

Autonomous Health Framework

User's Guide



23c
F47496-01
September 2023



Autonomous Health Framework User's Guide, 23c

F47496-01

Copyright © 2016, 2023, Oracle and/or its affiliates.

Primary Authors: Nirmal Kumar, Janet Stern

Contributing Authors: Aparna Kamath, Douglas Williams, Mark Bauer, Richard Strohm, Subhash Chandra

Contributors: Ankita Khandelwal, Arpit Shukla, Carol Colrain, Daniel Semler, Gareth Chapman, Girdhari Ghantiyala, Girish Adiga, Jesus Guillermo Munoz Nunez, Macharapu Prasanth, Mark Scardina, Pallavi Kamath, Robert Caldwell, Sahil Kumar, Troy Anthony, Vern Wagman, Walter Battistella

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, software documentation, data (as defined in the Federal Acquisition Regulation), 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," "commercial computer software documentation," or "limited rights data" 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®, Java, and MySQL 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	vi
Documentation Accessibility	vi
Related Documentation	vi
Conventions	vii

1 Introduction to Oracle Autonomous Health Framework

1.1 Oracle Autonomous Health Framework Problem and Solution Space	1-1
1.1.1 Availability Issues	1-1
1.1.2 Performance Issues	1-3
1.2 Components of Autonomous Health Framework	1-3
1.2.1 Introduction to Oracle Autonomous Health Framework Configuration Audit Tools	1-4
1.2.2 Introduction to Cluster Health Monitor	1-4
1.2.3 Introduction to Oracle Trace File Analyzer	1-5
1.2.4 Introduction to Oracle Cluster Health Advisor	1-6
1.2.5 Introduction to Blocker Resolver	1-7

Part I Analyzing the Cluster Configuration

2 Proactively Detecting and Diagnosing Performance Issues for Oracle RAC

2.1 Oracle Cluster Health Advisor Architecture	2-2
2.2 Monitoring the Oracle Real Application Clusters (Oracle RAC) Environment with Oracle Cluster Health Advisor	2-3
2.3 Using Cluster Health Advisor for Health Diagnosis	2-4
2.4 Calibrating an Oracle Cluster Health Advisor Model for a Cluster Deployment	2-6
2.5 Viewing the Details for an Oracle Cluster Health Advisor Model	2-9
2.6 Managing the Oracle Cluster Health Advisor Repository	2-10
2.7 Viewing the Status of Cluster Health Advisor	2-11

Part II Automatically Monitoring the Cluster

3 Collecting Operating System Resources Metrics

3.1	Understanding Cluster Health Monitor Services	3-2
3.2	Collecting Cluster Health Monitor Data	3-3
3.3	Operating System Metrics Collected by Cluster Health Monitor	3-3
3.4	Detecting Component Failures and Self-healing Autonomously	3-11
3.5	Using Cluster Health Monitor from Enterprise Manager Cloud Control	3-13

4 Monitoring System Metrics for Cluster Nodes

4.1	Monitoring Oracle Clusterware with Oracle Enterprise Manager	4-1
4.2	Monitoring Oracle Clusterware with Cluster Health Monitor	4-3
4.3	Using the Cluster Resource Activity Log to Monitor Cluster Resource Failures	4-3

Part III Automatic Problem Solving

5 Resolving Database and Database Instance Delays

5.1	Blocker Resolver Architecture	5-1
5.2	Optional Configuration for Blocker Resolver	5-2
5.3	Blocker Resolver Diagnostics and Logging	5-3

Part IV Appendixes

A OCLUMON Command Reference

A.1	oclumon debug	A-1
A.2	oclumon dumpnodeview	A-2
A.3	oclumon dumpnodeview local	A-13
A.4	oclumon manage	A-18
A.5	oclumon version	A-20
A.6	oclumon chmdiag description	A-20
A.7	oclumon chmdiag query	A-21
A.8	oclumon chmdiag collect	A-23

B Managing the Cluster Resource Activity Log

B.1	crsctl query calog	B-1
B.2	crsctl get calog maxsize	B-8
B.3	crsctl get calog retentiontime	B-9
B.4	crsctl set calog maxsize	B-9
B.5	crsctl set calog retentiontime	B-10

C chactl Command Reference

C.1	chactl monitor	C-2
C.2	chactl unmonitor	C-3
C.3	chactl status	C-4
C.4	chactl config	C-5
C.5	chactl calibrate	C-6
C.6	chactl query diagnosis	C-7
C.7	chactl query model	C-10
C.8	chactl query repository	C-11
C.9	chactl query calibration	C-11
C.10	chactl remove model	C-14
C.11	chactl rename model	C-14
C.12	chactl export model	C-15
C.13	chactl import model	C-15
C.14	chactl set maxretention	C-15
C.15	chactl resize repository	C-16

D Behavior Changes, Deprecated and Desupported Features

D.1	Oracle Database Quality of Service (QoS) Management is Deprecated in Release 21c	D-1
-----	--	-----

Preface

Oracle Autonomous Health Framework User's Guide explains how to use the Oracle Autonomous Health Framework diagnostic components.

The diagnostic components include Oracle ORAchk, Oracle EXAchk, Cluster Health Monitor, Oracle Trace File Analyzer Collector, Oracle Cluster Health Advisor, and Blocker Resolver.

Oracle Autonomous Health Framework User's Guide also explains how to install and configure Oracle Trace File Analyzer Collector.

This Preface contains these topics:

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

Audience

Database administrators can use this guide to understand how to use the Oracle Autonomous Health Framework diagnostic components. This guide assumes that you are familiar with Oracle Database concepts.

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 Documentation

For more information, see the following Oracle resources:

Related Topics

- *Oracle Automatic Storage Management Administrator's Guide*
- *Oracle Database 2 Day DBA*
- *Oracle Database Concepts*
- *Oracle Database Examples Installation Guide*
- *Oracle Database Licensing Information User Manual*
- *Oracle Database Release Notes*
- *Oracle Database Upgrade Guide*
- *Oracle Grid Infrastructure Installation and Upgrade Guide*
- *Oracle Real Application Clusters Installation Guide for Linux and UNIX*
- *Oracle Real Application Clusters Installation Guide for Microsoft Windows*

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	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

Introduction to Oracle Autonomous Health Framework

Oracle Autonomous Health Framework is a collection of components that analyzes the diagnostic data collected, and proactively identifies issues before they affect the health of your clusters or your Oracle Real Application Clusters (Oracle RAC) databases.

Most of the Oracle Autonomous Health Framework components are already available in Oracle Database 12c release 1 (12.1). In Oracle Database 12c release 2 (12.2), the output of several components is consolidated in the Grid Infrastructure Management Repository (GIMR) and analyzed in real time to detect problematic patterns on the production clusters.

- [Oracle Autonomous Health Framework Problem and Solution Space](#)
Oracle Autonomous Health Framework assists with monitoring, diagnosing, and preventing availability and performance issues.
- [Components of Autonomous Health Framework](#)
This section describes the diagnostic components that are part of Oracle Autonomous Health Framework.

1.1 Oracle Autonomous Health Framework Problem and Solution Space

Oracle Autonomous Health Framework assists with monitoring, diagnosing, and preventing availability and performance issues.

System administrators can use most of the components in Oracle Autonomous Health Framework interactively during installation, patching, and upgrading. Database administrators can use Oracle Autonomous Health Framework to diagnose operational runtime issues and mitigate the impact of these issues.

- [Availability Issues](#)
Availability issues are runtime issues that threaten the availability of software stack.
- [Performance Issues](#)
Performance issues are runtime issues that threaten the performance of the system.

1.1.1 Availability Issues

Availability issues are runtime issues that threaten the availability of software stack.

Availability issues can result from either software issues (Oracle Database, Oracle Grid Infrastructure, operating system) or the underlying hardware resources (CPU, Memory, Network, Storage).

The components within Oracle Autonomous Health Framework address the following availability issues:

Examples of Server Availability Issues

Server availability issues can cause a server to be evicted from the cluster and shut down all the database instances that are running on the server.

