calculated if the available data is insufficient. The resource utilization for such database systems is categorized as 'Unknown'.

The bucket with 0–25% resource utilization is considered as the low utilization range, and the bucket with 75–100% utilization of resources is considered as the high utilization range.

You can drill down to view the CPU and memory breakdown by individual databases and for an aggregated trend and forecast view.

- **Database Systems Projected to Reach High Utilization**

This section shows the number of database systems that are forecasted to reach high resource utilization threshold, in the next 30, 90, 180, and 180+ days. This forecast cannot be calculated for some database systems if the available data is insufficient. The forecast for such database systems is categorized as 'Unknown'. The database systems fall into different buckets depending on the utilization threshold values that you set in the **IT Analytics Settings** page. See [Configure Capacity Thresholds for Your Environment](#).

You can analyze the CPU and memory resources that are capacity bottleneck by using the forecasted data. Use this information to identify the databases running on systems that are under-provisioned, resulting in performance degradation.

 **Note:**

You can also view the forecast of resource utilization of database systems for 365+ days. You can add this time range from the **IT Analytics Settings** page. See [Configure Capacity Thresholds for Your Environment](#).

You can drill down to check for the databases that are forecasted to run out of capacity in future.

The CPU, Memory, Allocated Storage, and I/O utilization charts show the overall usage of these resources. To view detailed information about each of your critical database resources, click the corresponding tile in the navigation pane.

You can view the following information for each resource (CPU, memory, storage, I/O):

- **Aggregate resource utilization across all databases:** View the summary of resource utilization across all databases, which helps identify points in time where there may have been unexpected spikes in resource utilization.
- **Resource breakdown by individual databases:** View the resource utilization for each database, which helps compare resource utilization between databases, and identify servers with underused or overused resources. The size of the cells shows the volume of resources used by the specific database. The color of the cells show the rate of change, where red implies that the database is growing rapidly, green

implies that the database is shrinking, and grey implies no change in the resources used by the database.

- **Trend and Forecast:** See the forecasted utilization for the particular resource, along with the current utilization trend. Use this information for capacity planning.

## Analyze Database CPU Usage

You can analyze your CPU utilization, using the Database Resource Analytics application.

### Topics:

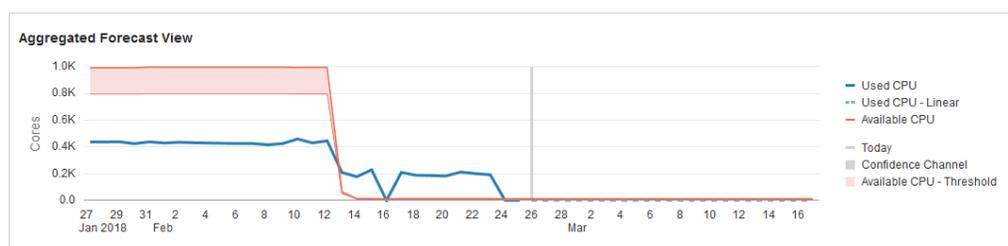
- [Analyze Available CPU Resources](#)
- [Compare CPU Utilization Between Databases](#)
- [Analyze Current CPU Trend and Forecast Utilization](#)

## Analyze Available CPU Resources

The **CPU** section under the **Overview** tab displays the available CPU capacity and the current utilization.

Alternatively, click the **CPU** tab in the left pane. The first chart (**Aggregate Forecast View**) shows you the actual CPU utilization (**Used CPU**), the available capacity (**Available CPU**), and the availability threshold (**Available CPU - Threshold**).

**Figure 3-1 Aggregate Forecast View for CPU Usage**



The difference between the values for **Used CPU** and **Available CPU** is your unused capacity. The smaller the difference between your actual resource usage and your resource capacity, the more urgent the need to expand your available resource capacity.

## Compare CPU Utilization Between Databases

The **CPU breakdown by individual databases** chart lets you compare how your different, individual databases are using their resources.

1. Click the **Group By** list to select the metrics by which you want to group the databases.
  - Life Cycle Status
  - Target Version

- Cost Center
  - Host Name
  - Contact
  - Database Type
2. To view the 10 databases consuming the most CPU resources, from the **Group by** list, select **Top 10**.
  3. To view the CPU consumption split on the basis of the pluggable databases, select **Split by Pluggable Databases (Top 10)**.
  4. Select the databases for which you want to view specific information.
  5. To view the list of databases, and their relevant information like the DB type, CPU used, change %, hostname, target version, life cycle, and cost center in a CSV format, click **Download CSV**.

The number of database processes run on the CPU, and the percentage change are represented visually, through the size and color of the cells, respectively. Cells that are larger in size run more database processes than smaller cells. The largest cell would be that of the databases using the most CPU resources. The color of the cells is determined by the percentage change in the CPU resources used by the databases.

## Analyze Current CPU Trend and Forecast Utilization

From the **Aggregated Forecast View** chart, you can view the current trend of CPU utilization, and use this trend to forecast your CPU utilization.

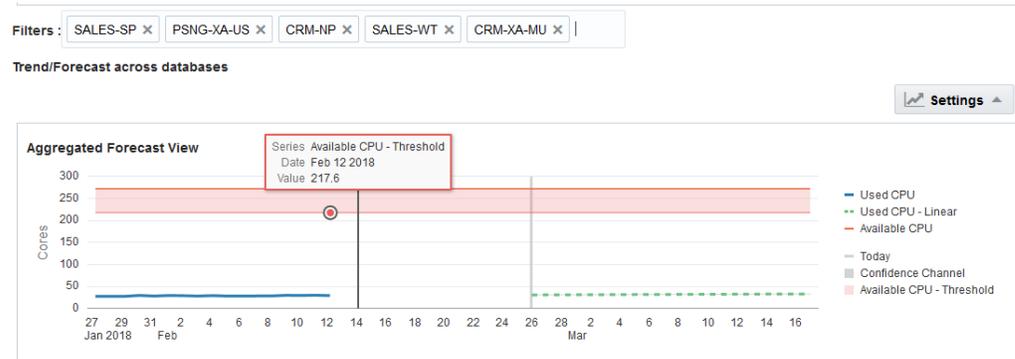
1. Under the section **Trend/Forecast across databases**, in the **Filters** field, enter the name of the database. You can specify up to five databases in the **Filters** field.  
  
Alternatively, in the **CPU Breakdown by Individual Databases** section, use **Ctrl + Shift** to select multiple databases. The data for the selected databases loads in the **Trend and Forecast** section.
2. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.
3. Click **Run**.

The information loads in the **Trend/Forecast across databases** section.

For the Database Resource Analytics application to forecast your resource utilization for a year, there must be at least 13 months of data stored in the IT Analytics warehouse. The forecast resource utilization is more accurate when there's more data stored in the IT Analytics warehouse.

You can plan your database capacity based on the current and forecast resource utilization.

**Figure 3-2 Trend and Forecast Chart for CPU Utilization**



In the image, the forecast CPU utilization increases slowly over time. This means that at some point in the future, the CPU capacity will be entirely used. You should plan to expand the CPU resource capacity before that point, to ensure that your databases have enough resources to function smoothly.

## Analyze Database Memory Usage

Use the Database Resource Analytics application to analyze how your databases use available memory resources.

 **Note:**

Analyze Database Memory Usage currently cannot be used with Autonomous Databases – Shared.

### Topics:

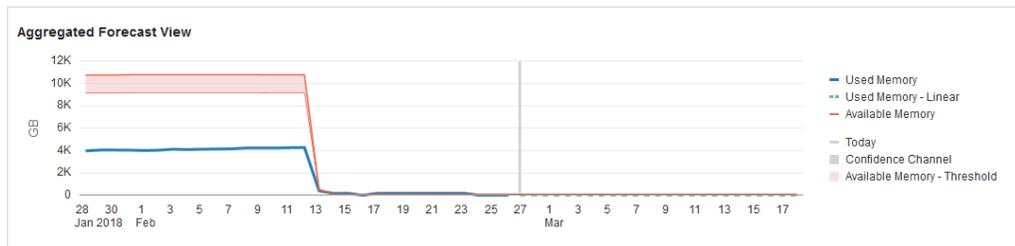
- [Analyze Available Memory Resources](#)
- [Compare Memory Utilization Between Databases](#)
- [Analyze Current Memory Usage Trend and Forecast Utilization](#)

## Analyze Available Memory Resources

The **Memory** section under the **Overview** tab displays the available memory capacity and the current utilization.

Alternatively, click the **Memory** tab in the left pane. The first chart (Aggregated Forecast View) shows you the database memory usage and the available memory resources.

**Figure 3-3 Aggregate Forecast Chart for Memory Usage**



The difference between the values for **Used Memory** and **Available Memory** is the unused capacity. The smaller the difference between your actual resource usage and your resource capacity, the more urgent the need to expand your available resource capacity.

## Compare Memory Utilization Between Databases

The **Memory breakdown by individual databases** chart lets you compare how your different, individual databases are using their resources.

1. Click the **Group by** list to select what metrics to group the databases by.
  - Life Cycle Status
  - Target Version
  - Cost Center
  - Host Name
  - Contact
  - Database Type
2. To view the 10 databases consuming the most memory resources, from the **Group By** list, select **Top 10**.
3. Select the databases for which you want to view specific information.
4. To view the list of databases, and their relevant information like the DB type, memory used, change %, hostname, target version, life cycle, and cost center in a CSV format, click **Download CSV**.

The amount of memory resources and the percentage change in the usage of memory resources are represented visually, through the size and color of the cells, respectively. The cells that are larger in size use more memory resources than the smaller cells. The largest cell would be that of the database using the most memory resources. The color of the cells is determined by the percentage change in the memory resources used by the databases.

## Analyze Current Memory Usage Trend and Forecast Utilization

From the **Aggregated Forecast View** chart, you can view the current trend of Memory usage, and use this trend to forecast your memory usage.

1. In the section **Trend/Forecast across databases**, in the **Filters** field, enter the name of the database. You can specify up to five databases in the Filters field.

Alternatively, in the **Memory breakdown by individual databases** section, use **Ctrl + Shift** to select multiple databases.

The data for the selected databases loads in the **Trend/Forecast across databases** section.

2. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.
3. Click **Run**.

The information loads in the **Trend/Forecast across databases** section.

For the Database Resource Analytics application to forecast your resource utilization for a year, there must be at least 13 months of data stored in the IT Analytics warehouse. The forecast resource utilization is more accurate when there's more data stored in the IT Analytics warehouse.

You can plan your database capacity based on the current and forecast resource utilization.

## Analyze Database I/O Activity

Analyze Input/Output (I/O) utilization across your databases, using the Database Resource Analytics application.

### Topics:

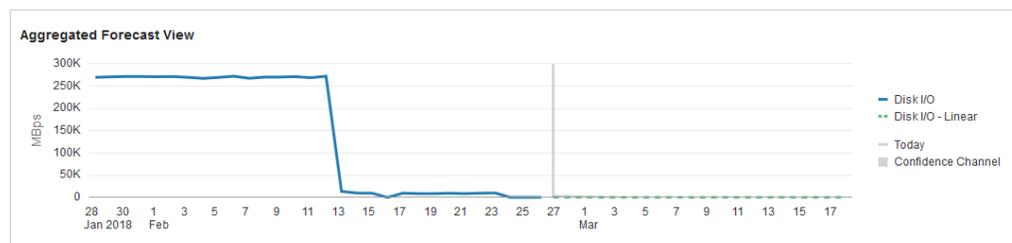
- [Analyze I/O Trend Over Time](#)
- [Compare I/O Activity Between Databases](#)
- [Analyze Current I/O Trend and Forecast Activity](#)

## Analyze I/O Trend Over Time

The **I/O** section under the **Overview** tab displays the current trend of disk I/O activity across the selected databases for the specified time period.

Alternatively, click the **I/O** tab in the navigation pane to view the current trend of I/O activity in terms of MB units, during a specified time and forecast future activity.

**Figure 3-4 Aggregate Forecast View for IO**



## Compare I/O Activity Between Databases

The **IO breakdown by individual databases** chart lets you compare the I/O activity of your different databases.

1. Click the **Group By** list to select what metrics to group the databases by.
  - Life Cycle Status
  - Target Version
  - Cost Center
  - Host Name
  - Contact
  - Database Type
2. To view the 10 databases with the highest I/O activity, from the **Group By** list, select **Top 10**.
3. To view the disk I/O activity split on the basis of the pluggable databases, select **Split by Pluggable Databases (Top 10)**.
4. Select the databases for which you want to view specific information.
5. To view the list of databases, and their relevant information like the DB type, disk I/O, change %, hostname, target version, life cycle, and cost center in a CSV format, click **Download CSV**.

The volume and percentage changes in the I/O activity are represented visually through the size and color of the cells, respectively. Cells that are larger have greater I/O activity than smaller cells. The largest cell would be that of the database with the highest I/O activity. The color of the cells is determined by the percentage change in the I/O activity during the specified time period.

## Analyze Current I/O Trend and Forecast Activity

From the **Aggregated Forecast View** chart, you can view the current trend of I/O activity, and use this trend to forecast future I/O activity.

1. In the section **Trend/Forecast across databases**, in the **Filters** field, enter the name of the database. You can enter up to five databases in the **Filters** field. Alternatively, in the **IO breakdown by individual databases** section, use **Ctrl + Shift** to select multiple databases.

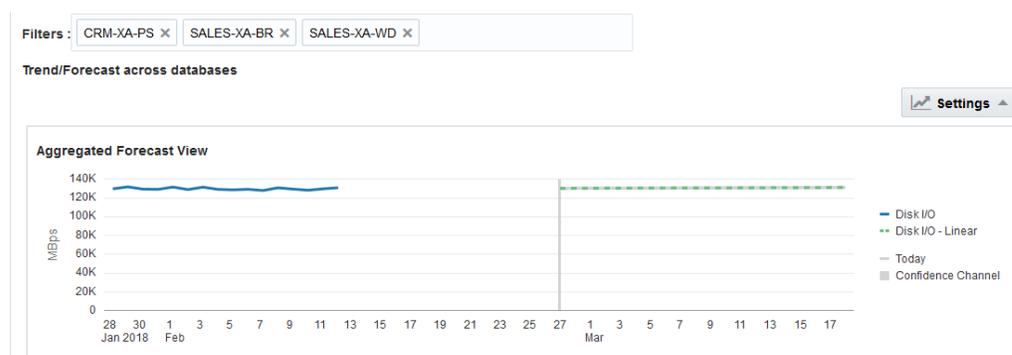
The data for the selected databases loads in the **Trend/Forecast across databases** section.

2. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.
3. Click **Run**.

The information loads in the **Trend and Forecast** section.

For the Database Resource Analytics application to forecast your resource utilization for a year, there must be at least 13 months of data stored in the Oracle IT Analytics warehouse. The forecast value is more accurate when there's more data stored in the Oracle IT Analytics warehouse.

You can use the current and forecast I/O activity to plan your database capacity.

**Figure 3-5 I/O Trend and Forecast Chart**

For example, if the forecast I/O activity is at a high level, then you must ensure that you have sufficient resources to support the elevated I/O activity.

## Analyze Database Storage Usage

You can analyze your storage utilization across databases, using the Database Resource Analytics application.

Oracle databases are made up of logical storage units called tablespaces. Using Data Explorer (see [Search, Explore, and Analyze Data Using Data Explorer](#)), you can view information about tablespace usage and allocation for your databases monitored by Oracle Management Cloud and Oracle Enterprise Manager.

For a list of related metrics, see [Storage Usage by Tablespace](#).

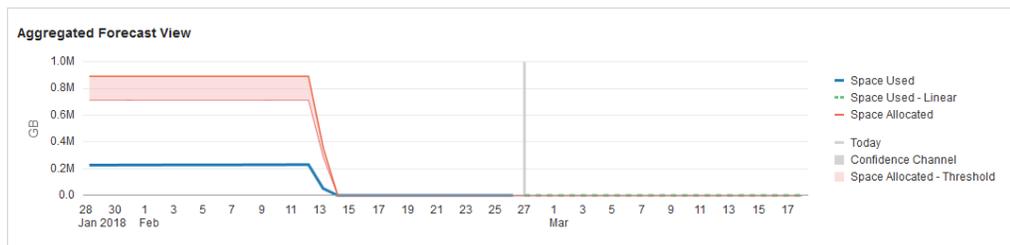
### Topics:

- [Analyze Available Storage Resources](#)
- [Compare Storage Utilization Between Databases](#)
- [Analyze Current Storage Trend and Forecast Utilization](#)
- [View Tablespace Breakdown by Database](#)

## Analyze Available Storage Resources

The **Allocated Storage** section under the **Overview** tab displays the current tablespace usage.

Alternatively, click the **Storage** tab in the navigation pane to view the aggregate storage utilization.

**Figure 3-6 Aggregated Forecast View**

From this chart, you can view the trend of space utilization.

## Compare Storage Utilization Between Databases

The **Storage breakdown by individual databases** chart lets you compare how your different databases use your available storage.

1. Click the **Group By** list to select what metrics to group the databases by.
  - Life Cycle Status
  - Target Version
  - Cost Center
  - Contact
  - Database Type
2. To view the 10 databases with the most storage resources, from the **Group By** list, select **Top 10**.
3. To view the storage resources split on the basis of pluggable databases, select **Split by Pluggable Databases (Top 10)**.
4. Select the databases for which you want to view specific information. Use **Ctrl + Shift** to select multiple databases.
5. To view the list of databases, and their relevant information such as the DB type, space used, change %, hostname, target version, life cycle, and cost center in a CSV format, click **Download CSV**.

The volume of storage resources used and the percentage change are represented visually, through the size and color of the cells, respectively. Cells that are larger in size use more storage resources than smaller cells. The largest cell would be that of the databases using the most storage resources. The color of the cells is determined by the percentage change in the storage resources used by the databases.

## Analyze Current Storage Trend and Forecast Utilization

From the **Trend/Forecast across databases** chart, you can view the current trend of storage utilization, and use this trend to forecast future resource utilization.

1. In the **Trend/Forecast across databases** section, in the **Filters** field, enter the name of the database. You can enter up to five databases in the Filters field.

Alternatively, you can select the database in the **Storage breakdown by individual databases** section.

The data for the selected database loads in the **Trend/Forecast across databases** section.

2. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.
3. Click **Run**.

The information loads in the **Trend/Forecast across databases** section.

For the application to forecast your resource utilization for a year, there must be at least 13 months of data stored in the Oracle IT Analytics warehouse. The forecast value is more accurate when there's more data stored in the Oracle IT Analytics warehouse.

You can use the current and forecast storage utilization to plan your database capacity.

**Figure 3-7 Trend/Forecast Across Databases Chart**



For example, if according to the forecast storage space utilization, the storage requirements will gradually increase, then you need to expand your storage capacity to allow for this forecast increase in storage needs.

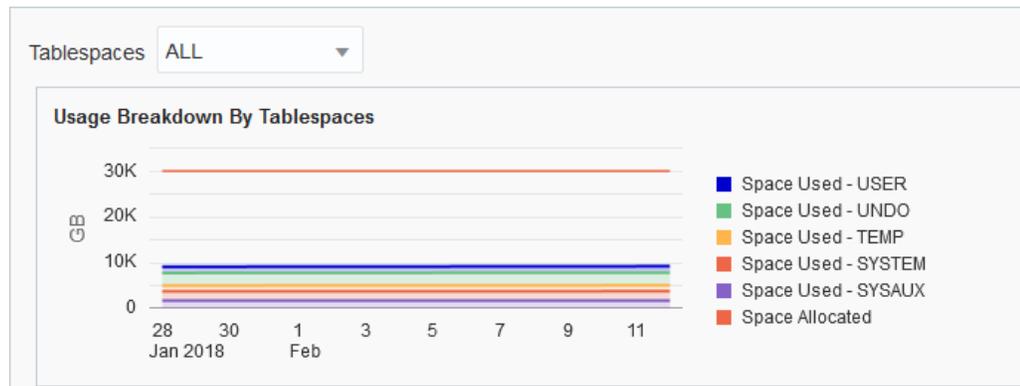
## View Tablespace Breakdown by Database

In addition to the consolidated trend and forecast for the selected databases, you can view:

- Space allocation breakdown by individual databases
- Trend and forecast per database
- Usage of each tablespace for an individual database

Click **more...** to view the usage breakdown for that database, by tablespace.

**Figure 3-8 Usage Breakdown by Tablespaces Chart**



## Analyze Database Interconnect Activity

You can analyze interconnect activity across databases, using the Database Resource Analytics application.



### Note:

Analyze Database Interconnect Activity cannot currently be used with Autonomous Databases – Shared.

### Topics:

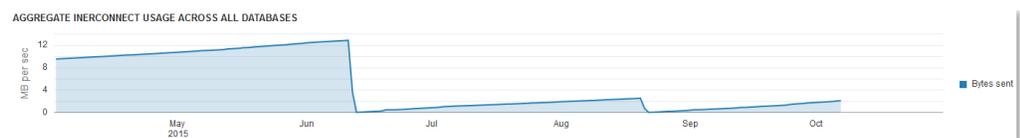
- [Analyze Interconnect Trend Over Time](#)
- [Compare Interconnect Activity Between Databases](#)
- [Analyze Current Interconnect Trend and Forecast Activity](#)

## Analyze Interconnect Trend Over Time

The **Interconnect** section under the **Overview** tab displays the current trend of interconnect activity in terms of bytes sent, across the selected databases for the specific time period.

Alternatively, click the **Interconnect** tile in the navigation pane to view the current trend of interconnect activity and forecast future activity.

**Figure 3-9 Aggregate Interconnect Across All Databases Chart**



## Compare Interconnect Activity Between Databases

The **Interconnect Breakdown by Individual Databases** chart lets you compare the interconnect activity of your different databases.