Examples of such issues are:

- **Issue:** Network congestion on the private interconnect can cause time-critical internode or storage I/O to have excessive latency or dropped packets. This type of failure typically builds up and can be detected early, and corrected or relieved.
Solution: If a change in the server configuration causes this issue, then Cluster Verification Utility (CVU) detects it if the issue persists for more than an hour. However, Oracle Cluster Health Advisor detects the issue within minutes and presents corrective actions.
- **Issue:** Network failures on the private interconnect caused by a pulled cable or failed network interface card (NIC) can immediately result in evicted nodes.
Solution: Although these types of network failures cannot be detected early, the cause can be narrowed down by using Cluster Health Monitor and Oracle Trace File Analyzer to pinpoint the time of the failure and the network interfaces involved.

Examples of Database Availability Issues

Database availability issues can cause an Oracle database or one of the instances of the database to become unresponsive and thus unavailable to users.

Examples of such issues are:

- **Issue:** Runaway queries or delays can deny critical database resources such as locks, latches, or CPU to other sessions. Denial of critical database resources results in database or an instance of a database being non-responsive to applications.
Solution: Blocker Resolver detects and automatically resolves these types of delays. Also, Oracle Cluster Health Advisor detects, identifies, and notifies the database administrator of such delays and provides an appropriate corrective action.
- **Issue:** Denial-of-service (DoS) attacks, vulnerabilities, or simply software bugs can cause a database or a database instance to be unresponsive.
Solution: Proactive recommendations of known issues and their resolutions provided by Oracle ORAchk can prevent such occurrences. If these issues are not prevented, then automatic collection of logs by Oracle Trace File Analyzer, in addition to data collected by Cluster Health Monitor, can speed up the correction of these issues.
- **Issue:** Configuration changes can cause database outages that are difficult to troubleshoot. For example, incorrect permissions on the `oracle.bin` file can prevent session processes from being created.
Solution: Use Cluster Verification Utility and Oracle ORAchk to speed up identification and correction of these types of issues. You can generate a diff report using Oracle ORAchk to see a baseline comparison of two reports and a list of differences. You can also view configuration reports created by Cluster Verification Utility to verify whether your system meets the criteria for an Oracle installation.

1.1.2 Performance Issues

Performance issues are runtime issues that threaten the performance of the system.

Performance issues can result from either software issues (bugs, configuration problems, data contention, and so on) or client issues (demand, query types, connection management, and so on).

Server and database performance issues are intertwined and difficult to separate. It is easier to categorize them by their origin: database server or client.

Examples of Database Server Performance Issues

- **Issue:** Deviations from best practices in configuration can cause database server performance issues.
Solution: Oracle ORAchk detects configuration issues when Oracle ORAchk runs periodically and notifies the database administrator of the appropriate corrective settings.
- **Issue:** A session can cause other sessions to slow down waiting for the blocking session to release its resource or complete its work.
Solution: Blocker Resolver detects these chains of sessions and automatically terminates the root holder session to relieve the bottleneck.
- **Issue:** Unresolved known issues or unpatched bugs can cause database server performance issues.
Solution: These issues can be detected through the automatic Oracle ORAchk reports and flagged with associated patches or workarounds. Oracle ORAchk is regularly enhanced to include new critical issues, either in existing products or in new product areas.

Examples of Performance Issues Caused by Database Client

- **Issue:** Misconfigured parameters such as SGA and PGA allocation, number of sessions or processes, CPU counts, and so on, can cause database performance degradation.
Solution: Oracle ORAchk and Oracle Cluster Health Advisor detect the settings and consequences respectively and notify you automatically with recommended corrective actions.

1.2 Components of Autonomous Health Framework

This section describes the diagnostic components that are part of Oracle Autonomous Health Framework.

- [Introduction to Oracle Autonomous Health Framework Configuration Audit Tools](#)
Oracle ORAchk and Oracle EXAchk provide a lightweight and non-intrusive health check framework for the Oracle stack of software and hardware components.
- [Introduction to Cluster Health Monitor](#)
Cluster Health Monitor is a component of Oracle Grid Infrastructure, which continuously monitors and stores Oracle Clusterware and operating system resources metrics.
- [Introduction to Oracle Trace File Analyzer](#)
Oracle Trace File Analyzer is a utility for targeted diagnostic collection that simplifies diagnostic data collection for Oracle Clusterware, Oracle Grid Infrastructure, and Oracle

Real Application Clusters (Oracle RAC) systems, in addition to single instance, non-clustered databases.

- [Introduction to Oracle Cluster Health Advisor](#)
Oracle Cluster Health Advisor continuously monitors cluster nodes and Oracle RAC databases for performance and availability issue precursors to provide early warning of problems before they become critical.
- [Introduction to Blocker Resolver](#)
Blocker Resolver is an Oracle Real Application Clusters (Oracle RAC) environment feature that autonomously resolves delays and keeps the resources available.

1.2.1 Introduction to Oracle Autonomous Health Framework Configuration Audit Tools

Oracle ORAchk and Oracle EXAchk provide a lightweight and non-intrusive health check framework for the Oracle stack of software and hardware components.

Oracle ORAchk and Oracle EXAchk:

- Automates risk identification and proactive notification before your business is impacted
- Runs health checks based on critical and reoccurring problems
- Presents high-level reports about your system health risks and vulnerabilities to known issues
- Enables you to drill-down specific problems and understand their resolutions
- Enables you to schedule recurring health checks at regular intervals
- Sends email notifications and diff reports while running in daemon mode
- Integrates the findings into Oracle Health Check Collections Manager and other tools of your choice
- Runs in your environment with no need to send anything to Oracle

You have access to Oracle ORAchk and Oracle EXAchk as a value add-on to your existing support contract. There is no additional fee or license required to run Oracle ORAchk and Oracle EXAchk.

Use Oracle EXAchk for Oracle Engineered Systems except for Oracle Database Appliance. For all other systems, use Oracle ORAchk.

Run health checks for Oracle products using the command-line options.

For more information, see *Oracle Autonomous Health Framework Checks and Diagnostics User's Guide*.

Related Topics

- [Oracle Autonomous Health Framework Checks and Diagnostics User's Guide](#)

1.2.2 Introduction to Cluster Health Monitor

Cluster Health Monitor is a component of Oracle Grid Infrastructure, which continuously monitors and stores Oracle Clusterware and operating system resources metrics.

Enabled by default, Cluster Health Monitor:

- Assists node eviction analysis
- Logs all process data locally
- Enables you to define pinned processes
- Listens to CSS and GIPC events
- Categorizes processes by type
- Supports plug-in collectors such as `traceroute`, `netstat`, `ping`, and so on
- Provides CSV output for ease of analysis

Cluster Health Monitor serves as a data feed for other Oracle Autonomous Health Framework components such as Oracle Cluster Health Advisor.

Related Topics

- [Collecting Operating System Resources Metrics](#)
CHM is a high-performance, lightweight daemon that collects, analyzes, aggregates, and stores a large set of operating system metrics to help you diagnose and troubleshoot system issues.

1.2.3 Introduction to Oracle Trace File Analyzer

Oracle Trace File Analyzer is a utility for targeted diagnostic collection that simplifies diagnostic data collection for Oracle Clusterware, Oracle Grid Infrastructure, and Oracle Real Application Clusters (Oracle RAC) systems, in addition to single instance, non-clustered databases.

Enabled by default, Oracle Trace File Analyzer:

- Provides comprehensive first failure diagnostics collection
- Efficiently collects, packages, and transfers diagnostic data to Oracle Support
- Reduces round trips between customers and Oracle

Oracle Trace File Analyzer reduces the time required to obtain the correct diagnostic data, which eventually saves your business money.

For more information, see *Oracle Autonomous Health Framework Checks and Diagnostics User's Guide*.

New Attention Log for Efficient Critical Issue Resolution

Diagnosability of database issues is enhanced through a new attention log, as well as classification of information written to database trace files. The new attention log is written in a structured format (XML or JSON) that is much easier to process or interpret and only contains information that requires attention from an administrator. The contents of trace files now contains information that enables much easier classification of trace messages, such as for security and sensitivity.