1. Click the **Group by** list to select what metrics to group the databases by.
  - Life Cycle Status
  - Target Version
  - Cost Center
  - Host Name
  - Contact
2. To view the 10 databases with the highest interconnect activity, from the **Group by** list, select **Top 10 databases**.
3. Select the databases for which you want to view specific information.

The volume and percentage changes of the interconnect activity are represented by the size and color of the cells, respectively. Cells that are larger have greater interconnect activity than smaller cells. The largest cell would be that of the database with the highest interconnect activity. The color of the cells is determined by the percentage change in the interconnect activity during the specified time period.

## Analyze Current Interconnect Trend and Forecast Activity

From the **Trend and Forecast** chart, you can view the current trend of interconnect activity, and use this trend to forecast future interconnect activity.

1. From the section **Trend and Forecast**, in the **Filters** field, enter the name of the database. You can enter multiple databases in the **Filters** field. Alternatively, in the **Interconnect Breakdown by Individual Databases** section, use **Ctrl + Shift** to select multiple databases.

The data for the selected databases loads in the **Trend and Forecast** section.

2. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.
3. Click **Run**.

For the Database Resource Analytics application to forecast your resource utilization for a year, there must be at least 13 months of data stored in the Oracle IT Analytics warehouse. The forecast value is more accurate when there's more data stored in the Oracle IT Analytics warehouse.

You can use the forecast and current interconnect activity for capacity planning.

**Figure 3-10 Trend and Forecast Chart for Interconnect Activity**



## Analyze App Server Resources

As an IT administrator or a capacity planner, you can use App Server Resource Analytics to analyze resources used by the app server deployments in your data center, and to understand resource constraints and requirements for their deployments.

Using the App Server Resource Analytics application, you can:

- Analyze the distribution of a assets across the enterprise based on resource usage.
- Assess aggregate resource usage versus configured capacity.
- Compare resource usage across multiple application server instances.
- Analyze historical trends in resource usage, and forecast future usage by using a variety of trending algorithms, such as linear regression.

### Topics:

- [Typical Tasks for Analyzing App Server Resources](#)
- [Analyze App Server Resource Usage](#)
- [Analyze Heap Usage](#)
- [Analyze Memory Usage](#)
- [Analyze CPU Usage](#)

## Typical Tasks for Analyzing App Server Resources

Here's the list of tasks to analyze app server resources in your data center.

Task	Description	More Information
Analyze app server resource usage across multiple app server instances.	Analyze how heap, memory, and CPU usage are distributed across multiple app server instances in your IT environment, and drill down to view more detailed information about a particular set of app server instances.	<a href="#">Analyze App Server Resource Usage</a>
Analyze and forecast heap usage.	View aggregate heap usage across selected app server instances, heap usage by servers, and trending and forecasting.	<a href="#">Analyze Heap Usage</a>
Analyze and forecast memory usage.	View aggregate memory usage across selected app server instances, memory usage by servers, and trending and forecasting.	<a href="#">Analyze Memory Usage</a>

Task	Description	More Information
Analyze and forecast CPU usage.	View aggregate CPU usage across selected app server instances, CPU usage by servers, and trending and forecasting.	<a href="#">Analyze CPU Usage</a>

## Analyze App Server Resource Usage

As an IT administrator, you can analyze the app server resource usage across multiple deployments, during a specified time period.

### Topics:

- [Analyze Resource Usage Across Application Server Instances](#)
- [Analyze Resource Usage by Date Ranges](#)
- [Analyze Resource Usage by Specific Time Periods](#)

## Analyze Resource Usage Across Application Server Instances

1. Click the **Overview** tab in the left pane.
2. Select an application server or an entity from the drop-down list and choose a time period to narrow down your analysis. You can also specify a custom time period.

The memory utilization is displayed under three categories:

- **Heap Usage:** The histogram shows the distribution of the instances for the selected app servers or entities in terms of their heap usage. The height of the bars indicates the number of app server deployments in the respective category for the selected time period.
- **Host Memory Usage:** The histogram shows the categorization of hosts for app server instances in terms of their memory usage. The height of the bars indicates the number of the app server hosts in the respective category for the selected time period.
- **Application Server CPU Usage:** The histogram shows the categorization of app server instances in terms of their CPU usage. The height of the bars indicates the number of app server instances in the respective category for the selected time period.

Click a bar or a group name to drill down to the respective usage page and view the usage details, and to forecast the usage for that group of app server instances. When you drill down from the histogram, a page level filter is set. You can clear this filter from the **Filters** option.

## Analyze Resource Usage by Date Ranges

Use the date-time list to select a time window for your analysis. The specified date range applies to all reports in App Server Resource Analytics.

You can use one of the following options as the time period for the analysis:

- Last 24 hours

- Last 7 days
- Last 14 days
- Last 30 days
- Last 90 days
- Last 12 months
- Custom: Select a start date and end date for the analysis period.

By default, middleware resource utilization is displayed for the last 30 days.

1. From the App Server Resource Analytics page, click the list next to **Last 30 days**.
2. From the list of options, select a predefined time period. Alternatively, click **Custom** to define a time period of your choice.

The selected date range applies to all reports in the application.

## Analyze Resource Usage by Specific Time Periods

You can use the Time Filter Settings option to specify time intervals for analysis and then analyze that usage.

For example, you can analyze usage during peak hours of the week or during a holiday season. Available options include:

- **Time:** Enter a time range (in terms of hours of the day) for your analysis. Permitted values follow a 24-hour format. For example, to analyze resource usage between 8 a.m. and 5 p.m., enter 8–17 in the space provided (note the use of 24-hour time, also referred to as military time; 5 p.m. is shown as 17 here).
- **Days:** Select the days of the week during which you want to analyze resource utilization. For example, select the weekday boxes to analyze resource usage from Monday to Friday.
- **Months:** Select specific months for your analysis. For example, if you want to analyze resource usage during the period between Thanksgiving and New Year's, then select the boxes for November and December.

To use the Time Filter Settings:

1. From the App Server Resource Analytics page, select **Custom** from the time period drop-down list, and then click the **Time Filter Settings** icon.
2. In the **Time** field, enter the business hours. For example, 8–17.
3. To specify days, select the corresponding options under **Days**.
4. To specify months for the analysis, select the corresponding options under **Months**.
5. Click **Apply**.

Data for the specified time period loads. The time period selection applies to all reports within the application.

## Analyze Heap Usage

The Java heap is a repository for active objects, dead objects, and free memory. The heap size determines the volume of objects stored in the repository. When the volume of the objects approaches the maximum heap size, the garbage collection happens.

If the heap size is large, then the full garbage collection is slower and happens less frequently.

**Topics:**

- [Monitor Total Heap Usage Across Servers](#)
- [Monitor Heap Usage By Server](#)
- [Monitor the Current Trend and Forecast Heap Usage](#)

## Monitor Total Heap Usage Across Servers

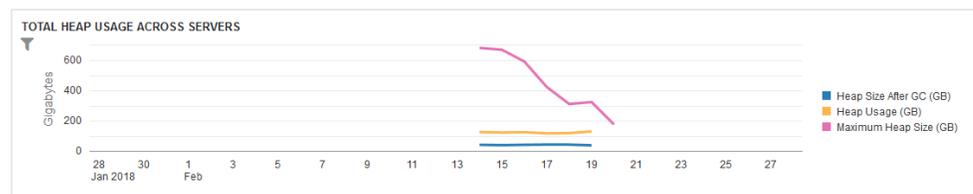
To monitor total heap usage across servers:

1. Click the **Heap** tab in the left pane.

The first chart (**Total Heap Usage Across Servers**) shows you the following metrics:

- **Heap Size After Garbage Collection:** This is the amount of Java heap memory that's currently being used by the Java Virtual Machine (JVM) after garbage collection has been completed.
- **Heap Usage:** This is the total amount of Java heap memory that's currently being used by the JVM.
- **Maximum Heap Size:** This is the amount of memory currently allocated to the JVM heap.

**Figure 3-11 Total Heap Usage Across All Server Chart**



2. From this chart:

- Monitor the heap usage with relation to the maximum heap size to identify whether the maximum heap size needs to be altered to accommodate an increase in heap usage.
- View the difference between the heap usage and the heap size after garbage collection to monitor the amount of garbage stored in the JVMs.

## Monitor Heap Usage By Server

1. Click the **Heap** tab in the left pane.

The second chart (**Heap Usage by Server**) lets you compare the heap usage between different app server instances. Each instance is represented by a cell. The size of each cell is determined by the heap size after garbage collection, while the color of each cell is determined by the percentage change in the heap usage. Shades of red indicate a growth in heap usage while shades of green indicate that the heap size has shrunk. Grey cells indicate no change in heap usage.

2. Click the **Group By** list to select the metrics to group the app server instances.
  - Domain
  - Target Version
  - Target Type
  - JDK Version
  - Host Name
  - Life Cycle Status
  - Cost Center
  - Contact
  - None
3. To view the top 10 app server instances with the most heap usage, from the **Group By** list, select **Top 10 Servers**.
4. Select the app server instances for which you want to view specific information by clicking the corresponding cell. If you've decided to group the app server instances, then you can select a group of instances by clicking a specific block.

Data about the current trend of heap usage and the forecasted heap usage for the selected app server instances loads in the **Trend and Forecast** chart.

## Monitor the Current Trend and Forecast Heap Usage

To monitor the current trend and forecast heap usage:

1. Click the **Heap** tab in the left pane.
2. From the **Heap Usage by Server** chart, select the app server instances or a group of app server instances for which you want to view the heap usage trend and forecast.

In the **Filters** field in the **Trend and Forecast** chart, enter the name of the app server instance.

3. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.

You can use both linear and nonlinear regression models when you're forecasting resource usage. The linear regression model is the default option, which is used when the relationship between the predictors and the target value can be approximated with a straight line in the graph.

When you're predicting resource usage, try linear regression first. If the default linear regression model can't model the specific curve that exists in your data, then you can't get an appropriate fit with linear regression. In such scenarios, try the other regression models to see if they produce a trend that fits the curve of your actual data better:

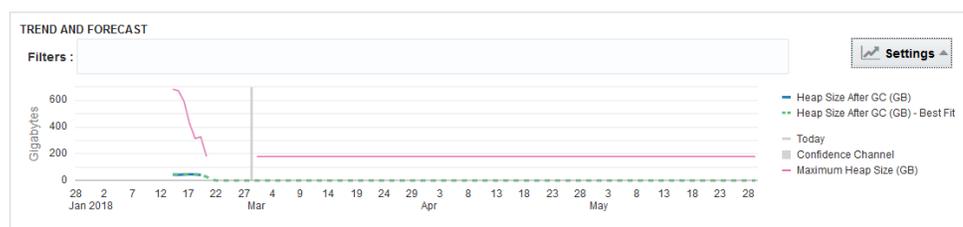
- **BestFit:** A line shown in the graph with the same number of data points on each side of the line. The line is in the median position, but it doesn't go from the first data point to the last data point.
- **Exponent:** A base and an exponent. For example,  $2^3$ . In this example, 2 is the base and 3 is the exponent.

- **NlogN**: A logarithm of a number that's the exponent to which the value of the base is increased to produce that number. It's also referred to as repeated multiplication.
- **Quadratic**: A curve created by the intersection of a cone and a plane parallel to the side of the cone.

From the **Trend and Forecast** chart, you can view the following metrics:

- Heap size after garbage collection
- Heap size after garbage collection according to the selected regression model
- Maximum heap size

**Figure 3-12 Trend and Forecast Chart for Heap Usage**



From this chart you can identify whether you need to expand resources to accommodate future heap usage.

## Analyze Memory Usage

Setting the maximum heap size according to memory needs ensures that garbage collection is faster and more frequent. Use App Server Resource Analytics to identify how your heap usage compares with the available memory.

### Topics:

- [Monitor Total Memory Usage Across Hosts](#)
- [Monitor Memory Usage by Host](#)
- [Monitor the Current Trend and Forecast Memory Usage](#)

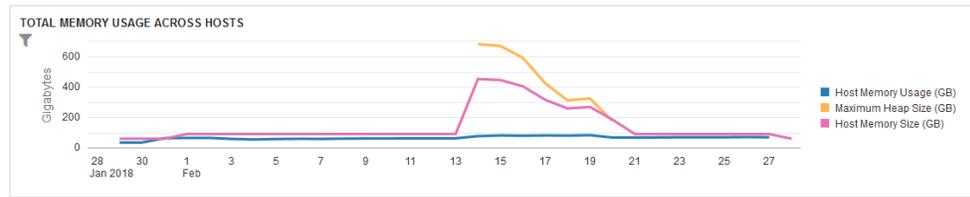
## Monitor Total Memory Usage Across Hosts

To monitor total memory usage across hosts:

1. Click the **Memory** tab in the left pane.

The first chart (**Total Memory Usage Across Hosts**) shows you the following metrics:

- **Host Memory Usage**: This is the amount of memory that's currently being used by the app server instances.
- **Maximum Heap Size**: This is the configured maximum heap size for your app server instances.
- **Host Memory Size**: This is the total amount of memory available for your app server instances.

**Figure 3-13 Host Memory Usage Chart**

2. From this chart:
  - Monitor the heap size and memory usage with relation to the available memory.
  - View the difference between the Host Memory Usage and the Host Memory Size to identify whether you should expand your memory resources.

## Monitor Memory Usage by Host

To monitor memory usage by host:

1. Click the **Memory** tab in the left pane.
 

The second chart (**Memory Usage by Host**) lets you compare the memory usage between different app server instances. Each instance is represented by a cell. The size of each cell is determined by the host memory usage, while the color of each cell is determined by the percentage change in the host memory usage. Shades of red indicate a growth in memory usage while shades of green indicate that the memory usage has shrunk. Grey cells indicate no change in memory usage.
2. Click the **Group By** list to select the metrics to group the app server instances.
  - Life Cycle Status
  - Cost Center
  - Contact
  - Domain
  - Host Version
  - Operating System
  - None
3. Select the app server instances for which you want to view specific information by clicking the corresponding cell. If you've decided to group the app server instances, then you can select a group of instances by clicking a specific block.

Data about the current trend of heap usage and the forecasted memory usage for the selected app server instances loads in the **Trend and Forecast** chart.

## Monitor the Current Trend and Forecast Memory Usage

To monitor the current trend and forecast memory usage:

1. Click the **Memory** tab in the left pane.
2. From the **Memory Usage by Host** chart, select the app server instances or a group of app server instances for which you want to view the memory usage trend and forecast.

In the **Filters** field of the **Trend and Forecast** section, specify the name of the app server instance.

- To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.

You can use both linear and nonlinear regression models when you're forecasting resource usage. The linear regression model is the default option, which is used when the relationship between the predictors and the target value can be approximated with a straight line in the graph.

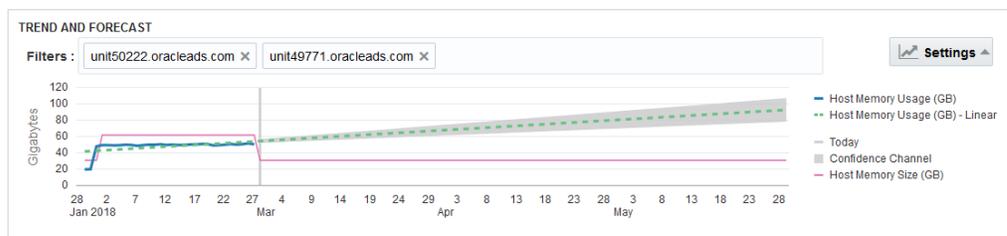
When you're predicting resource usage, try linear regression first. If the default linear regression model can't model the specific curve that exists in your data, then you can't get an appropriate fit with linear regression. In such scenarios, try the other regression models to see if they produce a trend that fits the curve of your actual data better:

- BestFit:** A line shown in the graph with the same number of data points on each side of the line. The line is in the median position, but it doesn't go from the first data point to the last data point.
- Exponent:** A base and an exponent. For example,  $2^3$ . In this example, 2 is the base and 3 is the exponent.
- NlogN:** A logarithm of a number that's the exponent to which the value of the base is increased to produce that number. It's also referred to as repeated multiplication.
- Quadratic:** A curve created by the intersection of a cone and a plane parallel to the side of the cone.

From the **Trend and Forecast** chart, you can view the following metrics:

- Host Memory Usage
- Host Memory Usage according to the selected regression model
- Host Memory Size

**Figure 3-14** Trend and Forecast Chart for Memory Usage



From this chart you can identify whether you need to expand resources to accommodate future memory usage.

## Analyze CPU Usage

High CPU utilization by your app server instances could result in a high response time. As an IT administrator, you can use App Server Resource Analytics to identify periods of high CPU utilization to let you conduct further analysis.

### Topics:

- [Monitor Average CPU Usage Across Servers](#)
- [Monitor CPU Usage By Server](#)
- [Monitor the Current Trend and Forecast CPU Usage](#)

## Monitor Average CPU Usage Across Servers

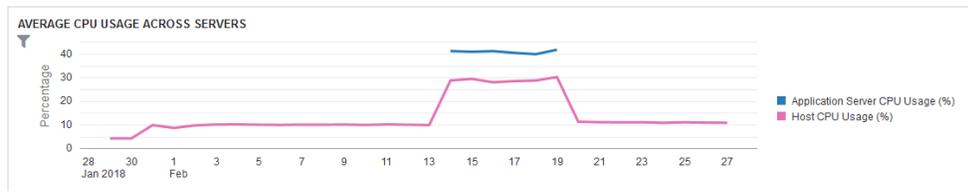
To monitor average CPU usage across servers:

1. Click the **CPU** tab in the left pane.

The first chart (**Average CPU Usage Across Servers**) shows you the following metrics:

- Application Server CPU Usage
- Host CPU Usage

**Figure 3-15 Average CPU Usage Across Servers Chart**



2. From this chart, monitor how your app server instances are using your CPU resources by studying the difference between app server CPU Usage and Host CPU Usage.

## Monitor CPU Usage By Server

To monitor CPU usage by server:

1. Click the **CPU** tab in the left pane.

The second chart (**CPU Usage by Server**) lets you compare the CPU usage between different app server instances. Each instance is represented by a cell. The size of each cell is determined by the app server CPU usage, while the color of each cell is determined by the percentage change in CPU usage by each app server instance. Shades of red indicate a growth in CPU usage, while shades of green indicate that the CPU usage has shrunk. Grey cells indicate no change in CPU usage.

2. Click the **Group By** list to select the metrics to group the app server instances.
  - Domain

- Target Version
  - Target Type
  - JDK Version
  - Host Name
  - Life Cycle Status
  - Cost Center
  - Contact
  - None
3. Select the app server instances for which you want to view specific information by clicking the corresponding cell. If you've decided to group the app server instances, then you can select a group of instances by clicking a specific block.

Data about the current trend of CPU usage and the forecasted CPU usage for the selected app server instances loads in the **Trend and Forecast** chart.

## Monitor the Current Trend and Forecast CPU Usage

To monitor the current trend and forecast CPU usage:

1. Click the **CPU** tab in the left pane.
2. From the **CPU Usage by Server** chart, select the app server instances or a group of app server instances for which you want to view the CPU usage trend and forecast.

Alternatively, in the **Filters** field in the **Trend and Forecast** chart, you can enter the name of the app server instances.

3. To set the Regression Model, Forecast Period, and Confidence (%), click **Settings**.

You can use both linear and nonlinear regression models when you're forecasting resource usage. The linear regression model is the default option, which is used when the relationship between the predictors and the target value can be approximated with a straight line in the graph.

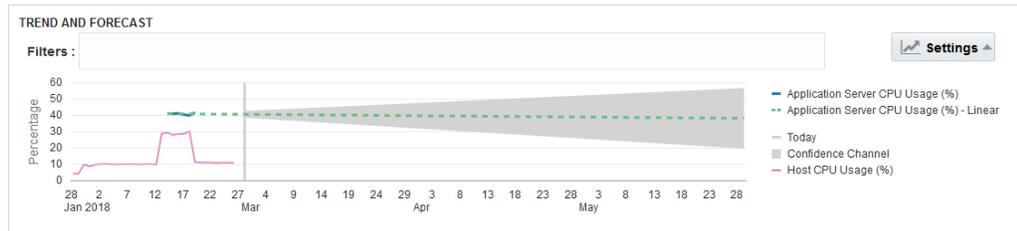
When you're predicting resource usage, try linear regression first. If the default linear regression model can't model the specific curve that exists in your data, then you can't get an appropriate fit with linear regression. In such scenarios, try the other regression models to see if they produce a trend that fits the curve of your actual data better:

- **BestFit:** A line shown in the graph with the same number of data points on each side of the line. The line is in the median position, but it doesn't go from the first data point to the last data point.
- **Exponent:** A base and an exponent. For example,  $2^3$ . In this example, 2 is the base and 3 is the exponent.
- **NlogN:** A logarithm of a number that's the exponent to which the value of the base is increased to produce that number. It's also referred to as repeated multiplication.
- **Quadratic:** A curve created by the intersection of a cone and a plane parallel to the side of the cone.

From the **Trend and Forecast** chart you can view the following metrics:

- Application Server CPU Usage
- Application Server CPU Usage according to the selected regression model
- Host CPU Usage

**Figure 3-16 Trend and Forecast Chart for CPU Usage**



From this chart you can identify whether you need to expand resources to accommodate future CPU usage.

## Analyze Host Resources

Host Resource Analytics lets you view and analyze utilization trends for critical server resources such as CPU, memory, and storage. You can also analyze, compare, and contrast resource usage across hosts.

Through Host Resource Analytics, Oracle IT Analytics lets you:

- Group host specifications to show a breakdown of resource utilization by operating system.
- Identify a server's capacity and utilization of resources. For example, high utilization is a warning to consider increasing the server's allocated resources to avoid negatively affecting the server's performance.
- Conduct a time-series analysis and view the trend of resource utilization of your server targets, and forecast values for a specific time period.
- View the aggregate of CPU and memory utilization for a specific time period.
- View the storage utilization split by operating system. You can see the list of hosts that have high and low resource utilization and also check the aggregate utilization across the selected time period.
- Identify the top server targets by a specific attribute. For example, you can group targets by operating system, CPU cores, lifecycle status, target version, cost center, or memory size.
- Conduct a categorical analysis to identify on which days a server has performed better or worse than others.
- View analytical information for host virtual machines and virtual platforms.

### Topics:

- [Typical Tasks for Analyzing Host and Virtual Machine Resources](#)
- [Analyze Host CPU and Memory Usage](#)

- [Analyze Host Storage Usage](#)
- [Analyze Virtual Machine CPU and Memory Usage](#)

## Typical Tasks for Analyzing Host and Virtual Machine Resources

Here's a list of tasks to analyze the utilization of your host and virtual machine resources using the Host Resource Analytics application.

Task	Description	More Information
Analyze host CPU and memory usage.	View the current aggregate CPU and memory utilization trend, compare resource utilization across hosts, and forecast future resource needs to plan your resources for a later time period. You can also check the usage pattern of CPU and memory or storage resources for a specific time period (for example, for last 30 days or last 12 months) together with common target properties.	<a href="#">Analyze Host CPU and Memory Usage</a>
Analyze server storage usage.	View the current space usage trend, compare storage utilization across servers, and forecast future storage resources for a later time period.	<a href="#">Analyze Host Storage Usage</a>
Analyze virtual machine CPU and memory usage.	View your virtual machine's aggregate CPU and memory utilization trends at a glance.	<a href="#">Analyze Virtual Machine CPU and Memory Usage</a>

## Analyze Host CPU and Memory Usage

You can analyze your host's CPU and memory utilization by using the Host Resource Analytics application to view and compare detailed information across one or more hosts.

From this application, you can:

- View usage analytics grouped by date and time.
- Identify a host's capacity and utilization of resources. For example, high utilization is a warning to consider increasing the host's allocated resources so as to avoid negatively impacting the host's performance.
- Group and sort hosts by a specific attribute. For example, you can group hosts by operating system, CPU cores, lifecycle status, target version, cost center, or memory size.
- View the current CPU and memory utilization of your hosts. Select one or more hosts for further analysis.
- Conduct time-series analysis and view the trend of resource utilization of your host targets, and forecast values for a specific time period.

- Download the data in CSV format.
- View CPU and memory utilization trends for each hosted entity in your hosts.
- Sort by entities that are top consumers of CPU and Memory resources.

 **Note:**

Data is displayed for the time selected from the global time selector. If no time range is selected, then data is displayed for the last 30 days. Use the time editor to alter the time period for which data is displayed. You can select a predefined time period or specify a custom one to filter the available data.

## Analyze Host Storage Usage

You can analyze your host's storage utilization by using the Host Resource Analytics application to view and compare detailed storage information across one or more hosts.

From this application, you can:

- Identify a host's resource utilization percentage. For example, high utilization is a warning to consider increasing the host's maximum storage capacity.
- Group and sort hosts by a specific attribute. For example, you can group hosts according to cost center, operating system, lifecycle status, or target version, and identify the hosts with the most used storage.
- View the current storage utilization of your hosts. Select one or more hosts for further analysis.
- Identify past trends and view a forecast of future storage needs for the hosts that you have selected.
- Download the data in CSV format.

 **Note:**

Data is displayed for the time selected from the global time selector. If no time range is selected, then data is displayed for the last 30 days. Use the time editor to alter the time period for which data is displayed. You can select a predefined time period or specify a custom one to filter the available data.

## Analyze Virtual Machine CPU and Memory Usage

You can analyze your virtual machine's CPU and memory utilization by using the Host Resource Analytics application to view and compare detailed CPU and memory information across one of more virtual machines.

From this application, you can:

- View usage analytics grouped by date and time.

- Identify a virtual machine's capacity and utilization of resources. For example, high utilization is a warning to consider increasing the virtual machine's allocated resources so as to avoid negatively impacting the virtual machine's performance.
- View the current CPU and memory utilization of your virtual machines. Select one or more for further analysis.
- Conduct time-series analysis and view the trend of resource utilization of your virtual machines, and forecast values for a specific time period.
- View the mapping from Virtual Platforms to VMs (that is, view which VMs are running on that particular Virtual Platform).

 **Note:**

Data is displayed for the time selected from the global time selector. If no time range is selected, then data is displayed for the last 30 days. Use the time editor to alter the time period for which data is displayed. You can select a predefined time period or specify a custom one to filter the available data.

## Analyze Exadata Resources

Get combined infrastructure resource utilization and workloads running across a set of Exadata systems (Exadata Database Machine, Exadata Cloud Service and Exadata Cloud at Customer). Analyze the capacity utilization for host and storage infrastructure, do a capacity planning for database workloads, and get insights into the database performance trends.

The Exadata Analytics application provides visibility into performance of critical workloads running on shared Exa infrastructure. A single integrated view of all Exadata systems across your enterprise helps identify issues and troubleshoot them quickly. You can also view the categorization of the Exadata systems based on the fleet type (Exadata Database Machine, Exadata CS, Exadata CC), rack type, the systems with low resource utilization, and the global resource usage in terms of CPU, memory, I/O, and storage.

As an IT administrator or a capacity planner, you can use this application to optimize capacity based on proactive planning.

### Topics:

- [Typical Tasks for Analyzing Exadata Resources](#)
- [Analyze Exadata Systems with the Integrated View of Logs and Analytics](#)
- [Analyze the Top Resource Consuming Exadata Systems](#)
- [Analyze the Critical Exadata Systems](#)
- [Analyze the Spare Capacity for New Workloads](#)

## Typical Tasks for Analyzing Exadata Resources

As an administrator or capacity planner, perform these tasks using Exadata Analytics to analyze how your different Exadata Database Machine, Exadata CS and Exadata CC resources are being used.

Task	Description	More Information
Monitor Exadata Database Machine, Exadata CS and Exadata CC systems with the integrated view of logs and analytics.	View and monitor the unified data about your enterprise-wide Exadata Database Machine, Exadata CS and Exadata CC systems to assess resource utilization and performance.	<a href="#">Analyze Exadata Systems with the Integrated View of Logs and Analytics</a>
Analyze the top resource consuming Exadata Database Machine, Exadata CS and Exadata CC systems.	View the Exadata Database Machine, Exadata CS and Exadata CC systems running with low resource utilization and high resource utilization. The different Exadata systems are categorized under 0–25%, 25–50%, 50–75%, and 75–100% utilization of resources.	<a href="#">Analyze the Top Resource Consuming Exadata Systems</a>
Analyze the critical Exadata Database Machine, Exadata CS and Exadata CC systems.	View the different Exadata Database Machine, Exadata CS and Exadata CC systems that are the top consumers of resources, or will reach the high utilization in the next 30, 90, 180, and 180+ days.	<a href="#">Analyze the Critical Exadata Systems</a>
Analyze the spare capacity for new workloads.	Analyze the different Exadata Database Machine, Exadata CS and Exadata CC systems with maximum resources available for the new workloads. These systems are forecasted to be using low resources for the next six months.	<a href="#">Analyze the Spare Capacity for New Workloads</a>

## Analyze Exadata Systems with the Integrated View of Logs and Analytics

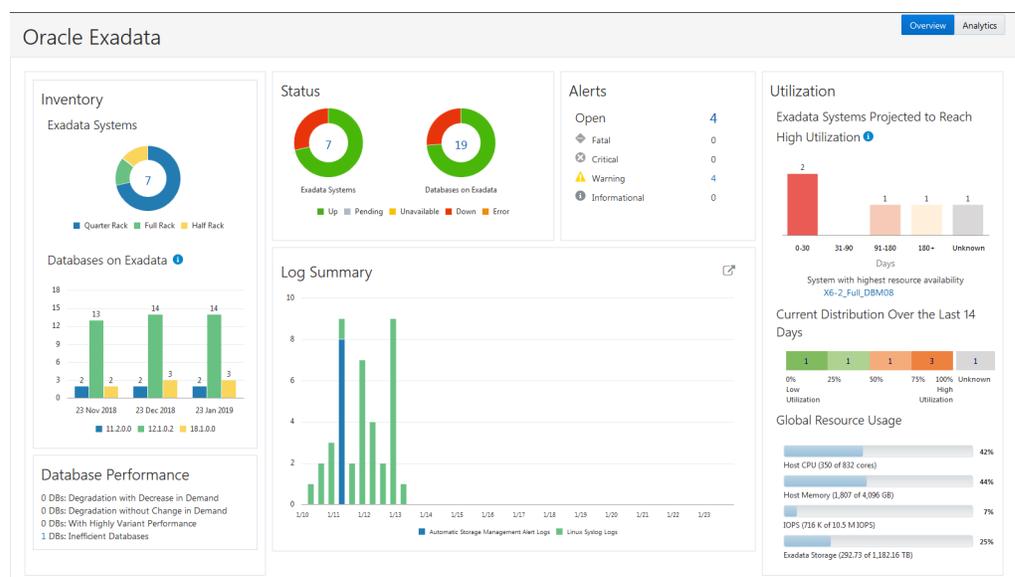
Oracle Management Cloud application provides a single, unified view of the different types of Exadata systems (Exadata Database Machine, Exadata CS and Exadata CC) across your enterprise. You can use the data analytics information from this view to identify current resource constraints and workloads to forecast your future resource requirements. You can also track down problems and resolve them quickly.

### Note:

This dashboard data assumes that your Exadata systems databases are managed by Oracle Enterprise Manager and data is made available to Oracle Management Cloud by using a Data Collector along with a management cloud agent. See Tasks 2 and 3 in [Perform the Prerequisite Tasks](#) for details on installing a data collector and a cloud agent.

To get a unified view of your enterprise-wide resource utilization and performance of Exadata Database Machine, Exadata CS and Exadata CC systems, on the Oracle Management Cloud home page, click the OMC Navigation Menu on the top-left corner, and then click **Exadata**. All aspects of monitoring (from Oracle Infrastructure Monitoring), logs (from Oracle Log Analytics) for troubleshooting, and performance analytics (from Database Performance Analytics) are integrated to provide this unified view.

**Figure 3-17 Exadata Overview Page**



The Overview page displays data for all Exadata systems. In addition, resource utilization and performance analytics is available for Exadata systems with Enterprise edition license type (that is, these Exadata systems are ITA enabled).

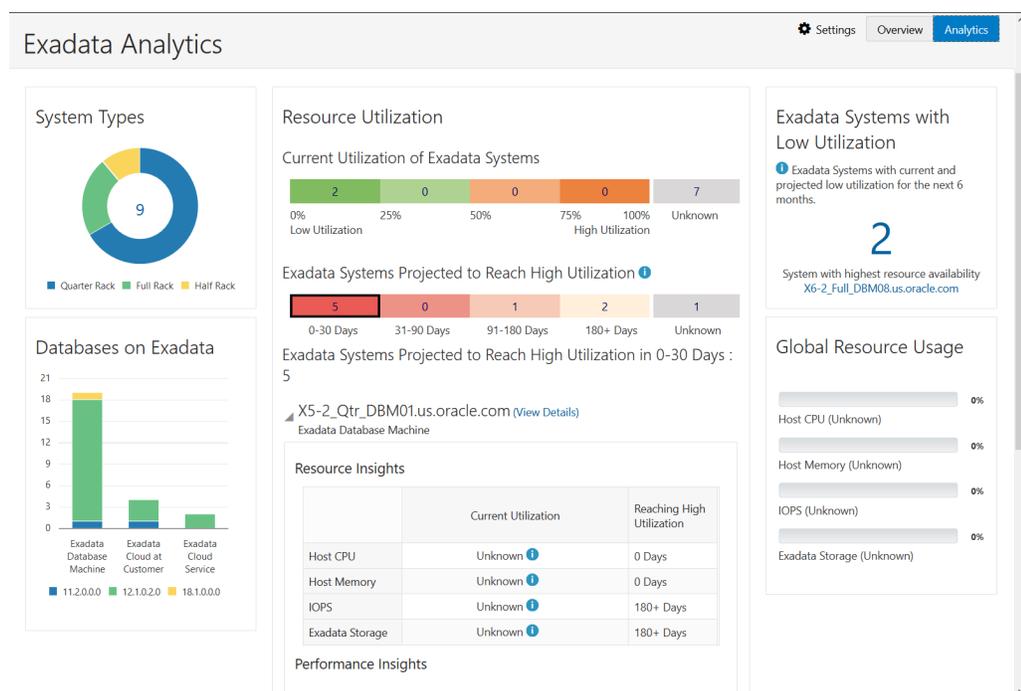
The unified view provides the following information:

- The categorization of the different Exadata systems based on type: Exadata Database Machine, Exadata CS and Exadata CC.
- The categorization of the different Exadata type systems across your enterprise, based on the rack type.
- The number of databases (including their version numbers) associated with Exadata systems for the past three months.
- Performance of your databases by listing the databases with performance degradation, databases with varying workload performance, and inefficient databases. You can drill down to further analyze the performance of these databases using the Database Performance Analytics application of Oracle IT Analytics.
- The status of each Exadata system along with the status of databases on these systems.
- The list of all open alerts grouped by their severity.
- The systems (with Enterprise edition license type) expected to reach high utilization in the next 30, 90, 180, and 180+ days. Based on the actual usage of

critical resources such as CPU, memory, IOPS, and storage over a period of time, ITA forecasts when these resources will reach their defined utilization threshold values (that you set in the IT Analytics Settings page) in the future.

- The number of Exadata systems running with 0–25%, 25– 50%, 50–75%, and 75–100% utilization of resources for the last 14 days. This categorization is based on the utilization of host CPU, host memory, IPOS, and Exadata storage.
- Summary of logs grouped by log source.
- The global resource usage for all Exadata systems independent of type (Exadata Database Machine, Exadata CS and Exadata CC) in terms of host CPU, host memory, IOPS, and Exadata storage.

**Figure 3-18 Exadata Analytics Dashboard**



The Analytics page displays data for all on-premise and cloud Exadata systems. These include:

- A breakdown of system types in a circular graph by Quarter Rack, Full Rack and Half Rack.
- Databases categorized by Exadata system (Exadata Database Machine, Exadata CC, Exadata CS), and by database version.
- The number of Exadata systems running with 0–25%, 25– 50%, 50–75%, and 75–100% utilization of resources for the last 14 days. This categorization is based on the utilization of host CPU, host memory, IPOS, and Exadata storage.
- The number of Exadata systems projected to reach a high utilization in the next 30 days. This categorization is based on the utilization of host CPU, host memory, IPOS, and Exadata storage.

- The global resource usage for all Exadata systems independent of type (Exadata Database Machine, Exadata CS and Exadata CC) in terms of host CPU, host memory, IOPS, and Exadata storage.

## Analyze the Top Resource Consuming Exadata Systems

You can analyze the resource utilization of the different types of Exadata systems (Exadata Database Machine, Exadata CS and Exadata CC) using Exadata Analytics, and thereby identify the spare capacity for new workloads.

The Current Utilization of Exadata Systems section shows the number of Exadata systems running with 0–25%, 25–50%, 50–75%, and 75–100% utilization of resources. This categorization is based on CPU, memory, I/O, and storage utilization. The current utilization of some Exadata systems cannot be calculated if the available data is insufficient. The resource utilization for such database systems is categorized as 'Unknown'.

The bucket with 0–25% resource utilization is considered as low utilization, and the bucket with 75–100% utilization of resources is considered as high utilization.

You can drill down to view the resource and performance insights for the different types of Exadata systems with the preferred range of resource utilization. Click the preferred range to view these details.

The **Resource Insights** table displays the current usage and the number of days after which the selected Exadata system is expected to reach capacity with respect to host CPU, host memory, IOPS, and storage.

The **Performance Insights** section lists the databases with performance degradation, databases with highly variant performance, and the inefficient databases in the selected Exadata system.

Click **View Details** for the selected Exadata system to view the following:

- Exadata system inventory
- Resource and performance insights of an Exadata system
- Top databases by the resource type CPU, memory, I/O, and storage
- Top hosts by the resource type CPU and memory
- Top Exadata storage servers by storage, I/O, and throughput

## Analyze the Critical Exadata Systems

You can analyze the critical Exadata systems (Exadata Database Machine, Exadata CC and Exadata CS) which may reach capacity in 30, 90, 180, 365, 365+ days, using Exadata Analytics.

The Exadata Systems Projected to Reach High Utilization section shows the number of all different types of systems that reach the high threshold in terms of capacity, in the next 30, 90, 180, 365, and 365+ days. You can analyze the CPU, memory, and I/O resources that are capacity bottleneck using the forecasted data. Use this information to identify the top databases consuming host CPU, memory, storage disk, or I/O resources.

You can drill down to check the different types of Exadata systems that are forecasted to run out of capacity in the future. Click on each of the different Exadata systems to view these details.

The **Resource Insights** table displays the current usage and the reaching capacity of the selected Exadata system with respect to CPU, memory, I/O, and Exadata storage.

The **Performance Insights** section lists the databases with performance degradation, databases with highly variant performance, and the inefficient databases in the selected Exadata system.

Click **View Details** for the selected Exadata system to view the following:

- Exadata inventory
- Resource and performance insights of an Exadata system
- Top databases by the resource type CPU, memory, I/O, and storage
- Top hosts by the resource type CPU and memory
- Top Exadata storage servers by the storage, I/O, and throughput

To resolve the capacity bottleneck:

1. Relocate the high resource consuming databases.
2. Find if the CPU usage increment is caused by bad SQL.
3. Consider adding more servers.

## Analyze the Spare Capacity for New Workloads

Analyze and identify the servers with most available capacity for placing new or migrating existing workloads.

The Exadata Database Machine, Exadata CS and Exadata CC systems with Low Utilization section shows the number of Exadata systems that currently have a low resource utilization, and are forecasted to maintain a low resource utilization for the next six months. The servers can be used for new workloads.

You can drill down to check for the databases, hosts, and Exadata storage servers with available resources.

# 4

## Maximize Performance and Availability

As an IT operations administrator, you can use Oracle IT Analytics to identify enterprise-wide systemic problems that affect the availability and performance of your applications.

Using Oracle IT Analytics, you can:

- Visualize top performance findings and trends across databases and middleware over time.
- Analyze activities of databases or middleware over time across related components or groups to view the effect of shared resources or application workloads.
- Analyze trends of top incoming web requests or SQL statements and their changes over time to detect changes in workload or degradation in response.

If you see a break in the graphs in the applications, then it could be because the entities weren't functioning during that time period, and so no data was collected. This is applicable to all entities and graphs.

### Note:

Oracle IT Analytics retains the raw data (the metric data that is collected and uploaded by agents on the target hosts) for eight days. Before purging the raw data, it rolls up the data into hourly and daily intervals. The rolled up data is then retained for 13 months to provide historical metric information. Data becomes inaccessible at the end of the 13-month period.

### Topics:

- [Read This Before You Begin](#)
- [Maximize Database Performance](#)
- [Analyze SQL Performance](#)
- [Maximize Application Performance](#)
- [Maximize Application Server Performance](#)
- [Analyze IT Server and Application Availability](#)

## Read This Before You Begin

Here are some of the common terms and basic concepts regarding the Performance Analytics applications.

**Automatic Workload Repository (AWR)**, part of an Oracle database, collects, processes, and maintains performance statistics for problem detection and self-tuning purposes. This data is both in memory and stored in the database.

**Garbage collection**, built into many programming languages, is an automatic way of managing the memory available to various objects. The garbage collection mechanism tracks objects still being used, marks the objects no longer in use as *garbage* and through the process of *garbage collection*, it automatically frees up the memory for use by other objects. In the case of Java, garbage collection is done within the Java Virtual Machine environment.

**Garbage collection overhead** refers to the additional resources, processing time, used by the garbage collection mechanism.

Software applications can choose to implement various **garbage collection types** depending on the performance goals. In Java, for example, developers can decide to implement garbage collection of types single-threaded (serial), multithreaded (parallel), concurrent mark then sweep (CMS) or parallel collection in multiple memory zones. Each method uses various resources and CPU time, affecting the overall application performance.

The **heap**, for example in the **Java heap**, is a repository of Java objects. The heap is the sum of active objects, dead objects (marked for garbage collection), and free memory. **Heap size tuning** in Java refers to minimizing the time that the JVM spends doing garbage collection while maximizing the number of clients that the application server can handle at a given time. A **heap map** provides the memory details of a given process.

## Maximize Database Performance

Oracle IT Analytics lets you get insights into the performance of your databases by viewing information about database performance degradation, databases with varying workload performance, database inefficiency, and Top SQL statements.

Problems in the performance of your databases affect your application performance. As an IT administrator, you can use Database Performance Analytics to analyze database performance issues across your enterprise, and identify the underlying systemic causes of the performance issues.

Using Oracle IT Analytics, you can find:

- Databases degrading in performance with decrease in demand
- Databases degrading in performance without change in demand
- Databases which are increasingly inefficient
- Top SQL statements
- SQL queries that are highly varied in performance

Oracle IT Analytics can be also be used to maximize database performance of both types of Autonomous Databases - Shared: Autonomous Data Warehouse –

Shared and Autonomous Transaction Processing – Shared. For more information on Autonomous Databases - Shared, see: [Discover Autonomous Databases](#) in *Using Oracle Database Management for Autonomous Databases*.