Enhanced diagnosability features simplify database administration and improve data security.

For more information, see Attention Log

Related Topics

- [Oracle Autonomous Health Framework Checks and Diagnostics User's Guide](#)

1.2.4 Introduction to Oracle Cluster Health Advisor

Oracle Cluster Health Advisor continuously monitors cluster nodes and Oracle RAC databases for performance and availability issue precursors to provide early warning of problems before they become critical.

Oracle Cluster Health Advisor is integrated into Oracle Enterprise Manager Cloud Control (EMCC) Incident Manager.

Oracle Cluster Health Advisor does the following:

- Detects node and database performance problems
- Provides early-warning alerts and corrective action
- Supports on-site calibration to improve sensitivity

In Oracle Database 12c release 2 (12.2.0.1), Oracle Cluster Health Advisor supports the monitoring of two critical subsystems of Oracle Real Application Clusters (Oracle RAC): the database instance and the host system. Oracle Cluster Health Advisor determines and tracks the health status of the monitored system. It periodically samples a wide variety of key measurements from the monitored system.

Over a hundred database and cluster node problems have been modeled, and the specific operating system and Oracle Database metrics that indicate the development or existence of these problems have been identified. This information is used to construct a trained, calibrated model that is based on a normal operational period of the target system.

Oracle Cluster Health Advisor runs an analysis multiple times a minute. Oracle Cluster Health Advisor estimates an expected value of an observed input based on the default model. Oracle Cluster Health Advisor then performs anomaly detection for each input based on the difference between observed and expected values. If sufficient inputs associated with a specific problem are abnormal, then Oracle Cluster Health Advisor raises a warning and generates an immediate targeted diagnosis and corrective action.

Oracle Cluster Health Advisor models are conservative to prevent false warning notifications. However, the default configuration may not be sensitive enough for critical production systems. Therefore, Oracle Cluster Health Advisor provides an onsite model calibration capability to use actual production workload data to form the basis of its default setting and increase the accuracy and sensitivity of node and database models.

Oracle Cluster Health Advisor stores the analysis results, along with diagnosis information, corrective action, and metric evidence for later triage, in the Grid Infrastructure Management Repository (GIMR). Oracle Cluster Health Advisor also sends warning messages to Enterprise Manager Cloud Control using the Oracle Clusterware event notification protocol.

You can also use Oracle Cluster Health Advisor to diagnose and triage past problems. You specify the past dates through Oracle Enterprise Manager Cloud Control (EMCC) Incident Manager or through the command-line interface CHACTL. Manage the capability of Oracle Cluster Health Advisor to review past problems by configuring the retention setting for Oracle Cluster Health Advisor's tablespace in the Grid Infrastructure Management Repository (GIMR). The default retention period is 72 hours.

1.2.5 Introduction to Blocker Resolver

Blocker Resolver is an Oracle Real Application Clusters (Oracle RAC) environment feature that autonomously resolves delays and keeps the resources available.

Enabled by default, Blocker Resolver:

- Reliably detects database delays and deadlocks
- Autonomously resolves database delays and deadlocks
- Logs all detections and resolutions
- Provides SQL interface to configure sensitivity (Normal/High) and trace file sizes

A database delays when a session blocks a chain of one or more sessions. The blocking session holds a resource such as a lock or latch that prevents the blocked sessions from progressing. The chain of sessions has a root or a final blocker session, which blocks all the other sessions in the chain. Blocker Resolver resolves these issues autonomously by detecting and resolving the delays.

Related Topics

- [Resolving Database and Database Instance Delays](#)
Blocker Resolver preserves the database performance by resolving delays and keeping the resources available.

Part I

Analyzing the Cluster Configuration

You can use tools in the Autonomous Health Framework to analyze your cluster configuration.

- [Proactively Detecting and Diagnosing Performance Issues for Oracle RAC](#)
Oracle Cluster Health Advisor provides system and database administrators with early warning of pending performance issues, and root causes and corrective actions for Oracle RAC databases and cluster nodes. Use Oracle Cluster Health Advisor to increase availability and performance management.

2

Proactively Detecting and Diagnosing Performance Issues for Oracle RAC

Oracle Cluster Health Advisor provides system and database administrators with early warning of pending performance issues, and root causes and corrective actions for Oracle RAC databases and cluster nodes. Use Oracle Cluster Health Advisor to increase availability and performance management.

Oracle Cluster Health Advisor estimates an expected value of an observed input based on the default model, which is a trained calibrated model based on a normal operational period of the target system. Oracle Cluster Health Advisor then performs anomaly detection for each input based on the difference between observed and expected values. If sufficient inputs associated with a specific problem are abnormal, then Oracle Cluster Health Advisor raises a warning and generates an immediate targeted diagnosis and corrective action.

Oracle Cluster Health Advisor stores the analysis results, along with diagnosis information, corrective action, and metric evidence for later triage, in the Grid Infrastructure Management Repository (GIMR). Oracle Cluster Health Advisor also sends warning messages to Enterprise Manager Cloud Control using the Oracle Clusterware event notification protocol.

The ability of Oracle Cluster Health Advisor to detect performance and availability issues on Oracle Exadata systems has been improved in this release.

With the Oracle Cluster Health Advisor support for Oracle Solaris, you can now get early detection and prevention of performance and availability issues in your Oracle RAC database deployments.

For more information on *Installing Grid Infrastructure Management Repository*, see *Oracle® Grid Infrastructure Grid Infrastructure Installation and Upgrade Guide 20c for Linux*.

- [Oracle Cluster Health Advisor Architecture](#)
Oracle Cluster Health Advisor runs as a highly available cluster resource, `ochad`, on each node in the cluster.
- [Monitoring the Oracle Real Application Clusters \(Oracle RAC\) Environment with Oracle Cluster Health Advisor](#)
Oracle Cluster Health Advisor is automatically provisioned on each node by default when Oracle Grid Infrastructure is installed for Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database.
- [Using Cluster Health Advisor for Health Diagnosis](#)
Oracle Cluster Health Advisor raises and clears problems autonomously and stores the history in the Grid Infrastructure Management Repository (GIMR).
- [Calibrating an Oracle Cluster Health Advisor Model for a Cluster Deployment](#)
As shipped with default node and database models, Oracle Cluster Health Advisor is designed not to generate false warning notifications.
- [Viewing the Details for an Oracle Cluster Health Advisor Model](#)
Use the `chactl query model` command to view the model details.

- [Managing the Oracle Cluster Health Advisor Repository](#)
Oracle Cluster Health Advisor repository stores the historical records of cluster host problems, database problems, and associated metric evidence, along with models.
- [Viewing the Status of Cluster Health Advisor](#)
SRVCTL commands are the tools that offer total control on managing the life cycle of Oracle Cluster Health Advisor as a highly available service.
- [Enhanced Cluster Health Advisor Support for Oracle Pluggable Databases](#)
The Cluster Health Advisor (CHA) diagnostic capabilities have been extended to support 4K PDBs, up from 256 in Oracle Database 23c.

Related Topics

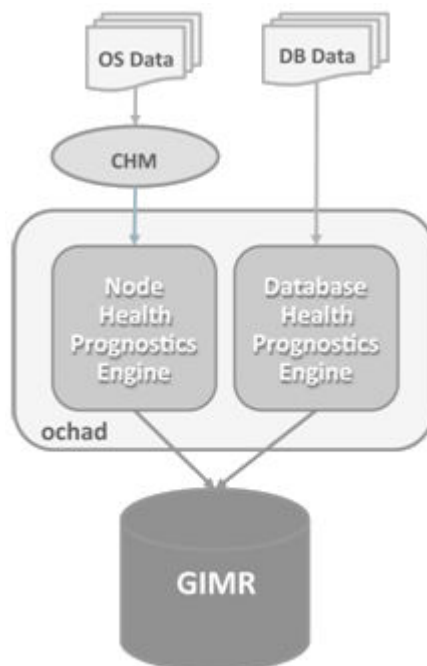
- [Introduction to Oracle Cluster Health Advisor](#)
Oracle Cluster Health Advisor continuously monitors cluster nodes and Oracle RAC databases for performance and availability issue precursors to provide early warning of problems before they become critical.
- [Installing Grid Infrastructure Management Repository](#)

2.1 Oracle Cluster Health Advisor Architecture

Oracle Cluster Health Advisor runs as a highly available cluster resource, `ochad`, on each node in the cluster.

Each Oracle Cluster Health Advisor daemon (`ochad`) monitors the operating system on the cluster node and optionally, each Oracle Real Application Clusters (Oracle RAC) database instance on the node.

Figure 2-1 Oracle Cluster Health Advisor Architecture



The `ochad` daemon receives operating system metric data from the Cluster Health Monitor and gets Oracle RAC database instance metrics from a memory-mapped file. The daemon does not require a connection to each database instance. This data, along with the selected model, is used in the Health Prognostics Engine of Oracle Cluster Health Advisor for both the node and each monitored database instance in order to analyze their health multiple times a minute.

2.2 Monitoring the Oracle Real Application Clusters (Oracle RAC) Environment with Oracle Cluster Health Advisor

Oracle Cluster Health Advisor is automatically provisioned on each node by default when Oracle Grid Infrastructure is installed for Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database.

Oracle Cluster Health Advisor does not require any additional configuration. The credentials of OCHAD daemon user in the Grid Infrastructure Management Repository (GIMR), are securely and randomly generated and stored in the Oracle Grid Infrastructure Credential Store.

When Oracle Cluster Health Advisor detects an Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database instance as running, Oracle Cluster Health Advisor autonomously starts monitoring the cluster nodes. Use `CHACTL` while logged in as the Grid user to turn on monitoring of the database.

To monitor the Oracle Real Application Clusters (Oracle RAC) environment:

1. To monitor a database, run the following command:

```
$ chactl monitor database -db db_unique_name
```

Oracle Cluster Health Advisor monitors all instances of the Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node database using the default model. Oracle Cluster Health Advisor cannot monitor single-instance Oracle databases, even if the single-instance Oracle databases share the same cluster as Oracle Real Application Clusters (Oracle RAC) databases.

Oracle Cluster Health Advisor preserves database monitoring status across cluster restarts as Oracle Cluster Health Advisor stores the status information in the GIMR. Each database instance is monitored independently both across Oracle Real Application Clusters (Oracle RAC) database nodes and when more than one database run on a single node.

2. To stop monitoring a database, run the following command:

```
$ chactl unmonitor database -db db_unique_name
```

Oracle Cluster Health Advisor stops monitoring all instances of the specified database. However, Oracle Cluster Health Advisor does not delete any data or problems until it is aged out beyond the retention period.

3. To check monitoring status of all cluster nodes and databases, run the following command:

```
$ chactl status
```

Use the `-verbose` option to see more details, such as the models used for the nodes and each database.

2.3 Using Cluster Health Advisor for Health Diagnosis

Oracle Cluster Health Advisor raises and clears problems autonomously and stores the history in the Grid Infrastructure Management Repository (GIMR).

The Oracle Grid Infrastructure user can query the stored information using CHACTL.

To query the diagnostic data:

1. To query currently open problems, run the following command:

```
chactl query diagnosis -db db_unique_name -start time -end time
```

In the syntax example, *db_unique_name* is the name of your database instance. You also specify the start time and end time for which you want to retrieve data. Specify date and time in the YYYY-MM-DD HH24:MI:SS format.

2. Use the `-htmlfile file_name` option to save the output in HTML format.

Example 2-1 Cluster Health Advisor Output Examples in Text and HTML Format

This example shows the default text output for the `chactl query diagnosis` command for a database named `oltpacbd`.

```
$ chactl query diagnosis -db oltpacdb -start "2016-02-01 02:52:50" -
end "2016-02-01 03:19:15"
2016-02-01 01:47:10.0 Database oltpacdb DB Control File IO
Performance (oltpacdb_1) [detected]
2016-02-01 01:47:10.0 Database oltpacdb DB Control File IO
Performance (oltpacdb_2) [detected]
2016-02-01 02:52:15.0 Database oltpacdb DB CPU Utilization
(oltpacdb_2) [detected]
2016-02-01 02:52:50.0 Database oltpacdb DB CPU Utilization
(oltpacdb_1) [detected]
2016-02-01 02:59:35.0 Database oltpacdb DB Log File Switch
(oltpacdb_1) [detected]
2016-02-01 02:59:45.0 Database oltpacdb DB Log File Switch
(oltpacdb_2) [detected]
```

```
Problem: DB Control File IO Performance
Description: CHA has detected that reads or writes to the control
files are slower than expected.
Cause: The Cluster Health Advisor (CHA) detected that reads or writes
to the control files were slow
because of an increase in disk IO.
The slow control file reads and writes may have an impact on
checkpoint and Log Writer (LGWR) performance.
Action: Separate the control files from other database files and move
them to faster disks or Solid State Devices.
```

```
Problem: DB CPU Utilization
```

Description: CHA detected larger than expected CPU utilization for this database.
Cause: The Cluster Health Advisor (CHA) detected an increase in database CPU utilization because of an increase in the database workload.
Action: Identify the CPU intensive queries by using the Automatic Diagnostic and Defect Manager (ADDM) and follow the recommendations given there. Limit the number of CPU intensive queries or relocate sessions to less busy machines. Add CPUs if the CPU capacity is insufficient to support the load without a performance degradation or effects on other databases.

Problem: DB Log File Switch

Description: CHA detected that database sessions are waiting longer than expected for log switch completions.
Cause: The Cluster Health Advisor (CHA) detected high contention during log switches because the redo log files were small and the redo logs switched frequently.
Action: Increase the size of the redo logs.

The timestamp displays date and time when the problem was detected on a specific host or database.



Note:

The same problem can occur on different hosts and at different times, yet the diagnosis shows complete details of the problem and its potential impact. Each problem also shows targeted corrective or preventive actions.

Here is an example of what the output looks like in the HTML format.

```
$ chactl query diagnosis -start "2016-07-03 20:50:00" -end "2016-07-04 03:50:00" -htmlfile ~/chaprob.html
```

Figure 2-2 Cluster Health Advisor Diagnosis HTML Output

Timestamp	Target Information	Event Name	Detected/Cleared
2016-07-03 01:49:30.0	Host rwsbi07	Host CPU Utilization	detected
2016-07-03 01:49:50.0	Host rwsbi06	Host CPU Utilization	detected
2016-07-03 05:54:55.0	Host rwsbi06	Host Memory Consumption	detected
2016-07-04 03:40:00.0	Host rwsbi07	Host CPU Utilization	cleared
2016-07-04 03:40:05.0	Host rwsbi06	Host CPU Utilization	cleared
2016-07-04 03:40:05.0	Host rwsbi06	Host Memory Consumption	cleared

Problem	Description	Cause	Action
Host CPU Utilization	CHA detected larger than expected CPU utilization on this node. The available CPU resource may not be sufficient to support application failover or relocation of databases to this node.	The Cluster Health Advisor (CHA) detected an unexpected increase in CPU utilization by databases or applications on this node.	Identify CPU intensive processes and databases by reviewing Cluster Health Monitoring (CHM) data. Relocate databases to less busy machines, or limit the number of connections to databases on this node. Add nodes if more resources are required.
Host Memory Consumption	CHA detected that more memory than expected is consumed on this server. The memory is not allocated by sessions of this database.	The Cluster Health Advisor (CHA) detected an increase in memory consumption by other databases or by applications not connected to a database on this node.	Identify the top memory consumers by using the Cluster Health Monitor (CHM).

Related Topics

- [chactl query diagnosis](#)
Use the `chactl query diagnosis` command to return problems and diagnosis, and suggested corrective actions associated with the problem for specific cluster nodes or Oracle Real Application Clusters (Oracle RAC) databases.