#### Topics:

- [Typical Tasks for Maximizing Database Performance](#)
- [Analyze Database Performance Degradation](#)
- [Analyze SQL Execution Performance](#)
- [Analyze Database Inefficiency](#)
- [Analyze SQL Performance](#)

## Typical Tasks for Maximizing Database Performance

Here's a list of tasks to analyze the performance of your databases using Database Performance Analytics.

Task	Description	More Information
Analyze database performance degradation and access workload stability.	<ul style="list-style-type: none"> <li>• Identify the databases that have degraded in performance.</li> <li>• Identify negative correlations between performance trends and demand.</li> <li>• Identify the databases that have varying performance.</li> </ul>	<a href="#">Analyze Database Performance Degradation</a>
Analyze the SQL statements with highly varying performance and the top SQL statements.	<ul style="list-style-type: none"> <li>• Identify SQL statements that are varying highly in performance.</li> <li>• Identify the top SQL statements across your entire databases environment.</li> </ul>	<a href="#">Analyze SQL Execution Performance</a>
Analyze workloads with unusual characteristics.	Analyze database inefficiency.	<a href="#">Analyze Database Inefficiency</a>
Analyze SQL performance	Identify degrading SQLs, variant SQLs, inefficient SQLs, and SQLs with multiple execution plans across databases.	<a href="#">Analyze SQL Performance</a>

## Analyze Database Performance Degradation

From the Performance Degradation section in Database Performance Analytics, you can identify the databases that have degraded in performance as follows:

- Degradation with decrease in demand
  - Change in activity greater than 20%
  - Change in demand less than -10%

- Degradation without change in demand
  - Change in activity greater than 20%
  - Change in demand between 10% and -10%

Database Performance Analytics lets you view the change in activity and demand for databases with the highest performance degradation.

Clicking the values for each category of databases (**Degradation with decrease in demand** and **Degradation without change in demand**) opens a list of the top databases within that specific category, along with the database type, change in activity, and change in demand for each. You can select a database from the list to view more information specific to that database. From the database drill-down page, you can:

- **View Performance Trend:** View the actual and forecasted response time by demand and activity or active sessions.
- **View Workload Performance Variability:** View the relative variability of SQL statements against response time variation. Variant SQLs are the ones whose relative variability is high. Highly variant SQLs are the ones that need attention.
- **View Inefficiency:** View a chart that provides a breakdown of inefficient waits, I/O, and CPU. See [Analyze Database Inefficiency](#).
- **Access Workload Stability:** You can identify databases with workloads running, and SQL statements with a high degree of performance variability. Varying SQL statements have a relative variability greater than 1.66, and highly varying SQL statements have a relative variability of more than 3. SQL statements that are highly varying in performance need attention to ensure they have a stable response time. To calculate the percentage of the workload caused by the varying SQL statements for each database:
  1. Calculate which SQL statements are varying over the time period.
  2. Calculate the elapsed time for each varying SQL statement.
  3. Divide the sum of the elapsed time for each varying SQL statement by the total elapsed time on the database.
- **Identify Top SQLs:** You can identify the top SQL statements across your entire database environment. Top SQL statements may consume an uneven amount of system resources. These SQL statements often cause a large effect on database performance and resource consumption. Identifying top SQL statements is an important SQL tuning activity that you must perform regularly.
- **Analyze workload with unusual characteristics:** Inefficient wait time is calculated using the I/O, CPU and inefficient waits time. Most databases are likely to have some inefficiency. However, over time, if the database is getting increasingly inefficient, and the Inefficient percentage is greater than 50%, then, as the database administrator, you must pay attention to this database to maximize its performance.

You can drill down further to view details of the SQL queries run on that particular database. See [Analyze SQL Execution Performance According to Database](#).

## Analyze SQL Execution Performance

Using Database Performance Analytics in Oracle IT Analytics, you can view information about the execution of SQL queries across your databases or for a specific database.

### Topics:

- [Analyze SQL Execution Performance According to Database](#)
- [Analyze SQL Execution Performance Across Databases](#)

## Analyze SQL Execution Performance According to Database

From the details page of the database, you can view the list of SQL queries run on that particular database. Use this information to identify the top 10 degrading SQL statements by response time and highly varying SQL statements by relative variability.

Click a query to drill down further. The application displays the following information:

- **Response Time:** View the maximum response time, minimum response time, average response time, and the percentage change in the response time.
- **SQL Performance:** View the trend of SQL performance over time, measured by response time, active sessions, or executions per hour. You can also view execution statistics (for example, elapsed time in seconds per execution and in total).

Click the value in the **SQL Text** field to view the complete text of the SQL query.

## Analyze SQL Execution Performance Across Databases

Click **Top SQL** in the **Databases with Variant Workload Performance** chart to view a load map of the SQL queries. Use the drop-down lists to specify what the size and color of each cell is determined by. Click a cell to view details specific to that query. To view the query text, click **SQL ID**. If you decide to group the SQL queries according to databases, then you can visually compare the details for the SQL query across databases.

## Analyze Database Inefficiency

Using Database Performance Analytics in Oracle IT Analytics, you can view information about the inefficiency of your databases.

- [Analyze Inefficiency Across Databases](#)
- [Analyze Inefficiency According to Database](#)

## Analyze Inefficiency Across Databases

The Database Performance Analytics application presents you with a list of your most inefficient databases. It also points out the increasingly inefficient databases among them. This lets you quickly identify which databases are most in need of attention.

From the Database Inefficiency chart, you can view the performance of your databases based on a database inefficiency percentage. Databases are grouped

according to their rate of inefficiency. Databases that are more than 50% inefficient and those that are increasingly inefficient are considered to be your most inefficient databases.

Click **Increasingly inefficient among the most inefficient databases** to view information specific to that database and to analyze it. See [Analyze Inefficiency According to Database](#).

## Analyze Inefficiency According to Database

In the Database Inefficiency chart, click **Increasingly inefficient among the most inefficient databases** to view detailed information about the performance of the selected database.

From the Activity chart, view the inefficiency of your database according to the number of active sessions involved in I/O, CPU, and inefficient waits. From here, you can determine whether the inefficient waits can be reduced to improve the performance of your application.

## Analyze SQL Performance

Using the SQL Analytics application of Oracle IT Analytics, you can analyze the SQL performance problems for enterprise-wide applications across a fleet of databases. In addition, Oracle IT Analytics provides trends and key insights to SQL performance issues thereby helping you to be proactive in avoiding future problems.



### Note:

You need to have Oracle Diagnostic Pack installed on your target database to be able to use SQL Analytics.

The SQL Analytics application home page displays the following sections:

- Performance Summary
- Performance Analysis

By using SQL Analytics, you can:

- View the degrading SQLs, variant SQLs, inefficient SQLs, and SQLs with plan changes.
- View the SQLs with the highest CPU and I/O usage.
- Analyze the performance of SQLs based on these categories – degrading, inefficient, variant, and with plan changes.
- Perform a SQL-wise performance analysis with the help of the average response time statistics.
- View the number of plans per SQL and get further insights to SQL performance based on details such as best performing plan, most executed plan, plan with the most I/O, and so on.

SQL Analytics can be also be used to analyze SQL performance of both types of Autonomous Databases - Shared: Autonomous Data Warehouse – Shared and Autonomous Transaction Processing – Shared. For more information on Autonomous

Databases - Shared, see: [Discover Autonomous Databases](#) in *Using Oracle Database Management for Autonomous Databases*.

**Topics:**

- [Typical Tasks for Analyzing SQL Performance](#)
- [Monitor SQL Performance Across Databases](#)
- [Analyze the Performance of SQL Statements](#)

## Typical Tasks for Analyzing SQL Performance

Here's a list of tasks to analyze the performance of your SQL requests by using the SQL Analytics application of Oracle IT Analytics.

Task	Description	More Information
Monitor the performance of SQLs based on the summary data, across databases.	Identify degrading SQLs, variant SQLs, inefficient SQLs, and SQLs with multiple execution plans across databases.	<a href="#">Monitor SQL Performance Across Databases</a>
Analyze the performance of SQLs on the basis of the performance category.	Analyze the performance of SQLs in different categories (degrading, variant, inefficient, and with plan changes) by using parameters such as percentage database time, number of SQLs, and size.	<a href="#">Analyze the Performance of SQL Statements</a>

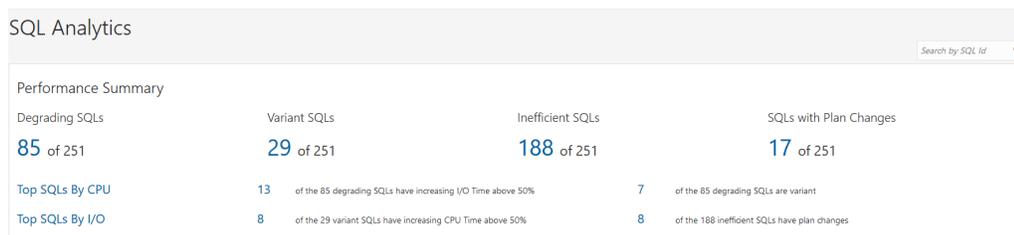
## Monitor SQL Performance Across Databases

The Performance Summary section of the SQL Analytics home page displays a snapshot of the performance categorization of SQLs across databases and applications.

The categorizations are:

- **Degrading SQLs:** SQLs with more than 20% increase in SQL response time, based on linear regression. The value of the SQL response time is derived from the total elapsed time divided by the total number of executions for the SQL.
- **Variant SQLs:** SQLs with a relative variability of more than 1.66. Relative variability of an SQL is measured by the standard deviation of the SQL response time divided by the average of the SQL response time. Those SQLs that have a relative variability of more than 3 are identified as SQLs with highly variant performance.
- **Inefficient SQLs:** SQLs with inefficiency of more than 20%. Inefficiency percentage of an SQL is derived from the inefficient wait time (wait time other than I/O, CPU, or idle wait time events) divided by the total database time.
- **SQLs with Plan Changes:** SQLs that utilize multiple execution plans.

**Figure 4-1 SQL Performance Summary Dashboard**



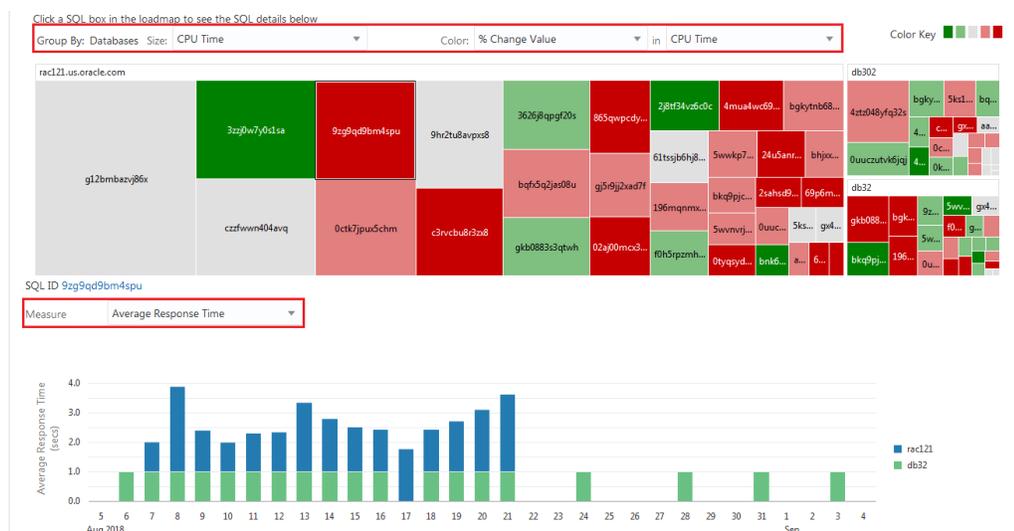
Clicking the SQL count of each categorization displays a heat map of all the SQLs across databases under the selected category. By default, the SQL with the maximum number of active sessions is selected. You can customize the heat map based on the following:

- **Size:** Customizes the size of the heat map segments based on Active Sessions, Average Response Time, Executions/Hour, I/O Time, and CPU Time.
- **Color:** Customizes the color coding of the heat map based on percentage change value or absolute value of Active Sessions, Average Response Time, Executions/Hour, I/O Time, and CPU Time.

When you select a SQL from the heat map, if the same SQL exists across databases, all SQLs get selected. This view helps you identify the SQLs that are common across databases. The SQL details (grouped by databases) are displayed in a bar chart in the section below the heat map. This display is based on the current time period. You can customize the bar chart display based on the following:

- Active Sessions
- Average Response Time
- Executions Per Hour
- I/O Time
- CPU Time

**Figure 4-2 SQL Performance Summary Chart**



The Performance Summary section also contains the following links to identify the high-load SQL statements across databases that consume a disproportionate amount of system resources and causes a large impact on the database performance:

- **Top SQLs By CPU:** SQLs with the highest growth in CPU usage
- **Top SQLs by I/O:** SQLs with the highest growth in I/O usage.

Clicking the links opens heat maps that display the SQLs with the highest growth in CPU and I/O usage differentiated by color keys.

Clicking each SQL in the heat map displays the SQL details (grouped by databases) in a bar chart in the section below the heat map. You can customize the bar chart display based on the following:

- Active Sessions
- Average Response Time
- Executions Per Hour
- I/O Time
- CPU Time
- Other Wait Time

## Findings and Performance Analysis

The goal of this section is to focus on the entire fleet of databases within your environment over a larger period of time, compared to other diagnostics tools that focus on immediate events. This tool tracks historical data, assisting in creating a more proactive action plan through custom analysis.

There are five main use cases that help you identify problem areas within your database fleet, these are viewed in the main **Performance Summary**.

**Figure 4-3 SQL Analytics Performance Summary Dashboard**

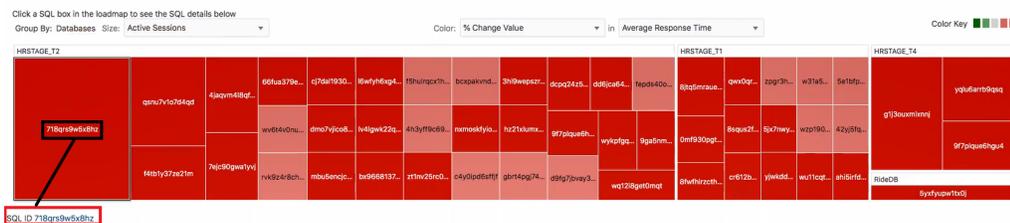


- **Degrading SQLs:**  
SQL Statements whose performance has degraded considerably over time.
- **Variant SQLs:**  
SQL Statements that have a significant variance in its executions times
- **Inefficient SQLs:**  
SQL Statements that are not optimized and/or have high wait times of I/O usage and high rates CPU utilization.
- **SQLs with Plan Changes**  
SQL Statements that are utilized by multiple plans with different results.
- **Tops SQLs by CPU and by I/O**

A ranking of all existing SQL Statements by CPU usage or by I/O utilization, ranking from highest to smallest usage percentage.

To begin your custom performance analysis, select the use case of interest by clicking on the impacted SQL statement number below each use case (Degrading, Variant, Inefficient, Plan Changes). This will take you to the heat map for each use case. Here you can analyze the performance of individual SQL Statements. The heat map is arranged in a friendly manner where you can easily view the SQL statements by database, size (represented by different sized squares) and rate of change (represented by a color key). When viewing the heat map click on a database of your choice to review; next, click on the matching *SQL ID* displayed below the heat map. This will load the data for that particular database; based on the use case selected.

**Figure 4-4 SQL Analytics Heat Map**



Once the SQL Statement has been selected you can review it's SQL Text and database location as well as it's key performance metrics under: **Performance Summary** and it's **Execution Plan Insights**

- Performance Summary:** This chart helps you view the trend in the SQL Statement performance for the current time period, based on the following measurement criteria:
  - The Average Response time in seconds.
  - The change percentage in response time. A negative value indicates an increase in the response time.
  - The number of SQL Statements that were executed per hour.
  - The variability value of the SQL Statement, which indicates the extent of variance in the SQL Statement's performance in the current time period.
  - The time range within which the maximum executions of the SQL Statement occurred.
  - The inefficiency percentage which is based on the idle wait time or non-productive time of the SQL Statement.
- Execution Plan Insights:** This section of the page is available only if the SQL Statement uses more than one execution plan. It provides the following information:
  - Plans Used
  - The best and worst performing plans
  - Plans with the most CPU and I/O usage
  - The most executed SQL Plan

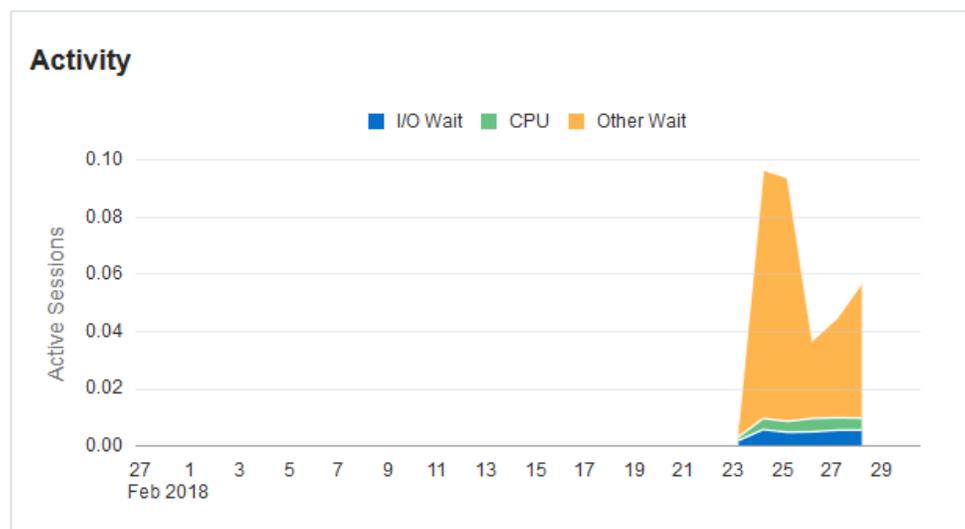
 **Note:**

The following screen descriptions are regardless of use case scenario. All use cases use the same second and third level screens.

In the **Home** tab of the page, you can view the representation of the performance data in the form of different charts.

- **Activity:**  
The Activity chart displays the session activity of the SQL Statement, for the current time period. It also shows how the SQL Statement activity is categorized into different wait classes indicated by the color, as described in the legend on the chart.

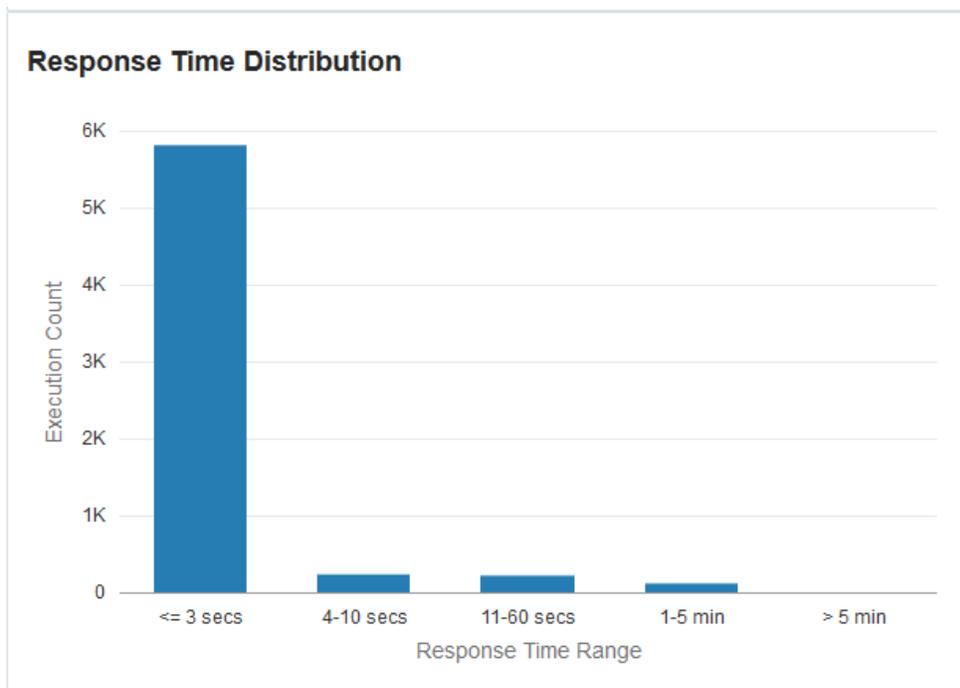
**Figure 4-5 Session Activity Chart for a SQL Statement**



For example, this chart indicates that the SQL Statement spent a significant portion of its idle time in the 'Other Wait' category and some portion in the 'CPU' and 'I/O' categories between 23 and 28 March 2018.

- **Response Time Distribution:** This bar chart plots the number of executions of the SQL Statement for a given response time range.

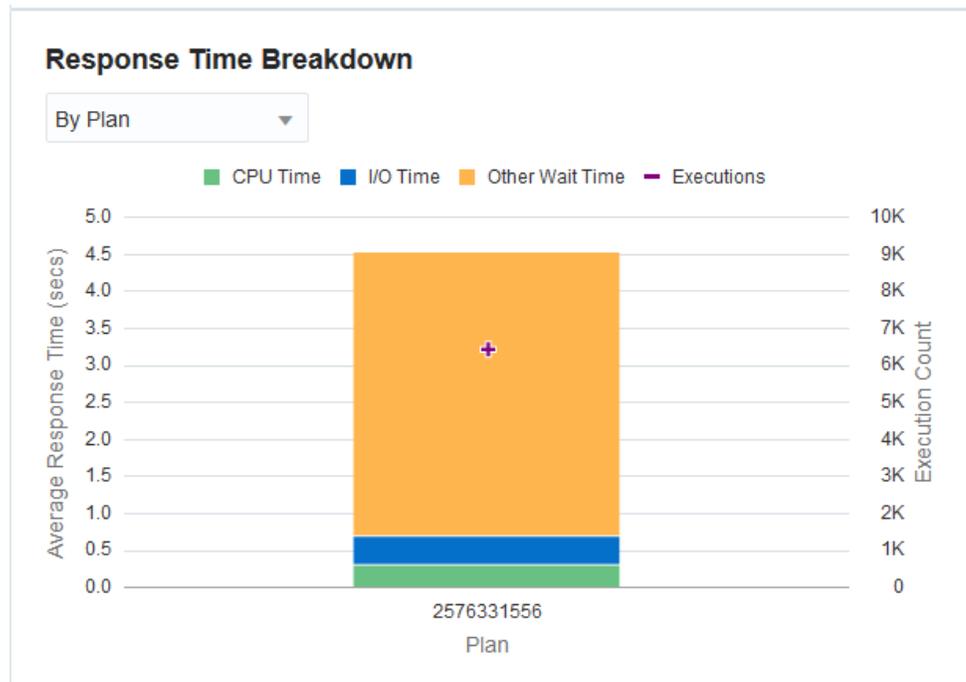
Figure 4-6 Response Time Distribution Chart



For example, this chart indicates that the maximum number of SQL executions occurred within three seconds. With this information, you can check on the conditions that favored a faster response time and analyze the reasons why other SQL executions had a relatively slow response time.

- **Response Time Breakdown:** The chart displays information by plan or by database and helps you identify where the SQL Statement spent most of its time by breaking down the average response time into different wait classes such as CPU Time, I/O time, and so on.

Figure 4-7 Response Time Breakdown Chart

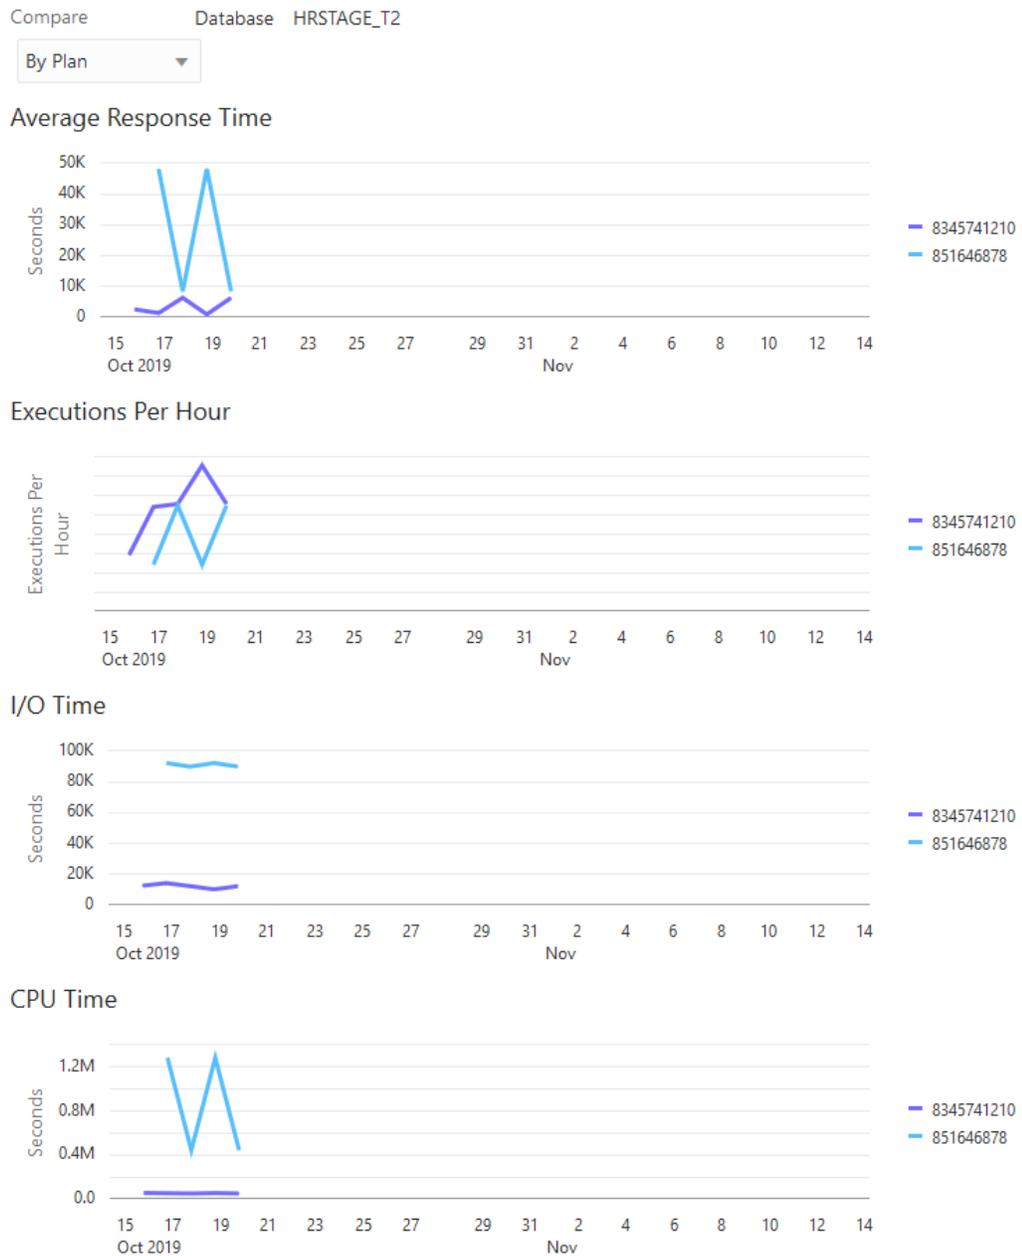


For example, this chart indicates that the SQL Statement used a single execution plan and a significant number of its executions spent a high time in the 'Other Wait Time' category and hence, has resulted in a slow response time.

In the **Compare** tab of the page, you can view the performance trend by comparing two or more plans or databases based on the following parameters:

- Average Response Time
- Executions Per Hour
- I/O Time
- CPU Time

**Figure 4-8 Compare Performance Trends**



In the **Plan** tab of the page, you can view a drill down version of any selected SQL execution plan performance.



**Note:**

To enable the **Plan** tab, first you must enable SQL Text and SQL Plan data collection, for more information please see: [Disable and Re-Enable SQL Execution Plan Collections](#).

The Execution Plan is presented in a top-down view by operation, you can further see the Operation Cost, Estimated Rows and Estimated bytes of each operation; allowing you to locate problematic operations. The higher the operation cost, the more problematic the operation can be.

**Figure 4-9 Execution Plan Breakdown**

Operation	Object	Information	Pruning	Operation Cost	Est.Rows	Est.Bytes
0	SELECT STATEMENT			0	0	0
1	VIEW			241.9M	665	30.6K
2	SORT GROUP BY			241.9M	665	111.1K
3	VIEW			201.9M	665	111.1K
4	HASH GROUP BY			201.9M	665	111.1K
5	VIEW			161.9M	665	111.1K
6	MERGE JOIN PARTITION OUTER			161.9M	665	196.8K
7	SORT JOIN			79.5M	5	645
8	VIEW			39.7M	5	645
9	HASH UNIQUE			39.7M	5	645
10	TABLE ACCESS FULL	PRODUCT_CATEG		36.5K	5	645
11	SORT PARTITION JOIN			82.4M	665	111.1K
12	VIEW			42.4M	665	111.1K
13	HASH GROUP BY			42.4M	665	162.9K
14	HASH JOIN			2.4M	665	162.9K

In this example we have highlighted in red a MERGE JOIN PARTITION OUTER command high in Cost, Rows and Bytes that is causing issues in SQL Plan 8345741210.

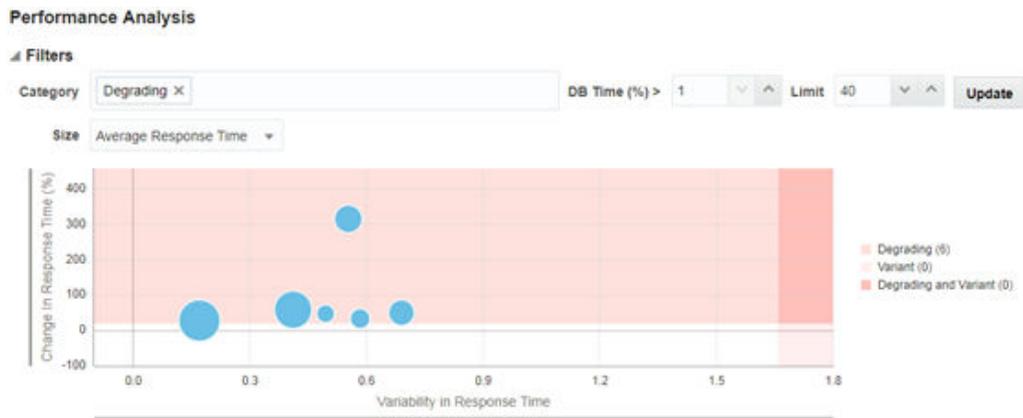
## Analyze the Performance of SQL Statements

The Performance Analysis section of the SQL Analytics home page displays a graph of the performance of a set of SQL Statements in their databases based on the percentage database time.

For example, to view a graph of the degrading SQL Statements among a set of 40 SQL Statements that are taking more than 1% database time in their respective databases:

1. Select **Degrading** in the **Category** field.
2. Enter 1 in the **DB Time (%) >** field.
3. Enter 40 in the **Limit** field. The **Limit** field specifies the number of SQL Statements that you want to base your graph on.
4. Click **Update**.
5. (Optional) Select **Average Response Time** in the **Size** drop-down list. The display of the graph changes based on the value that you select from the **Size** drop-down list. The other available values are **Active Sessions** and **Executions Per Hour**.

Figure 4-10 SQL Performance Analysis Graph



The bubbles on the chart denote the degrading SQL Statements. Hovering the mouse cursor over a blob displays the SQL details. For information about analyzing the performance of a SQL, see [Findings and Performance Analysis](#).

Similarly, you can view a graph with top SQL Statements that are variant, inefficient, and/or with plan changes by selecting the relevant values in the **Category** field.

## Maximize Application Performance

Oracle IT Analytics lets you monitor the performance of your applications by viewing the page performance, server request performance, page views, page load time, tier response breakdown, error distribution, and a profile of the users.

Using Application Performance Analytics, you can:

- View the number of pages that contribute to 60% of the total page views, or more.
- View the page load breakdown in terms of first byte time, interactive time, and page ready time.
- View the total number of server request calls, the response time for these calls, and the number of errors encountered during these calls.
- View the tier response breakdown in terms of database time, app server time, and external time.
- Look at the page views based on the user profile by region, the type of browser used, and the operating system of the device used to access the application.

You can also view these details for one or more entities. You can select the name of the entity from the drop-down list that is located next to the time period drop-down list.

### Topics:

- [Typical Tasks for Maximizing Application Performance](#)
- [Analyze Page Performance](#)
- [Analyze Server Request Performance](#)

## Typical Tasks for Maximizing Application Performance

Here's a list of tasks to analyze the performance of your applications using Application Performance Analytics.

Task	Description	More Information
Analyze page performance.	Analyze the performance of your application pages through page views, page load time, and page errors.	<a href="#">Analyze Page Performance</a>
Analyze server request performance.	Analyze the performance of your server requests through total calls, response time, and server request errors.	<a href="#">Analyze Server Request Performance</a>

### Analyze Page Performance

Using Application Performance Analytics, you can analyze the performance of the different pages of your application.

From the main application page, you can view the following:

- Total page views and the percentage increase or decrease in the number of page views during the current time period. A decrease in total page views indicates that either users are viewing the page less or fewer users are accessing your page. If the change in page views is a cause for concern, then the application will present it to you as a finding and lets you view detailed information.
- Page load time and the percentage increase or decrease in the page load time during the current time period. An increase in the page load time indicates that users have a longer wait before the page loads and they can see the information that they want. If the change in page load time is a cause for concern, then the application will present it to you as a finding and let you view detailed information.
- Total page errors and the increase or decrease in page errors during the current time period. If the change in page errors is a cause for concern, then the application presents it to you as a finding and lets you view detailed information.

 **Note:**

By default, the page details are displayed for the last 30 days. You can use the date-time list to select a different time window for your analysis. The specified date range applies to all the data that is displayed on the page. You can use one of the following options as the time period for the analysis:

- Last 24 hours
- Last 7 days
- Last 14 days
- Last 30 days
- Last 90 days
- Last 12 months
- Custom: Select a start date and end date for the analysis period.

Clicking a finding lets you view detailed information related to the finding.

For example, if your finding is that a page contributed to 80% and more of your total pages during the selected time period, then you can view a summary of the information related to the following in the **Summary** tab of the **Application Performance Analytics > Page Performance** page:

- General statistics related to total page views, page errors, load time, first byte time, and document interactive time, along with the percentage change since the previous time period for each. This information is similar to the information in the main page of the application.
- Trend for the page views, page errors, load time, first byte time, and interactive time for the page. The trend lets you identify time periods during which the values were unusually high or unusually low. From here, you can identify whether the change in page view percentage from the previous time period was due to consistently high or low values.
- User profile according to the region, browser used, and the operating system of the device used to access the application. From here, you can identify whether the increase or decrease in page performance was across regions, operating systems, and browsers, or whether specific regions, operating systems and browsers see a decrease in page performance.

This information lets you identify possible reasons for decreased page performance.

You can select a different finding from the drop-down list to view information specific to that finding.

In addition, the **Percentage Findings** option lets you check the details of pages that contribute towards a specific percentage of page views. For example, you may want to check the details of the four pages that have contributed towards 60% of the total page views. You can select the required parameter from the drop-down list. You can also perform a page-wise analysis. The **Select Pages** option lets you select the required pages from a heat map. Select the pages and click **Analyze**.

You can also perform a detailed analysis in the **Advanced** tab of the **Page Performance** page:

- View data distributed according to time, to identify specific days on which the application was accessed more than usual. With this information, the application owner can work with the capacity planner to ensure that the necessary resources are in place to handle the increased traffic on those days.

**Note:**

To perform a categorical analysis, you can also group data by different dimensions including time, browser, geographical region, and operating system.

- Identify correlation between the different metrics (for example, page views and first byte time) related to page performance.
- Forecast future metrics (such as page views) to decide whether to expand available resources to ensure that the application performs well.

In the **Report** tab of the **Page Performance** page, you can view the details of the analysis presented in a tabular format.

## Analyze Server Request Performance

Using Oracle Application Performance Analytics, you can analyze server request performance for your application.

From the main application page, you can view the following:

- Total calls and the percentage increase or decrease during the current time period. If the change in total calls is a cause for concern, then the application will present it to you as a finding.
- Total response time and the percentage increase or decrease during the current time period. An increase in the total response time means that the application takes longer to load information for the user. If the change in response time is a cause for concern, then the application will present it to you as a finding and let you view detailed information.
- Total server request errors and the percentage increase or decrease during the current time period. An increase in the total server request errors indicates that more errors are being generated. If the change in server request errors is a cause for concern, then the application will present it to you as a finding.

 **Note:**

By default, the server request details are displayed for the last 30 days. You can use the date-time list to select a different time window for your analysis. The specified date range applies to all the data that is displayed on the page. You can use one of the following options as the time period for the analysis:

- Last 24 hours
- Last 7 days
- Last 14 days
- Last 30 days
- Last 90 days
- Last 12 months
- Custom: Select a start date and end date for the analysis period.

Click a finding to view detailed information related to the finding.

For example, if your finding is that a server contributed to 40% of the total response time during the specified time period, then you can view the following information in the **Summary** tab of the **Application Performance Analytics > Server Request Performance** page:

- General statistics about the total calls, errors, response time, application server time, database time, and external time in the selected time period, along with the percentage change since the previous time period. This information is similar to the information in the main page of the application.
- Trend of the calls, errors, response time, application server time, database time, and external time for the server request.

In addition, the **Percentage Findings** option lets you check the details of server requests that contribute towards a specific percentage of calls. For example, you may want to check the details of the three server requests that have contributed towards 40% of the total calls, or more. You can select the required parameter from the drop-down list. By default, the **Server Performance** page displays the analysis for the top five server requests. However, you can also perform an analysis for a different set of server requests. The **Select Server Requests** option lets you select the required requests from a heat map. After you select the requests, click **Analyze** to view the details.

You can also perform a detailed analysis in the **Advanced** tab of the **Server Request Performance** page:

- View the server requests distributed according to time, to identify specific days of the week or hours during the day on which the number of requests were more or less than usual .

 **Note:**

To perform a categorical analysis, you can also group data by different dimensions including calls, errors, response time, application server time, database time, and external time.

- Identify correlation between the different metrics (for example, calls and database time) related to server request performance.
- Forecast future metrics (such as calls, response time, and so on) to decide on ways to improve server performance.

In the **Report** tab of the **Server Request Performance** page, you can view the details of the analysis presented in a tabular format.

## Maximize Application Server Performance

App Server Performance Analytics provides analytics insights for heterogeneous application servers including garbage collection overhead, workload, resource usage, availability, configuration profiles, and so on.

This application lets you analyze the performance of your application servers. You can analyze issues related to availability, garbage collection, Java Database Connectivity (JDBC) data source, heap, host memory, and issues based on insights. The application analyzes available data and presents you with findings or insights specific to the performance of your application servers. For example, the application presents the hour of the day with the peak load across your application servers.

You can access the application either from the Oracle Management Cloud home page or through the App Server Performance Analytics dashboard. You're directed to a dashboard where you can:

- View availability across your application servers and identify servers that are down.
- Perform an inventory analysis to identify the different operating systems and Java Development Kit (JDK) versions running across your application servers.
- Identify issues with host memory, which could lead to possible performance issues of your application servers.
- Monitor garbage collection overheads and view detailed information about Java Database Connectivity (JDBC) connection waits, requests and failures, and server requests.
- View the total server request count and total server request count time.
- Correlate metrics to identify the effect that changes on one metric have on another metric.

### Topics:

- [Typical Tasks for Maximizing Application Server Performance](#)
- [Analyze Application Server Performance](#)
- [Analyze Application Server Workload Performance](#)

## Typical Tasks for Maximizing Application Server Performance

Here's a list of tasks to analyze the performance of your application servers.

Task	Description	More Information
Analyze the performance of your application servers.	View the availability of your application server, identify any memory issues experienced by the host, and analyze the performance through distribution analysis and inventory analysis.	<a href="#">Analyze Application Server Performance</a>
Analyze the workload performance of your application servers.	Analyze the workload performance of your application servers through garbage collection and Java Database Connectivity (JDBC) Data source data. You can view the forecasted trend and correlation, and perform a distribution analysis for selected servers.	<a href="#">Analyze Application Server Workload Performance</a>

## Analyze Application Server Performance

App Server Performance Analytics provides you with key information and findings that let you quickly and effectively analyze the performance of your application servers.

When you go to the application from the Oracle Management Cloud home page or the corresponding dashboard, you can view information for the following:

- Garbage Collection Issues:** You can identify your server applications that are over the garbage collection limit. A long garbage collection period could affect the amount of time server requests have to wait in the queue before they're fulfilled by the application server. The threshold limit for garbage collection is set by the application, and is not user-defined. The application presents you with a finding based on the collected data. The following is an example of an application finding:

```
12 have JDK version 1.8.0_60.
23 have increasing Request Rate
```

According to this finding, based on the collected data, 12 of your application servers that are over the garbage collection limit have the JDK version 1.8.0\_60, which might be causing problems. Also, 23 of your application servers have an increasing request rate.

- Heap Issues:** You can view the application servers with a heap usage above a specified percentage. The application provides you with insights or findings based on the available data. The following is an example of an application insight:

```
21 have Heap Size above 40GB.
23 have an increasing Request Rate
```

According to this finding, based on the collected data, 21 of your application servers show a high heap size of above 40 GB. Also, 23 of your application servers have an increasing request rate.

- **Availability:** You can view the critical events across your application servers. If no server is down, then the application lets you know that there are no critical events across your application servers.
- **Memory Issues:** In some cases, issues with host performance can affect the performance of your application server.
- **Workload Performance:** You can drill down to view information specific to the performance of selected servers. The following is an example of an application finding related to workload performance:

```
12pm-1pm has the highest number of request 7.7M
Thursday has the highest number of request 30.3M
```

According to this finding, based on the collected data, the peak hour with the highest number of requests (7.7 M requests) was at 12 p.m. to 1 p.m.. Also, the day with the highest number of requests (30.3 M requests) was Thursday.

- **Distribution Analysis:** You can view information about the percentage of time spent in garbage collection and the total JDBC connection waits.
- **Inventory Analysis:** From here, you can view information about the JDK version, operating system vendor or version, and the XMX size across your application servers.

You can decide to view information for all your application servers or only for specific servers. You can also select the time period for which you want the application to load the information.

## Analyze Application Server Workload Performance

App Server Performance Analytics lets you analyze the workload performance of your application servers through measuring the volume of requests processed and the average response time.

Using App Server Performance Analytics, you can:

- Identify the peak time with a high volume of server requests.
- Analyze changes in workload performance brought about by changes in garbage collection overhead.
- Identify how changes in JDBC connection time affect workload performance.
- View information for the total request count and total request count time.