2.4 Calibrating an Oracle Cluster Health Advisor Model for a Cluster Deployment

As shipped with default node and database models, Oracle Cluster Health Advisor is designed not to generate false warning notifications.

You can increase the sensitivity and accuracy of the Oracle Cluster Health Advisor models for a specific workload using the `chactl calibrate` command.

Oracle recommends that a minimum of 6 hours of data be available and that both the cluster and databases use the same time range for calibration.

The `chactl calibrate` command analyzes a user-specified time interval that includes all workload phases operating normally. This data is collected while Oracle Cluster Health Advisor is monitoring the cluster and all the databases for which you want to calibrate.

1. To check if sufficient data is available, run the `query calibration` command.

If 720 or more records are available, then Oracle Cluster Health Advisor successfully performs the calibration. The calibration function may not consider some data records to be normally occurring for the workload profile being used. In this case, filter the data by using the `KPISET` parameters in both the `query calibration` command and the `calibrate` command.

For example:

```
$ chactl query calibration -db oltpacdb -timeranges
'start=2016-07-26 01:00:00,end=2016-07-26 02:00:00,start=2016-07-26
03:00:00,end=2016-07-26 04:00:00'
-kpiset 'name=CPUPERCENT min=20 max=40, name=IOTHROUGHPUT min=500
max=9000' -interval 2
```

2. Start the calibration and store the model under a user-specified name for the specified date and time range.

For example:

```
$ chactl calibrate cluster -model weekday -timeranges 'start=2016-07-03
20:50:00,end=2016-07-04 15:00:00'
```

After completing the calibration, Oracle Cluster Health Advisor automatically stores the new model in GIMR.

3. Use the new model to monitor the cluster as follows:

For example:

```
$ chactl monitor cluster -model weekday
```

Example 2-2 Output for the chactl query calibrate command

```
Database name : oltpacdb
Start time : 2016-07-26 01:03:10
End time : 2016-07-26 01:57:25
Total Samples : 120
Percentage of filtered data : 8.32%
The number of data samples may not be sufficient for calibration.
```

1) Disk read (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
4.96	0.20	8.98	0.06	25.68
<25	<50	<75	<100	>=100
97.50%	2.50%	0.00%	0.00%	0.00%

2) Disk write (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
27.73	9.72	31.75	4.16	109.39
<50	<100	<150	<200	>=200
73.33%	22.50%	4.17%	0.00%	0.00%

3) Disk throughput (ASM) (IO/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
2407.50	1500.00	1978.55	700.00	7800.00

<5000	<10000	<15000	<20000	>=20000
83.33%	16.67%	0.00%	0.00%	0.00%

4) CPU utilization (total) (%)

MEAN	MEDIAN	STDDEV	MIN	MAX
21.99	21.75	1.36	20.00	26.80

<20	<40	<60	<80	>=80
0.00%	100.00%	0.00%	0.00%	0.00%

5) Database time per user call (usec/call)

MEAN	MEDIAN	STDDEV	MIN	MAX
267.39	264.87	32.05	205.80	484.57

<10000000	<20000000	<30000000	<40000000	<50000000	<60000000
<70000000	>=70000000				
100.00%	0.00%	0.00%	0.00%	0.00%	0.00%
0.00%					

Database name : oltpacdb

Start time : 2016-07-26 03:00:00

End time : 2016-07-26 03:53:30

Total Samples : 342

Percentage of filtered data : 23.72%

The number of data samples may not be sufficient for calibration.

1) Disk read (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
12.18	0.28	16.07	0.05	60.98

<25	<50	<75	<100	>=100
64.33%	34.50%	1.17%	0.00%	0.00%

2) Disk write (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
57.57	51.14	34.12	16.10	135.29

<50	<100	<150	<200	>=200
49.12%	38.30%	12.57%	0.00%	0.00%

3) Disk throughput (ASM) (IO/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
5048.83	4300.00	1730.17	2700.00	9000.00

<5000	<10000	<15000	<20000	>=20000
63.74%	36.26%	0.00%	0.00%	0.00%

4) CPU utilization (total) (%)

MEAN	MEDIAN	STDDEV	MIN	MAX
------	--------	--------	-----	-----

```
23.10      22.80      1.88      20.00      31.40

<20       <40        <60       <80       >=80
0.00%     100.00%   0.00%     0.00%     0.00%
```

5) Database time per user call (usec/call)

```
MEAN      MEDIAN     STDDEV     MIN        MAX
744.39    256.47    2892.71    211.45     45438.35
```

```
<10000000 <20000000 <30000000 <40000000 <50000000 <60000000 <70000000
>=70000000
100.00%   0.00%     0.00%     0.00%     0.00%     0.00%     0.00%
```

Related Topics

- [chactl calibrate](#)
Use the `chactl calibrate` command to create a new model that has greater sensitivity and accuracy.
- [chactl query calibration](#)
Use the `chactl query calibration` command to view detailed information about the calibration data of a specific target.
- [chactl Command Reference](#)
The Oracle Cluster Health Advisor commands enable the Oracle Grid Infrastructure user to administer basic monitoring functionality on the targets.

2.5 Viewing the Details for an Oracle Cluster Health Advisor Model

Use the `chactl query model` command to view the model details.

- You can review the details of an Oracle Cluster Health Advisor model at any time using the `chactl query model` command.

For example:

```
$ chactl query model -name weekday
Model: weekday
Target Type: CLUSTERWARE
Version: OS12.2_V14_0.9.8
OS Calibrated on: Linux amd64
Calibration Target Name: MYCLUSTER
Calibration Date: 2016-07-05 01:13:49
Calibration Time Ranges: start=2016-07-03 20:50:00,end=2016-07-04 15:00:00
Calibration KPIs: not specified
```

You can also rename, import, export, and delete the models.

2.6 Managing the Oracle Cluster Health Advisor Repository

Oracle Cluster Health Advisor repository stores the historical records of cluster host problems, database problems, and associated metric evidence, along with models.

The Oracle Cluster Health Advisor repository is used to diagnose and triage periodic problems. By default, the repository is sized to retain data for 16 targets (nodes and database instances) for 72 hours. If the number of targets increase, then the retention time is automatically decreased. Oracle Cluster Health Advisor generates warning messages when the retention time goes below 72 hours, and stops monitoring and generates a critical alert when the retention time goes below 24 hours.

Use CHACTL commands to manage the repository and set the maximum retention time.

1. To retrieve the repository details, use the following command:

```
$ chactl query repository
```

For example, running the command mentioned earlier shows the following output:

```
specified max retention time(hrs) : 72
available retention time(hrs)      : 212
available number of entities       : 2
allocated number of entities       : 0
total repository size(gb)         : 2.00
allocated repository size(gb)      : 0.07
```

2. To set the maximum retention time in hours, based on the current number of targets being monitored, use the following command:

```
$ chactl set maxretention -time number_of_hours
```

For example:

```
$ chactl set maxretention -time 80
max retention successfully set to 80 hours
```

Note:

The `maxretention` setting limits the oldest data retained in the repository, but is not guaranteed to be maintained if the number of monitored targets increase. In this case, if the combination of monitored targets and number of hours are not sufficient, then increase the size of the Oracle Cluster Health Advisor repository.

3. To increase the size of the Oracle Cluster Health Advisor repository, use the `chactl resize repository` command.

For example, to resize the repository to support 32 targets using the currently set maximum retention time, you would use the following command:

```
$ chactl resize repository -entities 32  
repository successfully resized for 32 targets
```

2.7 Viewing the Status of Cluster Health Advisor

SRVCTL commands are the tools that offer total control on managing the life cycle of Oracle Cluster Health Advisor as a highly available service.

Use SRVCTL commands to check the status and configuration of Oracle Cluster Health Advisor service on any active hub or leaf nodes of the Oracle RAC cluster.



Note:

A target is monitored only if it is running and the Oracle Cluster Health Advisor service is also running on the host node where the target exists.