From the main page of the application, you can view workload performance data across your application servers, by:

- **Server Request Count:** This is the total number of requests serviced by the application server
- **Server Request Processing Time:** This is the time taken by your application server to process each request

You can drill down to view detailed information about the volume of processed requests and average response time through the following:

- A distribution chart to view the percentage of servers with increasing, decreasing, or constant number of server requests and processing time.
- A Distribution Analysis histogram that groups application servers according to percentage time in garbage collection and total JDBC connection waits. From this chart, you can drill down to view detailed information for the application servers within a selected bucket.

From the Workload Performance drill-down page, you can view a breakdown of the servers in the bucket that you had selected in the Distribution Analysis histogram. This information is presented in the form of a heap map, where each cell represents a server and the color of the servers is determined by the rate of change in the metric selected.



**Note:**

By default, information for the top five application servers is loaded, but you can choose to view information for all application servers in the selected bucket.

You can view the server breakdown according to the following:

- Server Request Count
- Server Request Processing Time
- Garbage Collection Invocation Count
- Percent Time in Garbage Collection
- JDBC Data source Wait Count
- JDBC Connection Request Count
- Heap Usage After Garbage Collection

Use the drop-down list to select what metrics to group application servers by.

Select a single server, multiple servers, or a group of servers to view detailed information for the following:

- **Total Server Request Count:** You can view the total number of server requests serviced by the selected application servers and the percentage change over the selected time period.
- **Average Server Request Processing Time:** You can view the average time taken by the selected application servers to process requests, and the percentage change over the selected time period.
- **Total Garbage Connection Invocation Count:** You can view the number of times that the selected servers called for garbage collection and the percentage change during the specified time period.
- **Total JDBC Datasource Wait Count:** You can view the total JDBC data source waits across the selected application servers, and the percentage change during the specified time period.
- **Total JDBC Connection Request Count:** You can view the total JDBC connection requests across the selected application servers, and the percentage change during the specified time period.

- **Total Heap Usage After Garbage Collection:** You can view the total heap usage for the selected application servers after garbage collection is completed and the percentage change during the specified time period.
- **Server Trend:** This chart displays the trend of the selected servers over the specified time period.

You can also perform an advanced analysis through the following:

- **Forecast Trend:** You can view the forecasted trend of server performance in terms of garbage collection overheads or JDBC data source count.
- **Distribution Analysis:** You can view the distribution of the percentage of time spent in garbage collection or the JDBC connection waits for each of the selected servers.
- **Correlation:** You can view the correlation between the metric that most affects the Server Request Processing Time, according to specific measures. For example, you can view the degree to which the Percent Time Spent in Garbage Collection correlates with the Server Request Processing Time, by Server Request Count.

## Analyze IT Server and Application Availability

The Availability Analytics application under Oracle Oracle IT Analytics provides you with visibility into the overall health of your IT systems and applications.

Availability Analytics provides insights into the following:

- Availability of your IT systems and applications through Health Score metrics, to determine whether your availability is getting worse over time.
- Systems and applications having frequent outages, to determine which system or application requires special attention.
- Average time required to recover, average time between outages, and how much time it took to recover from an outage, to help businesses take outages into consideration when planning to meet their service-level agreements.

The application analyzes the collected data and provides you with insights or findings, which you can use to identify possible issues with the availability of your IT servers and applications.

### Topics:

- [Typical Tasks for Analyzing IT Server and Application Availability](#)
- [Assess the Health of Your IT Servers and Applications](#)
- [Analyze the Availability of Your IT Servers and Applications](#)

## Typical Tasks for Analyzing IT Server and Application Availability

Here's a list of tasks to analyze the availability of your IT server and application availability.

Task	Description	More Information
Assess the health of your IT servers and applications.	Use the data chart and the displayed application finding to view the overall health score across your enterprise, and to identify the target with the lowest health score. You can also customize the criteria for calculating the health score of your targets.	<a href="#">Assess the Health of Your IT Servers and Applications</a>
Analyze the availability of your IT servers and applications.	Use the Availability drill-down page to view detailed information about the availability of your enterprise targets. You can view the targets with the lowest availability, targets with the most downtime events, the average time taken for your targets to recover after a downtime event, and the average time between successive downtime events.	<a href="#">Analyze the Availability of Your IT Servers and Applications</a>

## Assess the Health of Your IT Servers and Applications

Oracle IT Analytics rates the health of your IT servers and applications and provides you with valuable insights about the trend of your health score.

**Health Score** is used to assess whether the target is behaving as expected, and is determined by a weighted combination of the following:

- Total downtime for the target
- Number of times the target went down
- Number of fatal or critical events recorded on the target
- Duration of the fatal or critical events experienced by the target

 **Note:**

You can change the health score settings to decide what metrics should be considered when calculating the health score of your targets.

The application also displays the server with the lowest health score.

Clicking the **Health Score** tile on the main page of the application directs you to more information about the health of your servers and applications. You can view a summary of the health score across your servers and applications, or a report with your IT servers and applications arranged according to health score, mean time to recover, mean time between failures, total downtime, total down events, and availability (%).

From the main page of the application, you can also view the **Health Score Breakdown** according to your targets.

## Analyze the Availability of Your IT Servers and Applications

Availability Analytics assesses and displays the availability across your targets, along with a related finding.

From the main page of the application, you can view the **Availability Breakdown** according to your targets.

You can also view the **Availability** across your targets. Availability Analytics presents you with the target with the least availability.

Clicking the Availability tile directs you to more information about the availability of your targets. You can choose to view a summary of the metrics that affect the availability of your targets, or a report with your targets arranged according to Availability percentage.

The **Report** tab lists all your targets according to Availability percentage. Use the drop-down menu to change by what metrics to arrange the targets. Against each target under the **Summary** tab, you can view the mean time to recover (MTTR), mean time between failures, total down events and total downtime, and availability (%) and health score.

Click a target under the **Report** tab to view a target-specific **Availability Report**. From here, you can view calculations for Availability (%), Mean Time to Recover, and Mean Time before Failure. You can also view the Availability Divergence, and the summary of entity down events.

# 5

## Explore Data and Build Dashboards

Oracle IT Analytics lets you search and explore your IT operational data to find information specific to your needs, and to analyze the data by various dimensions. You can save important analyses as widgets, build dashboards, and present custom views for your stakeholders, such as line of business executives and analysts.

Data Explorer lets you perform custom and advanced analytical tasks that are beyond the possibilities offered to you by the other Oracle IT Analytics applications.

Through Data Explorer, you can:

- Search and explore real-time and historical IT operational data, including IT assets, configuration metrics, events, and topologies, by setting your own criteria.
- Conduct a time-series analysis and view the trend of resource utilization of your database and middleware targets, and forecast values for a specific time period.
- Identify the top database and middleware targets by a specific attribute. For example, you can view data for the top 10 Oracle WebLogic Server servers by heap.
- Conduct a categorical analysis to identify on which days your database and middleware targets have performed better or worse than on other days.
- Analyze events by searching for data specific to critical events and alerts, grouped by targets or target types.

### Note:

Oracle IT Analytics retains the raw data (the metric data that is collected and uploaded by agents on the target hosts) for eight days. Before purging the raw data, it rolls up the data into hourly and daily intervals. The rolled up data is then retained for 13 months to provide historical metric information. Data becomes inaccessible at the end of the 13-month period.

### Topics:

- [Read This Before You Begin](#)
- [Typical Tasks for Searching and Exploring Data](#)
- [Access Data Explorer](#)
- [Search, Explore, and Analyze Data Using Data Explorer](#)
- [Save Your Analysis](#)
- [Build Dashboards](#)

## Read This Before You Begin

Here are some basic concepts regarding exploring data and building dashboards.

A **data explorer**, in the context of Oracle Cloud services, is an interface that lets you explore collected IT operational data, perform forecasting, view trends, customize the display of this data, and save it for future use. In this case, views of saved searches are called **widgets** and a collection of custom widgets can form a **dashboard**.

A data explorer **target** is the entity whose data is being collected, for example the Oracle WebLogic Server. Targets can be of various **types**, for example the Oracle WebLogic Server is a Middleware type of target.

Data collected from various targets is in the form of **metrics**.

**Heat maps** (also known as treemaps) are graphical representations of data displayed as a combination of colored rectangular cells, each representing an element of data. Heat maps help visualize the state and effect of a large number of elements at once.

**Machine learning** algorithms are constantly looking at the data collected by Oracle Management Cloud. There are four major classes of machine learning algorithms:

- **Anomaly detection** is the capability to study behavior-specific metrics over time.
- **Clustering** is applied to unstructured data such as logs, to identify patterns, and then match the data to that pattern to quickly identify trends and outliers in logged data.
- **Prediction** is the capability to forecast future values of time-series metrics based on historical data and system behavior.
- **Correlation** is the ability to determine how sets of metrics correlate in behavior.

## Typical Tasks for Searching and Exploring Data

Here's a list of tasks to search and explore data using Data Explorer, and to save these searches to create dashboards in Oracle IT Analytics.

Task	Description	More Information
Access the data explorer.	Access the Data Explorer from a dashboard.	<a href="#">Access Data Explorer</a>
Search, explore, and analyze data using Data Explorer.	Search and explore data by specific attributes to view information specific only to the issue that you're troubleshooting. Analyze your data using the different visualization options.	<a href="#">Search, Explore, and Analyze Data Using Data Explorer</a>
Save your analysis	Save your analysis as a widget	<a href="#">Save Your Analysis</a>
Build dashboards	Use your saved searches to build dashboards for future use.	<a href="#">Build Dashboards</a>

## Access Data Explorer

You can access Data Explorer either directly from the Oracle Management Cloud home page, or through a dashboard.

When you access Data Explorer from a dashboard, the filters and visualization settings are already loaded. However, when you access Data Explorer directly, you have to build your analysis from the beginning.

## Search, Explore, and Analyze Data Using Data Explorer

Use the attributes in Data Explorer to filter real-time and historic data by targets and other attributes so that you can view information specific to the issue that you're troubleshooting.

Data Explorer lets you:

- Search and explore data by targets. For example, you can view the attributes specific to a target that's experiencing frequent unplanned downtime.
- Search and explore data for all targets by specific attributes. For example, you can view the resource utilization for all your databases, to identify which databases require more resources.

Through Data Explorer, you can search for data, select your visualization options, and analyze your data.

For example, you can use Data Explorer to:

- View the trend of CPU and memory usage for your database targets (such as Oracle Database, MySQL), for the past 13 months, and forecast the values for the next three months, to detect seasonal trends.
- Identify all targets with a CPU usage of more than 50% that are running on a specific software version.
- View the 25 cost centers with the highest CPU utilization.

### Data Explorer Visualization Options

Data Explorer lets you visualize data in the following ways:

Visualization Option	Description
Analytics Line	View the data as a line, with points marked out for analysis. Specific attributes are measured along the x and y axes, and a third attribute determines the color of the line.  Use this option for Trending and Forecasting

Visualization Option	Description
Area	<p>View the data as a filled-in area. Specific attributes are measured along the x and y axes. Data can be grouped according to a third attribute.</p> <p>Use this option for identifying the top or bottom X targets by attribute.</p>
Bar	<p>View the data as a bar diagram. Specific attributes are measured along the x and y axes, and a third attribute determines the color of the bars.</p> <p>Use this option for identifying the top or bottom X targets by attribute.</p>
Cluster	<p>View data as a cluster diagram, where information is clustered by a specific attribute.</p> <p>Use this option for identifying anomalies.</p>
Correlation	<p>View specific attributes as correlated data, with a specific attribute measured along the y axis. Time is measured along the x axis, based on the total time period selected using the time selector.</p> <p>Use this option for identifying the correlation between the behaviors of different attributes.</p>
Data Grid	<p>View data in the form of a grid. You can specify which columns are to be displayed.</p>
Donut	<p>Use this method to view the breakdown of an attribute.</p> <p>Use this option for identifying the top or bottom X targets by attribute.</p>
Gauge	<p>View a specific attribute for a circular or vertical gauge diagram.</p>
Histogram	<p>View the data as a histogram with a selected attribute measured along the x axis. The data can be grouped according to a third attribute.</p>
Label	<p>View a specific attribute as a label for your filtered targets.</p>

Visualization Option	Description
Line	<p>View the data as a line, with specific attributes measured along the x and y axes. Data can also be grouped according to a third attribute, while a fourth attribute can be used to determine the color of the line.</p> <p>Use this option for the following analytics tasks:</p> <ul style="list-style-type: none"><li>• Trending and forecasting</li><li>• Baselining</li><li>• Identifying the top or bottom X targets by attribute</li></ul>
Pareto Chart	<p>View data as a bar chart where the values are arranged in ascending or descending order. You can choose which attributes are measured along the x and y axes.</p>
Scatter Chart	<p>View data as a scatter chart with specific attributes measured along the x and y axes.</p> <p>Use this option for detecting anomalies.</p>
Table	<p>View the data as a table, with columns for each of the specified attributes. You can also choose to group the data by an attribute.</p> <p>Use this option for identifying the top or bottom X targets by attribute.</p>
Topology	<p>View data as a topology diagram.</p>
Treemap	<p>View the data as a map, showing the targets as cells, where the color of the cells is determined by a specific attribute.</p> <p>Use this option for identifying the top or bottom X targets by attribute.</p>

## Save Your Analysis

When you create an analysis (by filtering data and selecting necessary visualization options) that you want to retain to look at in the future, or that you want to share with others, then you can save it as a widget. Widgets can be used later to create dashboards.

When you save your analysis, you're saving:

- **Filters:** These are specific values or ranges for values that you select for attributes. Based on the values that you select, data is filtered out from the search results, for example, a specific target type or target name.
- **Visualization options:** This is the way to portray your data, for example, histogram, or treemap.
- **Analysis:** This is the mode of analysis that you select, for example, trending and forecasting.

## Build Dashboards

Dashboards are a collection of widgets on a single page that give you quick access to multiple data analyses.

A dashboard may contain one or more widgets. The dashboards appear on the Dashboards page of Oracle IT Analytics.

You can create and delete dashboards. If necessary, you can also edit dashboards by:

- Adding and deleting widgets
- Editing existing widgets on the dashboard
- Changing the layout

Exadata Database Machine entity type (on-premise Exadata) has an exclusive Exadata Health Dashboard. You can create similar out-of-the box Exadata Health Dashboards for ExaCS and ExaCC systems (on-cloud Exadata) to monitor your on cloud Exadata systems' health.

A dashboard can have widgets created from multiple data sources. For more information about using dashboards, see *Work with Dashboards* in *Working with Oracle Management Cloud*.

# A

## Data Collection and Storage

There are multiple metrics supported for collecting and storing data.

Oracle Oracle IT Analytics supports the following entities, only some of which are discussed in this appendix:

- Oracle Exadata
  - Oracle Exadata Database Machine
  - Oracle Exadata Rack
  - Oracle Engineered System ILOM Server
  - Oracle Engineered System PDU
  - Oracle Infiniband Network
  - [Oracle Exadata Storage Server](#)
  - Oracle Exadata Storage Server Grid
  - [Oracle Engineered System Cisco Switch](#)
  - [Oracle InfiniBand Switch](#)
- Oracle Database
  - [Oracle Database](#)
  - [Oracle Listener](#)
  - [Oracle Automatic Storage Management](#)
  - [Oracle Clusterware](#)
  - [Oracle Cluster Node](#)
- [Oracle as Host](#)
  - [Linux Host](#)
  - [Windows Host](#)
  - [Solaris Host](#)
  - [AIX Host](#)
  - [HP-UX Host](#)
- Oracle Hadoop
  - Oracle Hadoop Cluster
  - [Oracle Hadoop HDFS](#)
  - [Oracle Hadoop Yarn](#)
  - [Oracle Hadoop NameNode](#)
- [Oracle SOA Infrastructure and Service Bus](#)
- Oracle Siebel

- Oracle Siebel Component
- Oracle Siebel Enterprise
- Oracle Siebel Server
- Storage Manager
- Oracle Access Manager
- Oracle Internet Directory (OID)
- Oracle Home
- Oracle HTTP Server
- Oracle Traffic Director Configuration
- Oracle WebLogic Suite
  - Oracle WebLogic Cluster
  - WebLogic Domain
  - Oracle WebLogic Server
- Tomcat

## Linux Host

Here's a list of metrics, according to metric group.

- Capacity
- CPU
- Disk Activity Summary
- FileSystem
- FileSystem Usage
- Hardware Configuration
- Log Condition Match
- Memory
- Network Activity Summary
- Network Configuration
- Paging Activity
- Product
- Response

### Capacity

The following table lists the metrics and their descriptions.

Metric	Description
Total CPUs	Total number of CPUs available
Total Memory	Total amount of usable physical RAM available to the system in GB

## CPU

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
CPU in User Mode	Percentage of CPU time spent in User mode
CPU in System Mode	Percentage of CPU time spent in System mode
CPU Usage	Usage of all CPUs during the interval
CPU Utilization	Percentage of CPU utilized
CPU Load (1 minute)	Load on all CPUs in the last 1 minute
CPU Load (5 minutes)	Load on all CPUs in the last 5 minutes
CPU Load (15 minutes)	Load on all CPUs in the last 15 minutes
Interval	Interval of metric collection

## Disk Activity Summary

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
All Disks I/O Request Rate	Rate of read and write operations on all disks
Interval	Interval of metric collection

## FileSystem

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Mount Point	Mount point on the file system
Filesystem Name	Name of the file system
Filesystem Size	Size of the file system in GB

## FileSystem Usage

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Mount Point	Name of the system
Filesystem Usage	Space used in the file system in GB
Filesystem Space Available	Percentage of space available on the file system
Interval	Interval of metric collection

## Hardware Configuration

The following table lists the metric and its descriptions.

<b>Metric</b>	<b>Description</b>
CPU Architecture	Processor architecture used by the platform

### Log Condition Match

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Label Name	Log labels or tags, based on log condition matches in Log Processor
Time Window	Time window is which the tag or label count is summarized
odsUID	ODS unique ID, to distinguish the label or time window count map generated through different log processes
Count of Occurrence	Label or tag count per time window for the target or tenant

### Memory

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Memory Usage	Memory used in GB
Memory Utilization	Percentage of memory that's utilized
Memory Load	Load on memory in GB
Logical Memory Usage	Memory used in GB excluding buffers and cache
Logical Memory Utilization	Percentage of logical memory that's utilized
Interval	Interval of metric collection

### Network Activity Summary

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
All Network Interfaces Read Rate	Rate at which bytes are received across all network interfaces
All Network Interfaces Write Rate	Rate at which bytes are sent across all network interfaces
All Network Interface I/O Rate	Rate at which bytes are sent and received across all network interfaces
Interval	Interval of metric collection

### Network Configuration

The following table lists the metrics and their descriptions.

Metric	Description
Interface Name	Name of the network interface card
IP Address	IP address of the network card
MAC Address	MAC address of the network interface card

### Paging Activity

The following table lists the metrics and their descriptions.

Metric	Description
Pages Paged-in Rate	Disk read rate to resolve major page faults
Pages Paged-out Rate	Disk write rate to resolve major page faults
Interval	Interval of metric collection

### Product

The following table lists the metrics and their descriptions.

Metric	Description
Vendor	Vendor of the product
Name	Name of the product
Version	Version of the product

### Response

The following table lists the metric and its descriptions.

Metric	Description
Status	Status of the target

## Oracle Access Manager

Here's a list of metrics, according to metric group.

- [Authentication Aggregates](#)
- [Authorization Aggregates](#)
- [LDAP Operation Aggregates](#)
- [Response](#)

### Authentication Aggregates

The following table lists the metrics and their descriptions.

Metric	Description
Average Authentication Latency	Average authentication latency

Metric	Description
Authentication Requests	Authentication requests per second
Authentication Success Rate	Authentication success rate

### Authorization Aggregates

The following table lists the metrics and their descriptions.

Metric	Description
Average Authorization Latency	Average authorization latency
Authorization Requests	Authorization requests per second
Authorization Success Rate	Authorization success rate

### LDAP Operation Aggregates

The following table lists the metrics and their descriptions.

Metric	Description
Operation Name	Operation name
Average LDAP Operation Latency	Average Lightweight Directory Access Protocol (LDAP) operation latency (in milliseconds)
LDAP Success Rate	Lightweight Directory Access Protocol (LDAP) success rate percentage
LDAP Operations	Lightweight Directory Access Protocol (LDAP) operations per second

### Response

The following table lists the metric and its descriptions.

Metric	Description
Status	Status of Oracle Access Manager

## Oracle Automatic Storage Management

Here's a list of metrics for the entity Oracle Automatic Storage Management (ASM), according to metric group.

- [ASM Disks](#)
- [ASM Sparse Disk](#)
- [Database Disk Group Performance](#)
- [Database Disk Group Usage](#)
- [Disk Group Usage](#)
- [Disk Groups](#)
- [Sparse Database Disk Group Usage](#)

- [Sparse Disk Group Usage](#)
- [Volumes](#)

### ASM Disks

The following table lists the metrics and their descriptions.

Metric	Description
Disk Name	Name of the ASM Disk
Disk Group Name	Name of the ASM disk group to which the disk belongs
Failure Group Name	Name of the ASM failure group to which the disk belongs
Disk Path	Path of the disk service
Library	Software library that discovered the disk
Cell IP	IP address of the cell
Failure Group Type	Type of the failure group to indicate a regular or quorum disk
Redundancy of the disk	External redundancy of the disk
Total Size	Total size of the disk in GB

### ASM Sparse Disk

The following table lists the metrics and their descriptions.

Metric	Description
Disk Name	Name of the ASM Disk
Disk Group Name	Name of the ASM disk group to which the disk belongs
Physical Size	Physical size of the disk in GB

### Database Disk Group Performance

The following table lists the metrics and their descriptions.