1. To check the status of Oracle Cluster Health Advisor service on all nodes in the Oracle RAC cluster:

```
srvctl status cha [-help]
```

For example:

```
# srvctl status cha  
Cluster Health Advisor is running on nodes racNode1, racNode2.  
Cluster Health Advisor is not running on nodes racNode3, racNode4.
```

2. To check if Oracle Cluster Health Advisor service is enabled or disabled on all nodes in the Oracle RAC cluster:

```
srvctl config cha [-help]
```

For example:

```
# srvctl config cha  
Cluster Health Advisor is enabled on nodes racNode1, racNode2.  
Cluster Health Advisor is not enabled on nodes racNode3, racNode4.
```

2.8 Enhanced Cluster Health Advisor Support for Oracle Pluggable Databases

The Cluster Health Advisor (CHA) diagnostic capabilities have been extended to support 4K PDBs, up from 256 in Oracle Database 23c.

Going forward, this is crucial for Oracle Autonomous Database deployments. CHA's problem detection and root cause analysis will be improved by considering DB events such as reconfiguration. This improves detection, analysis, and targeted preventative actions for problems such as instance evictions.

Part II

Automatically Monitoring the Cluster

You can use components of Autonomous Health Framework to monitor your cluster on a regular basis.

- [Collecting Operating System Resources Metrics](#)
CHM is a high-performance, lightweight daemon that collects, analyzes, aggregates, and stores a large set of operating system metrics to help you diagnose and troubleshoot system issues.
- [Monitoring System Metrics for Cluster Nodes](#)
This chapter explains the methods to monitor Oracle Clusterware.

3

Collecting Operating System Resources Metrics

CHM is a high-performance, lightweight daemon that collects, analyzes, aggregates, and stores a large set of operating system metrics to help you diagnose and troubleshoot system issues.

You can now configure Oracle Cluster Health Monitor to operate in local mode to report the operating system metrics using the `oclumon dumpnodeview local` command even if you have not deployed GIMR.

In local mode, you can get only the local node data. In earlier releases, Oracle Cluster Health Monitor required GIMR to report the operating system metrics using the `oclumon dumpnodeview` command.

Supported Platforms

Linux, Microsoft Windows, Solaris, AIX, IBM Z Series, and ARM

Why CHM is unique

CHM	Typical OS Collector
Last man standing - daemon runs memory locked, RT scheduling class ensuring consistent data collection under system load.	Inconsistent data dropouts due to scheduling delays under system load.
High fidelity data sampling rate, 5 seconds. Very low resource usage profile at 5-second sampling rates.	Running multiple utilities creates additional overhead on the system being monitored, and worsens with higher sampling rates.
High Availability daemon, collated data collections across multiple resource categories. Highly optimized collector (data read directly from the operating system, same source as utilities).	Set of scripts/command-line utilities, for example, <code>top</code> , <code>ps</code> , <code>vmstat</code> , <code>iostat</code> , and so on re-directing their output to one or more files for every collection sample.
Collected data is collated into a system snapshot overview (Nodeview) on every sample, Nodeview also contains additional summarization and analysis of the collected data across multiple resource categories.	System snapshot overviews across different resource categories are very tedious to collate.
Significant inline analysis and summarization during data collection and collation into the Nodeview greatly reduces tedious, manual, time-consuming analysis to drive meaningful insights.	The analysis is time-consuming and processing-intensive as the output of various utilities across multiple files needs to be collated, parsed, interpreted, and then analyzed for meaningful insights.
Performs Clusterware-aware specific metrics collection (Process Aggregates, ASM/OCR/VD disk tagging, Private/Public NIC tagging). Also provides an extensive toolset for in-depth data analysis and visualization.	None

- [Understanding Cluster Health Monitor Services](#)
Cluster Health Monitor uses system monitor (`osysmond`) and cluster logger (`ologgerd`) services to collect diagnostic data.
- [Collecting Cluster Health Monitor Data](#)
Collect Cluster Health Monitor data from any node in the cluster.
- [Operating System Metrics Collected by Cluster Health Monitor](#)
Review the metrics collected by CHM.
- [Detecting Component Failures and Self-healing Autonomously](#)
Improved ability to detect component failures and self-heal autonomously improves business continuity.
- [Using Cluster Health Monitor from Enterprise Manager Cloud Control](#)
Histograms presented in real-time and historical modes enable you to understand precisely what was happening at the time of degradation or failure.

Related Topics

- [Introduction to Cluster Health Monitor](#)
Cluster Health Monitor is a component of Oracle Grid Infrastructure, which continuously monitors and stores Oracle Clusterware and operating system resources metrics.

3.1 Understanding Cluster Health Monitor Services

Cluster Health Monitor uses system monitor (`osysmond`) and cluster logger (`ologgerd`) services to collect diagnostic data.

About the System Monitor Service

The system monitor service (`osysmond`) is a real-time monitoring and operating system metric collection service that runs on each cluster node. The system monitor service is managed as a High Availability Services (HAS) resource. The system monitor service forwards the collected metrics to the cluster logger service, `ologgerd`. The cluster logger service stores the data in the Oracle Grid Infrastructure Management Repository database.

In addition, `osysmond` persists the collected operating system metrics under a directory in `ORACLE_BASE`.

Metric Repository is auto-managed on the local filesystem. You can change the location and size of the repository.

- Nodeview samples are continuously written to the repository (JSON record)
- Historical data is auto-archived into hourly zip files
- Archived files are automatically purged once the default retention limit is reached (default: 200 MB)

About the Cluster Logger Service

The cluster logger service (`ologgerd`) is responsible for preserving the data collected by the system monitor service (`osysmond`) in the Oracle Grid Infrastructure Management Repository database. In a cluster, there is one cluster logger service (`ologgerd`) per 32 nodes. More logger services are spawned for every additional 32

nodes. The additional nodes can be a sum of Hub and Leaf Nodes. Oracle Clusterware relocates and starts the service on a different node, if:

- The logger service fails and is not able to come up after a fixed number of retries
- The *node* where the cluster logger service is running, is down

Support for Deploying Grid Infrastructure Management Repository (GIMR) into a Separate Oracle Home

Starting with Oracle Grid Infrastructure 20c, you must configure the Grid Infrastructure Management Repository (GIMR) in a separate Oracle home, instead of in the Grid home. This option is available when you configure GIMR during a fresh Oracle Grid Infrastructure installation or you add a GIMR to an existing deployment. It is mandatory to configure GIMR in a separate Oracle home when you upgrade Oracle Grid infrastructure with an existing GIMR deployed in it.

A separate Oracle home for the GIMR ensures faster rolling upgrades, less errors, and fewer rollback situations. The Oracle Grid Infrastructure installation owner user must own the GIMR home.

For more information, see [Installing Grid Infrastructure Management Repository](#)

Remote GIMR Support for Oracle Standalone Clusters

The remote Grid Infrastructure Management Repository (GIMR) feature for Oracle Standalone Cluster enables you to use a centralized GIMR. This feature does not require local cluster resources to host the GIMR.

The remote GIMR feature provides access to a persistent data store that significantly enhances the proactive diagnostic functionality of Cluster Health Monitor, Cluster Health Advisor, and Autonomous Health Framework clients. The remote GIMR feature saves cost by freeing up local resources and licensed database server resources.

For more information, see [Creating GIMR Credentials File for Oracle Standalone Clusters With Remote GIMR](#)

3.2 Collecting Cluster Health Monitor Data

Collect Cluster Health Monitor data from any node in the cluster.

Oracle recommends that you run the `tfactl diagcollect` command to collect diagnostic data when an Oracle Clusterware error occurs.

3.3 Operating System Metrics Collected by Cluster Health Monitor

Review the metrics collected by CHM.

Overview of Metrics

CHM groups the operating system data collected into a **Nodeview**. A **Nodeview** is a grouping of metric sets where each metric set contains detailed metrics of a unique system resource.

Brief description of metric sets are as follows:

- **CPU metric set:** Metrics for top 127 CPUs sorted by usage percentage
- **Device metric set:** Metrics for 127 devices that include ASM/VD/OCR along with those having a high average wait time
- **Process metric set:** Metrics for 127 processes
 - Top 25 CPU consumers (idle processes not reported)
 - Top 25 Memory consumers (RSS < 1% of total RAM not reported)
 - Top 25 I/O consumers
 - Top 25 File Descriptors consumers (helps to identify top inode consumers)
 - Process Aggregation: Metrics summarized by foreground and background processes for all Oracle Database and Oracle ASM instances
- **Network metric set:** Metrics for 16 NICS that include public and private interconnects
- **NFS metric set:** Metrics for 32 NFS ordered by round trip time
- **Protocol metric set:** Metrics for protocol groups TCP, UDP, and IP
- **Filesystem metric set:** Metrics for filesystem utilization
- **Critical resources metric set:** Metrics for critical system resource utilization
 - CPU Metrics: system-wide CPU utilization statistics
 - Memory Metrics: system-wide memory statistics
 - Device Metrics: system-wide device statistics distinct from individual device metric set
 - NFS Metrics: Total NFS devices collected every 30 seconds
 - Process Metrics: system-wide unique process metrics

CPU Metric Set

Contains metrics from all CPU cores ordered by usage percentage.

Table 3-1 CPU Metric Set

Metric Name (units)	Description
system [%]	Percentage of CPU utilization occurred while executing at the system level (kernel).
user [%]	Percentage of CPU utilization occurred while executing at the user level (application).
usage [%]	Total utilization (system[%] + user[%]).
nice [%]	Percentage of CPU utilization occurred while executing at the user level with nice priority.
ioWait [%]	Percentage of time that the CPU was idle during which the system had an outstanding disk I/O request.
steal [%]	Percentage of time spent in involuntary wait by the virtual CPU while the hypervisor was servicing another virtual processor.

Device Metric Set

Contains metrics from all disk devices/partitions ordered by their service time in milliseconds.

Table 3-2 Device Metric Set

Metric Name (units)	Description
ioR [KB/s]	Amount of data read from the device.
ioW [KB/s]	Amount of data written to the device.
numIOs [#s]	Average disk I/O operations.
qLen [#]	Number of I/O queued requests, that is, in a wait state.
aWait [msec]	Average wait time per I/O.
svcTm [msec]	Average service time per I/O request.
util [%]	Percent utilization of the device (same as '%util' metric from the <code>iostat -x</code> command. Represents the percentage of time device was active).

Process Metric Set

Contains multiple categories of summarized metric data computed across all system processes.

Table 3-3 Process Metric Set

Metric Name (units)	Description
pid	Process ID.
pri	Process priority (raw value from the operating system).
psr	The processor that process is currently assigned to or running on.
pPid	Parent process ID.
nice	Nice value of the process.
state	State of the process. For example, R->Running, S->Interruptible sleep, and so on.
class	Scheduling class of the process. For example, RR->RobinRound, FF->First in First out, B->Batch scheduling, and so on.
fd [#]	Number of file descriptors opened by this process, which is updated every 30 seconds.
name	Name of the process.
cpu [%]	Process CPU utilization across cores. For example, 50% => 50% of single core, 400% => 100% usage of 4 cores.
thrds [#]	Number of threads created by this process.
vmem [KB]	Process virtual memory usage (KB).
shMem [KB]	Process shared memory usage (KB).
rss [KB]	Process memory-resident set size (KB).

Table 3-3 (Cont.) Process Metric Set

Metric Name (units)	Description
ioR [KB/s]	I/O read in kilobytes per second.
ioW [KB/s]	I/O write in kilobytes per second.
ioT [KB/s]	I/O total in kilobytes per second.
cswch [#s]	Context switch per second. Collected only for a few critical Oracle Database processes.
nvcswh [#s]	Non-voluntary context switch per second. Collected only for a few critical Oracle Database processes.
cumulativeCpu [ms]	Amount of CPU used so far by the process in microseconds.

NIC Metric Set

Contains metrics from all network interfaces ordered by their total rate in kilobytes per second.

Table 3-4 NIC Metric Set

Metric Name (units)	Description
name	Name of the interface.
tag	Tag for the interface, for example, public , private , and so on.
mtu [B]	Size of the maximum transmission unit in bytes supported for the interface.
rx [Kbps]	Average network receive rate.
tx [Kbps]	Average network send rate.
total [Kbps]	Average network transmission rate (rx[Kb/s] + tx[Kb/s]).
rxPkt [#s]	Average incoming packet rate.
txPkt [#s]	Average outgoing packet rate.
pkt [#s]	Average rate of packet transmission (rxPkt[#s] + txPkt[#s]).
rxDscrd [#s]	Average rate of dropped/discarded incoming packets.
txDscrd [#s]	Average rate of dropped/discarded outgoing packets.
rxUnicast [#s]	Average rate of unicast packets received.
rxNonUnicast [#s]	Average rate of multicast packets received.
dscrd [#s]	Average rate of total discarded packets (rxDscrd + txDscrd).
rxErr [#s]	Average error rate for incoming packets.
txErr [#s]	Average error rate for outgoing packets.
Err [#s]	Average error rate of total transmission (rxErr[#s] + txErr[#s]).

NFS Metric Set

Contains top 32 NFS ordered by round trip time. This metric set is collected once every 30 seconds.

Table 3-5 NFS Metric Set

Metric Name (units)	Description
op [#s]	Number of read/write operations issued to a filesystem per second.
bytes [#sec]	Number of bytes read/write per second from a filesystem.
rtt [s]	This is the duration from the time that the client's kernel sends the RPC request until the time it receives the reply.
exe [s]	This is the duration from that NFS client does the RPC request to its kernel until the RPC request is completed, this includes the RTT time above.
retrains [%]	This is the retransmission's frequency in percentage.

Protocol Metric Set

Contains specific metrics for protocol groups TCP, UDP, and IP. Metric values are cumulative since the system starts.

Table 3-6 TCP Metric Set

Metric Name (units)	Description
failedConnErr [#]	Number of times that TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times that TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
estResetErr [#]	Number of times that TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
segRetransErr [#]	Total number of TCP segments retransmitted.
rxSeg [#]	Total number of TCP segments received on TCP layer.
txSeg [#]	Total number of TCP segments sent from TCP layer.

Table 3-7 UDP Metric Set

Metric Name (units)	Description
unkPortErr [#]	Total number of received datagrams for which there was no application at the destination port.

Table 3-7 (Cont.) UDP Metric Set

Metric Name (units)	Description
rxErr [#]	Number of received datagrams that could not be delivered for reasons other than the lack of an application at the destination port.
rxPkt [#]	Total number of packets received.
txPkt [#]	Total number of packets sent.

Table 3-8 IP Metric Set

Metric Name (units)	Description
ipHdrErr [#]	Number of input datagrams discarded due to errors in their IPv4 headers.
addrErr [#]	Number of input datagrams discarded because the IPv4 address in their IPv4 header's destination field was not a valid address to be received at this entity.
unkProtoErr [#]	Number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
reasFailErr [#]	Number of failures detected by the IPv4 reassembly algorithm.
fragFailErr [#]	Number of IPv4 discarded datagrams due to fragmentation failures.
rxPkt [#]	Total number of packets received on IP layer.
txPkt [#]	Total number of packets sent from IP layer.

Filesystem Metric Set

Contains metrics for filesystem utilization. Collected only for **GRID_HOME** filesystem.

Table 3-9 Filesystem Metric Set

Metric Name (units)	Description
mount	Mount point.
type	Filesystem type, for example, etx4 .
tag	Filesystem tag, for example, GRID_HOME .
total [KB]	Total amount of space (KB).
used [KB]	Amount of used space (KB).
avbl [KB]	Amount of available space (KB).
used [%]	Percentage of used space.
ifree [%]	Percentage of free file nodes.

System Metric Set

Contains a summarized metric set of critical system resource utilization.