Metric	Description
Instance ID	ASM Instance ID
Disk Group Name	Name of the ASM disk group
Database Name	Name of the database client
Response Time	Time taken by the disk group to process an IO request
Throughput	Total I/O throughput of the disk group for the database
IO Per Second	Total I/O processed per second by the disk group for the database
Size Per IO	Total size in MB processed per I/O by the disk group for the database

### Database Disk Group Usage

The following table lists the metrics and their descriptions.

Metric	Description
Disk Group Name	Name of the ASM disk group
Database Name	Name of the database client
Physical Allocated Size	Allocated size of the database in the disk group in MB

### Disk Group Usage

The following table lists the metrics and their descriptions.

Metric	Description
Disk Group Name	Name of the ASM disk group
Free Size	Free size of the disk group in GB
Usable Free Size	Usable free size of the disk group in GB
Disk Group Used Percent	Percentage of space used in a disk group
Usable Total Size	Usable total size, subject to redundancy, of the disk group in GB
Used Percent of Safely Usable	Percentage of safely usable space used in a disk group

### Disk Groups

The following table lists the metrics and their descriptions.

Metric	Description
Disk Group Name	Name of the ASM disk group
Redundancy	Redundancy type of the ASM disk group
Disks Count	Number of member disks in the disk group
Allocation Unit Size	Size in MB of one contiguous block in the disk group
Contains Voting Files	Indicates whether the disk contains voting files
Total Size	Total size of the disk group (in GB)

### Sparse Database Disk Group Usage

The following table lists the metrics for Sparse Disk Group Usage, and their descriptions.

Metric	Description
Disk Group Name	Name of the ASM disk group
Database Name	Name of the database client
Physical Allocated Size	Physical allocated size of the database in the disk group in MB

### Sparse Disk Group Usage

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Disk Group Name	Name of the ASM disk group
Physical Total Size	Total physical size of the disk group in GB
Physical Free Size	Free physical size of the disk group in GB
Physical Used Percent	Percentage of disk group physical space used

### Volumes

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Disk Group Name	Name of the ASM disk
Volume Device	OS path of the volume device
Volume Name	Name of the ASM volume
Usage	Usage of the volume
Redundancy	Redundancy of the volume
Total Size	Allocated size of the volume in GB
Disk Group Allocated Space	Total size of the volume in GB

## Oracle Cluster Node

Here's a list of metrics, according to metric group.

### Status of the Cluster Stack

The following table lists the metric and its description.

<b>Metric</b>	<b>Description</b>
Detailed Status of Cluster Stack	Detailed status of the cluster stack

## Oracle Clusterware

Here's a list of metrics, according to metric group.

- [Database, ASM, Listener, and Services Managed by the Cluster](#)
- [Detailed Status of the Cluster Stack](#)
- [Network Interfaces Used by the Cluster](#)
- [Scan Configuration](#)

### Database, ASM, Listener, and Services Managed by the Cluster

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Name of the Application	Name of the application
Type of Application	Type of the application
Cardinality	Cardinality
Degree	Degree
Node Name for a Node Application	Node name for a node application
Oracle Home for the Resource	Oracle home for the resource
Unique Parameters for the Application	Unique parameters for the application

### Detailed Status of the Cluster Stack

The following table lists the metric and its descriptions.

<b>Metric</b>	<b>Description</b>
Detailed Status of Cluster Stack	Detailed status of the cluster stack

### Network Interfaces Used by the Cluster

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Network Interface Name	Name of the network interface
Network Interface Subnet	Network interface subnet
Network Interface Type	Type of the network interface

### Scan Configuration

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Scan Name	Scan name configured for the cluster
Scan Port Configured	Whether the scan port has been configured
VIP Configured for Scan	When there scan VIP has been configured
Scan Listener Name	Scan Listener name
VIP Name	VIP name
VIP IP Address	VIP IP address
VIP Subnet	VIP subnet

## Oracle Database

Here's a list of parameters for the entity Oracle Database, according to metric group.

- [Capacity](#)
- [Character Set](#)

- CPU
- Creation Date
- Database Activity
- Efficiency
- I/O Bandwidth Usage
- I/O Per Second
- I/O Per Second Instance Level
- Instance CPU
- Instance Efficiency
- Instance Load
- Instance Memory Usage
- Instance Response
- Interconnect Usage Statistics
- Latency
- Limit Utilization
- Load
- Memory
- Memory Component Sizes
- Operating System Statistics
- Parameter
- Product
- Response
- SQL Statistics
- Storage Space Usage
- Storage Usage by Tablespace
- System Wait Classes
- Top SQL Executions
- Top Wait Events
- Total Memory Capacity

### Capacity

The following table lists the metrics and their descriptions.

Metric	Description
CPU Capacity	CPU capacity
Memory Capacity	Memory capacity (in GB)

### Character Set

The following table lists the metric and its descriptions.

Metric	Description
Character Set	Character set

### CPU

The following table lists the metrics and their descriptions.

Metric	Description
CPU Usage	CPU usage (in seconds)
CPU Utilization	Percentage of CPU utilization
CPU Load	Total CPU load
Interval	Interval for metric collection

### Creation Date

The following table lists the metric and its descriptions.

Metric	Description
Creation Data	Date of Creation

### Database Activity

The following table lists the metrics and their descriptions.

Metric	Description
DB Time	Amount of elapsed time (in seconds) spent on performing database user-level calls
Background CPU	Amount of CPU time (in seconds) consumed by database background processes
SQL Execution	Amount of elapsed time (in seconds) that SQL statements are executing
Instance Name	RDB instance
Host Name	Host name of the instance
CPU Time	CPU time (in seconds) of foreground sessions
Wait Time	Wait time (in seconds) of foreground sessions
I/O Time	I/O time (in seconds) of foreground sessions

### Efficiency

The following table lists the metrics and their descriptions.

Metric	Description
Sorts in Memory	Percentage of sorts in memory

<b>Metric</b>	<b>Description</b>
Soft Parses	Percentage of soft parses

### I/O Bandwidth Usage

The following table lists the metric and its description.

<b>Metric</b>	<b>Description</b>
I/O Bandwidth Usage	Total I/O bandwidth usage

### I/O Per Second

The following table lists the metric and its description.

<b>Metric</b>	<b>Description</b>
I/O Per Second	I/O per second

### I/O Per Second Instance Level

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	RDB Instance
Host Name	Host name of the instance
I/O Per Second	I/O per second for this instance
I/O Bandwidth Usage	I/O bandwidth usage during the last collection interval for this instance

### Instance CPU

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	RDB Instance
Host Name	Host name of the instance
CPU Usage	Usage of all CPUs for this instance, during the interval
CPU Utilization Interval	Percentage of CPU utilized in this interval
Interval	Interval of metric collection

### Instance Efficiency

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance ID	Instance ID

<b>Metric</b>	<b>Description</b>
PGA Cache Hit Percent	Percentage of PGA cache hits

### Instance Load

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	RDB Instance
Host Name	Host name of the instance
Logons	The number of logons during the last collection interval
Logons Rate	Number of logons per second during the last collection interval
Transactions	Number of transactions during the last collection interval
Transactions Rate	Number of transactions per second during the last collection interval
Executions	Number of executions during the last collection interval
Executions Rate	Number of executions per second during the last collection interval
Network I/O	Number of bytes sent and received over the network during the last collection interval
Network I/O Rate	Network I/O in bytes per second during the last collection interval
User Calls	User calls
Physical Read I/O Requests	Physical read I/O requests
Physical Read Bytes	Physical read bytes
Physical Write I/O Requests	Physical write I/O requests
Physical Write Bytes	Physical write bytes
Redo Bytes	Redo bytes
Redo Writes	Redo writes

### Instance Memory Usage

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	RDB instance
Host Name	Host name of the instance
Instance Memory Usage	Memory used in GB
Memory Utilization	Percentage of memory that's utilized in this instance
Interval	Interval of metric collection

## Instance Response

The following table lists the metrics and their descriptions.

Metrics	Description
Instance ID	ID of the instance
Status	Status of the instance

## Interconnect Usage Statistics

The following table lists the metrics and their descriptions.

Metric	Description
Instance Name	Database instance name
Host Name	Host name for instance
Bytes Sent	Number of bytes sent over interconnect
Bytes Received	Number of bytes received over interconnect

## Latency

The following table lists the metrics and their descriptions.

Metrics	Description
SQL Service Response Time	SQL server response time
Average Synchronous Single-Block Read Latency	Average synchronous single-block read latency
Average Wait Time for Log File Synchronization	Average wait time for log file synchronization
Average Wait Time for Sequential Read of Datafile	Average wait time for a sequential read of the data files

## Limit Utilization

The following table lists the metrics and their descriptions.

Metric	Description
Process Limit Utilization	Percentage of the process limit that has been utilized
Session Limit Utilization	Percentage of the session limit that has been utilized

## Load

The following table lists the metrics and their descriptions.

Metric	Description
Average Active Sessions	Average active sessions

<b>Metric</b>	<b>Description</b>
Logons	Total number of logons
Logons Rate	Number of logons per second
Transactions	Total number of transactions
Transactions Rate	Number of transactions per second
Executions	Total number of executions
Executions Rate	Number of executions per second
Network I/O	Total network I/O (in bytes)
Network I/O Rate	I/O per second

### Memory

The following table lists the metrics and their description.

<b>Metric</b>	<b>Description</b>
Memory Usage	Memory usage (in GB)
Memory Utilization	Percentage of memory utilization
Memory Load	Total memory load
Interval	Interval for metric collection

### Memory Component Sizes

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	Database instance name
Host Name	Host name for instance
Memory Type	Memory type
Memory Pool Name	Memory pool name
Allocated Size	Allocated memory size in bytes

### Operating System Statistics

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	RDB Instance
Host Name	Host name of the instance
Total CPUs	Total number of CPUs allocated or available in this instance
Total Memory	Total amount of usable physical RAM allocated or available in this instance

## Parameter

The following table lists the metrics and their descriptions.

Metric	Description
Instance ID	Instance ID
Container ID	Container ID
Parameter Name	Name of the parameter
Parameter Value	Value of the parameter

## Product

The following table lists the metrics and their descriptions.

Metric	Description
Vendor	Product vendor
Name	Product name
Version	Product version

## Response

The following table lists the metric and its descriptions.

Metric	Description
Status	Response status

## SQL Statistics

The following table lists the metrics and their descriptions.

Metric	Description
Disk Reads	Total number of disk reads
Buffer Gets	Total number of buffer gets
CPU Time	Total time (in milliseconds) spent on CPU
User I/O Wait	Total time (in milliseconds) spent waiting for events in the User I/O class
Cluster Wait	Total time (in milliseconds) spent waiting for events in the Cluster class
Application Wait	Total time (in milliseconds) spent waiting for events in the Application class
Concurrency Wait	Total time (in milliseconds) spend waiting for events in the Concurrency class
Direct Writes	Number of direct writes
Physical Read Requests	Number of physical read I/O requests issued by the SQL
Physical Read Bytes	Number of bytes read from disks by the SQL

<b>Metric</b>	<b>Description</b>
Physical Write Requests	Number of physical write I/O requests issued by the SQL
Physical Write Bytes	Number of bytes written to disks by the SQL
SQL ID	SQL ID
Plan Hash Value	Numerical representation of the SQL plan for the cursor
Executions	Number of times the SQL statement was executed
Elapsed Time	Total time (in milliseconds) used by this SQL for parsing, executing, and fetching

### Storage Space Usage

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Storage Space Allocated	Total allocated storage space (in GB)
Storage Space Utilization	Percentage of allocated storage space that has been utilized
Space Used	Total space used (in GB)

### Storage Usage by Tablespace

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Storage Tablespace Name	Tablespace name
Total Space Allocated	Total size of a tablespace's data files in GB
Space Utilization Percentage	Percent of storage space utilized
Space Used	Tablespace used in GB

### System Wait Classes

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Instance Name	Database instance name
Host Name	Host name of instance
Wait Class	Wait class
Time Waited	Total time waited in seconds by both foreground and background sessions
Time Waited Foreground	Total time waited in seconds by foreground sessions only

### Top SQL Executions

The following table lists the metrics and their descriptions.

Metric	Description
SQL ID	ID of the SQL query
CPU Time	CPU time used to process the query
Elapsed Time	Time elapsed before the query results were returned
Top SQL Text	Text of the top SQL queries
Execution Count	Execution count (in seconds)

### Top Wait Events

The following table lists the metrics and their descriptions.

Metric	Description
Wait Event Name	Name of the wait event
Wait Class Name	Name of the wait class
Total Waits	Total number of wait events
Total Wait Time	Total wait time
Average Wait Time	Average wait time
Total Foreground Waits	Total number of foreground waits
Total Foreground Wait Time	Total time of foreground waits
Average Foreground Wait Time	Average time of foreground waits

### Total Memory Capacity

The following table lists the metric and its descriptions.

Metric	Description
Host Memory	Total host memory (in GB)

## Oracle Engineered System Cisco Switch

Here's a list of metrics, according to the metric group.

- [Component Details](#)
- [Component State](#)
- [Network Management Interface](#)
- [Network Port Details](#)
- [Network Port Statistics](#)
- [Performance](#)
- [Summary Details](#)

### Component Details

The following table lists the metrics and their descriptions.

Metric	Description
Component Name	Name of the component
Component Manufacturer	Manufacturer of the component
Component Model	Model of the component

### Component State

The following table lists the metrics and their descriptions.

Metric	Description
Component Name	Name of the component
Component State	State of the component

### Network Management Interface

The following table lists the metrics and their descriptions.

Metric	Description
IP Address	IP address of the switch
Physical Address	Physical address of the switch

### Network Port Details

The following table lists the metrics and their descriptions.

Metric	Description
Port Name	Name of the port
Port Physical Address	Physical address of the port
Port Type	Type of the port
Peer Port Physical Address	Physical address of the peer port
Peer Port Type	Type of the peer port
Port Media Type	Port media type

### Network Port Statistics

The following table lists the metrics and their descriptions.

Metric	Description
Port Name	Name of the port
Port Status	Status of the port
Inbound Octets Rate	Inbound octets rate
Outbound Octets Rate	Outbound octets rate

<b>Metric</b>	<b>Description</b>
Total Octets Rate	Total octets rate

### Performance

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Ports Free	Unplugged ports
Ports Used	Plugged ports

### Summary Details

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Part Number	The part number of the switch
Serial Number	The serial number of the switch
Model	The model of the switch
Architecture	The architecture of the switch

## Oracle Exadata Storage Server

Here's a list of metrics, according to the metric group.

- [Exadata Capacity Metric](#)
- [Exadata Cell Configuration](#)
- [Exadata Cell Metric](#)
- [Exadata Diskgroup Capacity Metric](#)
- [Exadata Flash Cache IORM Database Metric](#)
- [Exadata Flash Cache Metric](#)
- [Exadata Flash Log Metric](#)
- [Exadata FlashDisk and HardDisk Metric](#)
- [Exadata Sparse Diskgroup Capacity Metric](#)

### Exadata Capacity Metric

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Disk Type	Disk type
Disk Size	Disk size
Usage	Usage

### Exadata Cell Configuration

The following table lists the metrics and their descriptions.

Metric	Description
Cell ID	Cell ID
Cell Version	Cell version
Release Version	Release version
Kernel Version	Kernel version
Make Model	Make model
Maximum Physical Disk IOPS	Maximum physical disk I/O operation per second
Maximum Flash Disk IOPS	Maximum flash disk I/O operation per second
Maximum Physical Disk MBPS	Maximum physical disk in MBs per second
Maximum Flash Disk MBPS	Maximum flash disk in MBs per second

### Exadata Cell Metric

The following table lists the metrics and their descriptions.

Metric	Description
Infiniband Network Received	The number of megabytes received by the InfiniBand interfaces per second.
Infiniband Network Sent	The number of megabytes transmitted by the InfiniBand interfaces per second.

### Exadata Diskgroup Capacity Metric

The following table lists the metrics and their descriptions.

Metric	Description
ASM Instance Name	ASM instance name
Diskgroup Name	Name of the ASM disk group
Diskgroup Size	ASM disk group size

### Exadata Flash Cache IORM Database Metric

The following table lists the metrics and their descriptions.

Metric	Description
Database Name	Exadata flash cache I/O resource management database metric
Space Used in Flash Cache	Space used in flash cache

### Exadata Flash Cache Metric

The following table lists the metrics and their descriptions.

Metric	Description
Space Used	Flash cache space used
Space Used for Keep Objects	Space used for keep objects

### Exadata Flash Log Metric

The following table lists the metrics and their descriptions.

Metric	Description
Effective Size	Effective flash log size
Flash Log Object Size	Flash log object Size

### Exadata FlashDisk and HardDisk Metric

The following table lists the metrics and their descriptions.

Metric	Description
Disk Type	Exadata FlashDisk and HardDisk metric
Small Read I/O per second	Small read I/O per second
Large Read I/O per sec	Large read I/O per second
Small Read Throughput	Small read throughput
Large Read Throughput	Large read throughput
Read Response Time	Average read response time
Write Response Time	Average write response time
IO Utilization	Average I/O utilization

### Exadata Sparse Diskgroup Capacity Metric

The following table lists the metrics and their descriptions.

Metric	Description
ASM Instance Name	ASM instance name
Diskgroup Name	Name of the ASM disk group
Diskgroup Size	ASM disk group size
Virtual Size	Size of the virtual disk group

## Oracle Hadoop HDFS

Here's a list of metrics, according to metric group.

- [Application Metrics](#)
- [Block Metrics](#)

### Application Metrics

The following table lists the metrics and their descriptions.

Metric	Description
Create Rate	Number of create operations
Delete Rate	Number of delete operations
Rename Rate	Number of rename operations
Sync Rate	Number of synchronization operations

### Block Metrics

The following table lists the metrics and their descriptions.

Metric	Description
Block Capacity	The block capacity of the NameNode
Block With Corrupt Replicas	Blocks with corrupt replicas
Excess Blocks	The total number of excess blocks
Missing Blocks	Missing blocks
Pending Deletion Blocks	The number of blocks pending deletion
Pending Replication Blocks	The number of blocks with replication pending
Scheduled Replication Blocks	The number of blocks with replication currently scheduled
Under Replicated Blocks	Under-replicated blocks
Expired Heartbeats	The number of expired heartbeats

## Oracle Hadoop NameNode

Here's a list of the metrics, according to metric group.

- [Application Metrics](#)
- [Block Metrics](#)
- [Memory Resource Metrics](#)
- [Read Write Metrics](#)
- [Response Metrics](#)

### Application Metrics

The following table lists the metrics and their description.

Metric	Description
Create Rate	Number of create operations
Delete Rate	Number of delete operations
Rename Rate	Number of rename operations
Syncs Rate	Number of synchronization operations

### Block Metrics

The following table lists the metrics and their description.

<b>Metric</b>	<b>Description</b>
Block Capacity	Block capacity if Hadoop Distributed File System target
Blocks Total	Total blocks of Hadoop Distributed File System target
Blocks With Corrupt Replicas	Number of blocks with corrupt replicas
Excess Blocks	Total number of excess blocks
Missing Blocks	Total number of missing blocks

### Memory Resource Metrics

The following table lists the metrics and their description.

<b>Metric</b>	<b>Definition</b>
DFS Capacity	Total configured Hadoop Distributed File System storage capacity
DFS Capacity Free	Free storage space
DFS Capacity Used	Storage space used by Hadoop Distributed File System files
DFS Capacity Used Non HDFS	Storage space used by non-Hadoop Distributed File System files
JVM Blocked Threads	Blocked threads
JVM Heap Used	Total amount of used heap memory
JVM Max Memory	Maximum allowed memory
JVM New Threads	New threads
JVM Runnable Threads	Runnable threads
JVM Terminated Threads	Terminated threads
JVM Times Waiting Threads	Timed waiting threads
JVM Waiting Threads	Waiting threads
Resident Memory	Resident memory used
Virtual Memory	Virtual memory used

### Read Write Metrics

The following table lists the metric and their description.

<b>Metric</b>	<b>Description</b>
Time Since Last Checkpoint	The time since the last file system checkpoint
RPC Call Queue Length	RPC call queue length
RPC Open Connections	Open RPC connections
Transactions Since Last Checkpoint	Transactions since last log checkpoint
Transactions Since Last Log Roll	Transactions since last log roll

## Response Metrics

The following table lists the metrics and its description.

Metric	Description
Response	Status of the target

## Oracle Hadoop Yarn

Here's a list of metrics, according to metric group.

- [Alert Metrics](#)
- [Application Counter Metrics](#)

### Alert Metrics

The following table lists the metric and its description.

Metric	Description
Alert Metrics	Alert events generated in the collection time interval

### Application Counter Metrics

The following table lists the metrics and their descriptions.

Metric	Description
Application Running	Application running on the service target during the collection time interval
Applications Accepted	Applications accepted on the service target during the collection time interval
Applications Succeeded	Applications successfully completed on the service target during the collection time interval
Applications Failed	Applications failed on the service target during the collection time interval

## Oracle Home

Here's a list of metrics, according to metric group.

### Components

The following table lists the metrics and their descriptions.

Metric	Description
Component Name	Name of the component
Base Version	Base version of the component

Metric	Description
Current Version	Current version (affected by patches)
Install Time	Installation time of the patch
Is it a top level component	Whether the component is top-level
External Name	External name of the component
Description	Description of the component

## Oracle as Host

Here's a list of metrics, according to metric group.

- [CPU](#)
- [File System](#)
- [File System Usage](#)
- [Hardware Configuration](#)
- [Memory](#)
- [Network Configuration](#)
- [Resource Allocation](#)

### CPU

The following table lists metrics and their descriptions.

Metric	Description
CPU Utilization (User Mode)	Percentage of CPU time spent in User mode across all CPUs
CPU Utilization (System Mode)	Percentage of CPU time spent in System mode across all CPUs
CPU Usage	Total CPU time spent across all CPUs
CPU Utilization	Percentage of total CPU time spent across all CPUs
CPU Load (1 minute)	Average number of processes in memory and subject to be run in the last 1 minute
CPU Load (5 minutes)	Average number of processes in memory and subject to be run in the last 5 minutes
CPU Load (15 minutes)	Average number of processes in memory and subject to be run in the last 15 minutes
Interval	Interval of metric collection

### File System

The following table lists metrics and their descriptions.

Metric	Description
Mount Point	Mount point of the file system

<b>Metric</b>	<b>Description</b>
File System Name	Name of the file system
File System Size	Size of the file system in GB

### File System Usage

The following table lists metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Mount Point	Mount point of the file system
File System Usage	Space used in the file system in GB
File System Space Available	Space available on the file system in %
Interval	Interval of metric collection

### Hardware Configuration

The following table lists the metric and its descriptions.

<b>Metric</b>	<b>Description</b>
CPU Architecture	Processor architecture used by the platform

### Memory

The following table lists metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Memory Usage	Memory used in GB
Memory Utilization	Percentage of memory that's utilized
Memory Load	Sum of memory used and page faults generated in GB
Logical Memory Usage	Memory used in GB excluding buffers and cache
Logical Memory Utilization	Percentage of logical memory that's utilized
Interval	Interval of metric collection

### Network Configuration

The following table lists metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Pages Paged-in Rate	Disk read rate to resolve major page faults
Pages Paged-out Rate	Disk write rate to resolve major page faults
Interval	Interval of metric collection

## Resource Allocation

The following table lists metrics and their descriptions.

Metric	Description
Total CPUs	Total number of CPUs available
Total Memory	Total amount of usable physical RAM available to the system (in GB)

# Oracle HTTP Server

Here's a list of metrics, according to metric group.

- [Capacity](#)
- [CPU](#)
- [Datasource Metric](#)
- [Listen Addresses](#)
- [Memory](#)
- [Modules](#)
- [Product](#)
- [Response](#)
- [Routing](#)
- [Servers](#)
- [Virtual Hosts](#)
- [Web Application Component](#)
- [Web Requests](#)

## Capacity

The following table lists the metrics and their descriptions.

Metric	Description
CPU Capacity	Total number of CPUs available
Memory Capacity	Total amount of usable physical RAM available to the system
Name	Name
Host	Host
Process	Process
Server Name	Name of the server
Total Declined Requests	Total number of declined requests

## CPU

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
CPU Usage	Total CPU usage during the interval
CPU Utilization	Percentage of CPU utilized
CPU Load	Average load on all CPUs in the last 1 minute
Interval	Interval of metric collection

### Datasource Metric

The following table lists the metric and its descriptions.

<b>Metric</b>	<b>Description</b>
Connections in Use	Number of connections in use

### Listen Addresses

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
IP	IP address
Port	Port number
Protocol	Protocol followed.

### Memory

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Memory Usage	Memory used in GB
Memory Utilization	Percentage of memory that's utilized
Memory Load	Load on memory
Interval	Interval of metric collection

### Modules

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Total Handled Requests	Number of requests handled
Total Handling Time	Total time taken to handle requests
Average Handling Time	Average time to handle requests
Minimum Handling Time	Minimum time taken to handle requests
Maximum Handling Time	Maximum time taken to handle requests
Handling Throughput	Handling throughput
Current Handling Time	Current time taken to handle requests

## Product

The following table lists the metrics and their descriptions.

Metric	Description
Vendor	Vendor of the product
Name	Name of the product
Version	Version of the product

## Response

The following table lists the metric and its description.

Metric	Description
Status	Status of the target

## Routing

The following table lists the metrics and their descriptions.

Metric	Description
Virtual Host	Name of virtual host
App Path	Path of the application
WebLogic Host	Oracle WebLogic Server host
WebLogic Port	Oracle WebLogic Server port
WebLogic Cluster	Oracle WebLogic Server cluster
WebLogic IOT Timeout Time	Time of the Oracle WebLogic Server IoT timeout
ssl	Secure Sockets Layer
connectRetry Time	Time take to retry the connection
fileCaching	File caching

## Servers

The following table lists the metrics and their descriptions.

Metric	Description
Name	Name of server
Host	Host
Process	Process
Total Get Requests	Total number of Get requests
Total Post Requests	Total number of Post requests
Total Errors	Total number of errors
Time of Last Configuration File Change	Time when the last configuration file was changed

<b>Metric</b>	<b>Description</b>
Total Response Data Size	Total size of response data
Total Number of Child Process	Total number of child processes
Total Number of Idle Process	Total number of idle processes
Total Number of Busy Process	Total number of busy processes
Total HTTP Connections	Total HTTP connections
Total Time HTTP Connections have been Open	Total time for which the HTTP connections have been open
Active HTTP Connections	Number of active HTTP connections
Average Time an HTTP Connection has been Open	Average time for which an HTTP connection has been open
Minimum Time an HTTP Connection has been Open	Minimum time for which an HTTP connection has been open
Maximum Time an HTTP Connection has been Open	Maximum time for which an HTTP connection has been open
Total HTTP Requests Processed	Total number of HTTP requests that have been processed
Total HTTP Processing Time	Total time taken to process HTTP requests
Active Requests	Total number of active requests
Average HTTP Request Processing Time	Average HTTP request processing time
Minimum HTTP Request Processing Time	Minimum time to process an HTTP request
Maximum HTTP Request Processing Time	Maximum time to process an HTTP request
Total Invocations of a Module Handler	Total calls of a module handler
Total Module Handler Time	Total module handler time
Active Handlers	Total number of active handlers
Average Module Handler Time	Average module handler time
Minimum Module Handler Time	Minimum module handler time
Maximum Module Handler Time	Maximum module handler time
Loaded Modules	Loaded modules
Total Internal Redirects	Total number of internal redirects
Current Request Throughput	Current request throughput
Current Request Processing Time	Current request processing throughput
Error Rate	Error rate
Connection Duration	Total connection duration
Size of Response Data Processed	Size of the processed response data
Response Data Throughput	Response data throughput
Total Number of Busy Threads	Total number of busy threads
Total Number of Threads	Total number of threads
Percentage of Busy Threads	Percentage of busy threads
Total Number of Busy Threads	Total number of busy threads
Percentage of Busy Threads	Percentage of busy threads
Total Number of Idle Threads	Total number of idle threads

<b>Metric</b>	<b>Description</b>
Percentage of Busy Processes	Percentage of busy processes

### Virtual Hosts

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
IP	IP address of the virtual host
Port	Port number
Server Name	Server name
Server Type	Server type
Name	Name of the server
Host	Host
Process	Process
Active Requests	Number of active requests
Average Processing Time	Average time taken to process requests
Completed Requests	Number of completed requests
Maximum Request Processing Time	Maximum time taken to process a request
Minimum Request Processing Time	Minimum time taken to process a request
Total Request Processing Time	Total time taken to process requests
Total Response Data Size	Total size of the response data
Virtual Host Type	Virtual Host type
OHS Server	Oracle HTTP Server server
Virtual Host Set	Virtual host set
Current Request Throughput	Total current request throughput
Current Processing Time	Total current processing time
Current Response Data Size	Total current response data size
Current Response Throughput	Total current response throughput

### Web Application Component

The following table lists the metric and its description.

<b>Metric</b>	<b>Description</b>
WebApp Open Sessions	Current total number of open sessions for all the web applications on this web application server

### Web Requests

The following table lists the metrics and their descriptions.

Metric	Description
Web Request Processing Time	Average processing time per web request since the last collection
Request Rate	Number of web requests per minute since the last collection

## Oracle InfiniBand Switch

Here's a list of metrics, according to the metric group.

- [Subnet Manager Details](#)
- [Subnet Manager State](#)

### Subnet Manager Details

The following table lists the metrics and their descriptions.

Metric	Description
Subnet Manager IP Address	IP address of the subnet manager
Subnet Manager LID	Subnet manager hosting HCA LID
Subnet Manager GUID	Subnet manager hosting HCA GUID

### Subnet Manager State

The following table lists the metrics and their descriptions.

Metric	Description
Subnet Manager State	Subnet manager state
Subnet Manager Priority	Subnet manager priority
Subnet Manager Activity Count	Subnet manager activity count

## Oracle Listener

Here's a list of metrics, according to metric group.

- [Listener EndPoints](#)
- [Listener Load](#)

### Listener EndPoints

The following table lists the metrics and their descriptions.

Metric	Description
Server	Hosted machine name
Protocol	Protocol followed (for example, IPC, TCP, and TCPS)

Metric	Description
Protocol Value	Key, in the case of IPC, and Port in the case of TCP and TCPS.

### Listener Load

The following table lists the metrics and their descriptions.

Metric	Description
Connections Established (per minute)	Rate of connections established
Connections Refused (per minute)	Rate of connections refused

## Oracle Siebel Component

Here's a list of metrics, according to metric group.

- [Siebel Component Overview](#)
- [Siebel Component Processes](#)

### Siebel Component Overview

The following table lists the metrics and their descriptions.

Metric	Description
Average Object Manager Response Time	Average Object Manager response time (in milliseconds)
Average number of requests per Object Manager session	Average number of requests per Object Manager session
Average size of request messages	Average size of request messages (in bytes)
Average size of reply messages	Average size of reply messages (in bytes)
Average Connect time for Object Manager sessions	Average connect time for Object Manager sessions (in seconds)
Number Component Process Failures	Number of component process failures

### Siebel Component Processes

The following table lists the metric and its description.

Metric	Description
Process ID	ID of the Siebel component process

## Oracle Siebel Enterprise

Here's a list of metrics, according to metric group.

### Siebel Enterprise Database Configuration

The following table lists the metrics and their descriptions.

Metric	Description
Siebel Database DSN Information	Information about the Siebel database DSN
Host DSN Information	Information about the host DSN
Siebel Database Host Name	Name of the Siebel database host
Siebel Database SID	Siebel database SID

## Oracle Siebel Server

Here's a list of metrics, according to metric group.

- [Siebel Server Overview](#)
- [Siebel Server Processes](#)

### Siebel Server Overview

The following table lists the metrics and their descriptions.

Metric	Description
Average Connect Time for Object Manager Sessions	Average time to connect for object manager sessions (in seconds)
Average Object Manager Response Time	Average response time of the object manager (in milliseconds)
Average Size of Request Messages	Average size of request messages (in bytes)
Average Size of Reply Messages	Average size of reply messages (in bytes)
Average Number of Requests per Object Manager Session	Average number of requests per Object Manager session
Total CPU time for component tasks	Total CPU time for completing component tasks (in seconds)
Number of Retries due to Deadlock Rollbacks	Number of retries because of deadlock rollbacks
Number of Component Process Failures	Total number of component process failures

### Siebel Server Processes

The following table lists the metric and its description.

Metric	Description
Process ID	ID of the Siebel server process

# Oracle SOA Infrastructure

Here's a list of metrics, according to metric group.

- [Service Engine](#)
- [SOA Composite Application](#)
- [SOA Infrastructure](#)

## Service Engine

The following table lists the metrics and their descriptions.

Metric	Description
Engine Name	Name of the SOA service engine
Throughput	Messages processed by the service engine
Errors	Errors generated by the service engine
Error Rate	Percentage of errors out of total messages processed by the service engine

## SOA Composite Application

The following table lists the metrics and their descriptions.

Metric	Description
Revision	Revision of the SOA composite application
Name	Name of the SOA composite application
SOA Partition	Name of the SOA partition
System Faults	System faults generated by the SOA composite application
Business Faults	Business faults generated by the SOA composite application
Policy Violations	Policy violations generated by the SOA composite application
Error Rate	Percentage of errors out of total messages processed by the SOA composite application
Synchronous Response Time	Average time for processing synchronous transactions
Asynchronous Response Time	Average time for posting asynchronous transactions
Errors	Errors generated by the SOA composite application
Throughput	Messages processed by the SOA composite application

## SOA Infrastructure

The following table lists the metrics and their descriptions.

Metric	Description
Throughput	Messages processed across all the SOA composite applications
Errors	Errors generated across all the SOA composite applications

## Oracle Traffic Director Configuration

Here's a list of metrics, according to metric group.

- [Failover Groups](#)
- [Listeners](#)
- [Origin Server](#)
- [Origin Server Pools](#)
- [Performance Statistics for all Instances](#)
- [Virtual Server](#)

### Failover Groups

The following table lists the metrics and their descriptions.

Metric	Description
Failover Virtual IP	The virtual IP address for the failover group
Network Prefix	The subnet mask in terms of the number of bits that is used to identify the network
Primary Node	Name of the primary node
Backup Node	Name of the backup node
Primary Node NIC	Network interface card (NIC) for the primary node
Backup Node NIC	Network interface card (NIC) for the backup node
Primary Instance	Name of the primary Oracle Traffic Director instance
Backup Instance	Name of the backup Oracle Traffic Director instance
Primary Instance NIC	Network interface card (NIC) for the primary Oracle Traffic Director instance
Backup Instance NIC	Network interface card (NIC) for the backup Oracle Traffic Director instance
Router ID	Router ID for the failover group

### Listeners

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Listener	Listener name
Listener Enabled	Whether the listener is enabled
Listener IP Address	Server listening IP address
Listener Port	Server listening port
Server Name	Server name used in server-generated URLs
Default Virtual Server / TCP Proxy	Default virtual server or TCP proxy for the listener
Acceptor Threads	Number of threads dedicated to accepting connections received by this listener
Protocol Family (Default/inet/inet6/inet-sdp)	The network address family for the listener. 'Default' implies that the family will be derived from the IP address.
Listen Queue Size	Maximum size of the operating-system listen queue backlog
Protocol Mismatch Handled	Indicates whether the server redirects clients that use wrong protocol
Blocking IO Enabled	Whether the server is using blocking I/O

### Origin Server

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Origin Server	Name of the origin server
Server Pool	Name of the origin server pool
Status	Whether the back-end server is online, offline, or unreferenced
Active Connections	Number of active connections across all Oracle Traffic Director instances
Idle Connections	Number of idle connections across all Oracle Traffic Director instances
Request Rate	Number of requests processed per second across all Oracle Traffic Director instances
Sticky Request Rate	Number of sticky requests processed per second across all Oracle Traffic Director instances
Error Rate	Number of errors returned per second across all Oracle Traffic Director instances
Data Transmission Rate	Amount of data transmitted per second across all Oracle Traffic Director instances
Data Receive Rate	Amount of data received per second across all Oracle Traffic Director instances

### Origin Server Pools

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Server Pool	Server pool
Origin Server Host	Origin server host
Origin Server Port	Origin server listening port
Enabled	Whether the origin server is enabled
Backup Server	Whether the origin server is the backup server
Received Requests Ratio	Ratio of requests (with respect to other origin servers) received by this origin server
Maximum Connections	Maximum number of connections that the origin server can handle concurrently
Ramp Up Time	Time for Oracle Traffic Director to ramp up the request sending rate to the full capacity of this origin server

### Performance Statistics for all Instances

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Request Rate	Number of requests processed per second across all Oracle Traffic Director instances
Request Processing Time	Average response time for the requests across all Oracle Traffic Director instances
Error Rate	Number of errors returned per second across all Oracle Traffic Director instances
Data Transmission Rate	Size of data transmitted per second across all Oracle Traffic Director instances
Data Receive Rate	Size of data received per second across all Oracle Traffic Director instances

### Virtual Server

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Virtual Server	Virtual server name
Request Rate	Number of requests processed per second across all Oracle Traffic Director instances
Error Rate	Number of errors returned per second across all Oracle Traffic Director instances
Data Receive Rate	Amount of data received per second across all Oracle Traffic Director instances
Data Transmission Rate	Amount of data transmitted per second across all Oracle Traffic Director instances

## Oracle WebLogic Server

Here's a list of metrics, according to metric group.

- [Data Source Metrics](#)
- [Work Manager Metrics](#)

### Data Source Metrics

The following table lists the metrics and their descriptions.

Metric	Description
JDBC Connection Creation Rate	Data source connection create operation throughput
JDBC Connection Leak Rate	Rate of connection leaks
JDBC Connection Create Delay	Average latency creating a data source connection
JDBC Connection Request Rate	Rate of incoming requests for Java Database Connectivity (JDBC) connections
JDBC Successful Connection Request Wait Rate	Rate of requests that had to wait before getting a Java Database Connectivity (JDBC) connection and eventually succeeded in getting a connection
JDBC Connection Requests Waiting	Number of Java Database Connectivity (JDBC) connection requests waiting for a database connection
JDBC Connection Request Wait Rate	Rate of requests that had to wait before getting a Java Database Connectivity (JDBC) connection including those that eventually got a connection and those that didn't get a connection

### Work Manager Metrics

The following table lists the metric and its description.

Metric	Description
Work Manager Stuck Threads	Count of work manager stuck threads

## Service Bus

Here's a list of metrics, according to metric group.

- [Service Bus Metrics](#)
- [Service Bus Service Details Statistics Metrics](#)

### Service Bus Metrics

The following table lists the metrics and their descriptions.

Metric	Description
Throughput	Messages processed across all the proxy services
Errors	Errors generated across all the proxy services

### Service Bus Service Details Statistics Metrics

The following table lists the metrics and their descriptions.

Metric	Description
Project	Name of the service bus project
Service Name	Name of the service
Service Type	Type of the service (for example, Proxy Service, Pipeline, and Split-Join)
Throughput	Messages processed by the service
Errors	Errors generated by the service
WS Security Errors	Web service security errors generated by the service
Average Response Time	Average time for processing the transaction by the service
Error Rate	Percentage of errors out of total messages processed by the service

## Storage Manager

Here's a list of the metrics, according to metric group.

- [Storage Capacity](#)
- [Storage Utilization](#)

### Storage Capacity

The following table lists the metrics and their descriptions.

Metric	Description
Total Size	Total size of the managed storage
Physical Size	Physical size of the managed storage

### Storage Utilization

The following table lists the metrics and their descriptions.

Metric	Description
Used Size	Used size of the managed storage
Physical Used Size	Physical used size of the managed storage
Storage Used	Percentage of space used in the storage

Metric	Description
Physical Storage Used	Percentage of physical space used in the storage

## Windows Host

Here's a list of metrics, according to metric group.

- [Capacity](#)
- [CPU](#)
- [Disk Activity Summary](#)
- [File System](#)
- [File System Usage](#)
- [Hardware Configuration](#)
- [Log Condition Match](#)
- [Memory](#)
- [Network Activity Summary](#)
- [Network Configuration](#)
- [Paging Activity](#)
- [Product](#)
- [Response](#)

### Capacity

The following table lists the metrics and their descriptions.

Metric	Description
CPU Capacity	Total number of CPUs available
Memory Capacity	Total amount of usable physical RAM available to the system in GB

### CPU

The following table lists the metrics and their descriptions.

Metric	Description
CPU in User Mode	Percentage of CPU time spent in User mode
CPU is Sysytem Mode	Percentage of CPU time spent in System mode
CPU Usage	Usage across all CPUs during the interval
CPU Utilization	Percentage of CPU that was utilized
Interval	Interval of metric collection

## Disk Activity Summary

The following table lists the metrics and their descriptions.

Metric	Description
All Disks I/O Request Rate	Rate of read and write operations on all disks
Interval	Interval of metric collection

## File System

The following table lists the metrics and their descriptions.

Metric	Description
Mount Point	Mount Point of the file system
Filesystem Name	Name of the file system
Filesystem Size	Size of the file system in GB

## File System Usage

The following table lists the metrics and its descriptions.

Metric	Description
Mount Point	Name of the file system
Filesystem Usage	Space used in the file system in GB
Filesystem Space Available	Percentage of space available on the file system
Interval	Interval of metric collection

## Hardware Configuration

The following table lists the metrics and their descriptions.

Metric	Description
CPU Architecture	Processor architecture used by the platform

## Log Condition Match

The following table lists the metrics and their descriptions.

Metric	Description
Label Name	Log labels or tags, based on log condition matches in log processor
timeWindow	The time window in which the tag or label count is summarized
odsUID	ODS unique ID to distinguish the label or the time window count map, generated through different log processors

<b>Metric</b>	<b>Description</b>
Count of Occurrence	Label or tag count per time window for the target or tenant

### Memory

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Memory Usage	Memory used in GB
Memory Utilization	Percentage of memory that's utilized
Memory Load	Load on memory in GB
Interval	Interval of metric collection

### Network Activity Summary

The following table lists the metric and their description.

<b>Metric</b>	<b>Description</b>
All Network Interfaces Read Rate	Rate at which bytes are received on all network interfaces
All Network Interfaces Write Rate	Rate at which bytes are sent on all network interfaces
All Network Interfaces I/O Rate	Rate at which bytes are sent and received on all network interfaces
Interval	Interval of metric collection

### Network Configuration

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Interface Name	Name of the network interface card
IP Address	IP address of the network card
MAC Address	MAC address of the network interface card

### Paging Activity

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Pages Paged-in Rate	Disk read rate to resolve major page faults
Pages Paged-out Rate	Disk write rate to resolve major page faults
Interval	Interval of metric collection

---

## Product

The following table lists the metrics and their descriptions.

<b>Metric</b>	<b>Description</b>
Vendor	Vendor of the product
Name	Name of the product
Version	Version of the product

## Response

The following table lists the metric and its description.

<b>Metric</b>	<b>Description</b>
Stats	Status of the target