Table 3-10 CPU Metrics

Metric Name (units)	Description
pCpus [#]	Number of physical processing units in the system.
Cores [#]	Number of cores for all CPUs in the system.
vCpus [#]	Number of logical processing units in the system.
cpuHt	CPU Hyperthreading enabled (Y) or disabled (N).
osName	Name of the operating system.
chipName	Name of the chip of the processing unit.
system [%]	Percentage of CPUs utilization that occurred while executing at the system level (kernel).
user [%]	Percentage of CPUs utilization that occurred while executing at the user level (application).
usage [%]	Total CPU utilization (system[%] + user[%]).
nice [%]	Percentage of CPUs utilization occurred while executing at the user level with NICE priority.
ioWait [%]	Percentage of time that the CPUs were idle during which the system had an outstanding disk I/O request.
Steal [%]	Percentage of time spent in involuntary wait by the virtual CPUs while the hypervisor was servicing another virtual processor.
cpuQ [#]	Number of processes waiting in the run queue within the current sample interval.
loadAvg1	Average system load calculated over time of one minute.
loadAvg5	Average system load calculated over of time of five minutes.
loadAvg15	Average system load calculated over of time of 15 minutes. High load averages imply that a system is overloaded; many processes are waiting for CPU time.
Intr [#s]	Number of interrupts occurred per second in the system.
ctxSwitch [#s]	Number of context switches that occurred per second in the system.

Table 3-11 Memory Metrics

Metric Name (units)	Description
totalMem [KB]	Amount of total usable RAM (KB).
freeMem [KB]	Amount of free RAM (KB).
avblMem [KB]	Amount of memory available to start a new process without swapping.
shMem [KB]	Memory used (mostly) by tmpfs .
swapTotal [KB]	Total amount of physical swap memory (KB).
swapFree [KB]	Amount of swap memory free (KB).

Table 3-11 (Cont.) Memory Metrics

Metric Name (units)	Description
swpIn [KB/s]	Average swap in rate within the current sample interval (KB/sec).
swpOut [KB/s]	Average swap-out rate within the current sample interval (KB/sec).
pgIn [#s]	Average page in rate within the current sample interval (pages/sec).
pgOut [#s]	Average page out rate within the current sample interval (pages/sec).
slabReclaim [KB]	The part of the slab that might be reclaimed such as caches.
buffer [KB]	Memory used by kernel buffers.
Cache [KB]	Memory used by the page cache and slabs.
bufferAndCache [KB]	Total size of buffer and cache (buffer[KB] + Cache[KB]).
hugePageTotal [#]	Total number of huge pages present in the system for the current sample interval.
hugePageFree [KB]	Total number of free huge pages in the system for the current sample interval.
hugePageSize [KB]	Size of one huge page in KB, depends on the operating system version. Typically the same for all samples for a particular host.

Table 3-12 Device Metrics

Metric Name (units)	Description
disks [#]	Number of disks configured in the system.
ioR [KB/s]	Aggregate read rate across all devices.
ioW [KB/s]	Aggregate write rate across all devices.
numIOs [#s]	Aggregate I/O operation rate across all devices.

Table 3-13 NFS Metrics

Metric Name (units)	Description
nfs [#]	Total NFS devices.

Table 3-14 Process Metrics

Metric Name (units)	Description
fds [#]	Number of open file structs in system.
procs [#]	Number of processes.
rtProcs [#]	Number of real-time processes.
procsInDState	Number of processes in uninterruptible sleep.
sysFdLimit [#]	System limit on a number of file structs.

Table 3-14 (Cont.) Process Metrics

Metric Name (units)	Description
procsOnCpu [#]	Number of processes currently running on CPU.
procsBlocked [#]	Number of processes waiting for some event/resource becomes available, such as for the completion of an I/O operation.

Process Aggregates Metric Set

Contains aggregated metrics for all processes by process groups.

Table 3-15 Process Aggregates Metric Set

Metric Name (units)	Description
DBBG	User Oracle Database background process group.
DBFG	User Oracle Database foreground process group.
MDBBG	MGMTDB background processes group.
MDBFG	MGMTDB foreground processes group.
ASMBG	ASM background processes group.
ASMFG	ASM foreground processes group.
IOXBG	IOS background processes group.
IOXFG	IOS foreground processes group.
APXBG	APX background processes group.
APXFG	APX foreground processes group.
CLUST	Clusterware processes group.
OTHER	Default group.

For each group, the below metrics are aggregated to report a group summary.

Metric Name (units)	Description
processes [#]	Total number of processes in the group.
cpu [%]	Aggregated CPU utilization.
rss [KB]	Aggregated resident set size.
shMem [KB]	Aggregated shared memory usage.
thrds [#]	Aggregated thread count.
fds [#]	Aggregated open file-descriptor.
cpuWeight [%]	Contribution of the group in overall CPU utilization of the machine.

3.4 Detecting Component Failures and Self-healing Autonomously

Improved ability to detect component failures and self-heal autonomously improves business continuity.

Cluster Health Monitor introduces a new diagnostic feature that identifies critical component events that indicate pending or actual failures and provides recommendations for corrective action. These actions may sometimes be performed autonomously. Such events and actions are then captured and admins are notified through components such as Oracle Trace File Analyzer.

Terms Associated with Diagnosability

CHMDiag: CHMDiag is a python daemon managed by `osysmond` that listens for events and takes actions. Upon receiving various events/actions, CHMDiag validates them for correctness, does flow control, and schedules the actions for runs. CHMDiag monitors each action to its completion, and kills an action if it takes longer than pre-configured time specific to that action.

This JSON file describes all events/actions and their respective attributes. All events/actions have uniquely identifiable IDs. This file also contains various configurable properties for various actions/events. CHMDiag loads this file during its startup.

CRFE API: CRFE API is used by all C clients to send events to CHMDiag. This API is used by internal clients like components (RDBMS/CSS/GIPC) to publish events/actions.

This API also provides support for both synchronous and asynchronous publication of events. Asynchronous publication of events is done through a background thread which will be shared by all CRFE API clients within a process.

CHMDIAG_BASE: This directory resides in `ORACLEB_BASE/hostname/crf/chmdiag`. This directory path contains following directories, which are populated or managed by CHMDiag.

- **ActionsResults:** Contains all results for all of the invoked actions with a subdirectory for each action.
- **EventsLog:** Contains a log of all the events/actions received by CHMDiag and the location of their respective action results. These log files are also auto-rotated after reaching a fixed size.
- **CHMDiagLog:** Contains CHMDiag daemon logs. Log files are auto-rotated and once they reach a specific size. Logs should have sufficient debug information to diagnose any problems that CHMDiag could run into.
- **Config:** Contains a run sub-directory for CHMDiag process pid file management.

New commands to query, collect, and describe CHMDiag events/actions sent by various components:

- **oclumon chmdiag description:** Use the `oclumon chmdiag description` command to get a detailed description of all the supported events and actions.
- **oclumon chmdiag query:** Use the `oclumon chmdiag query` command to query CHMDiag events/actions sent by various components and generate an HTML or a text report.
- **oclumon chmdiag collect:** Use the `oclumon chmdiag collect` command to collect all events/actions data generated by CHMDiag into the specified output directory location.

Related Topics

- [oclumon chmdiag description](#)
Use the `oclumon chmdiag description` command to get a detailed description of all the supported events and actions.
- [oclumon chmdiag query](#)
Use the `oclumon chmdiag query` command to query CHMDiag events/actions sent by various components and generate an HTML or a text report.
- [oclumon chmdiag collect](#)
Use the `oclumon chmdiag collect` command to collect all events/actions data generated by CHMDiag into the specified output directory location. This command will primarily be used by Oracle Trace File Analyzer to collect all events/actions that fall within a problematic window.

3.5 Using Cluster Health Monitor from Enterprise Manager Cloud Control

Histograms presented in real-time and historical modes enable you to understand precisely what was happening at the time of degradation or failure.

The metric data from Cluster Health Monitor is available in graphical display within Enterprise Manager Cloud Control. Complete cluster views of this data are accessible from the cluster target page. Selecting the **Cluster Health Monitoring** menu item from the **Cluster** menu presents a log-in screen prompting for the Cluster Health Monitor credentials. There is a fixed EMUSER and the password is user-specified. Once the credentials are saved, you then can view Cluster Health Monitor data for the last day in overview format for the entire cluster. Metric categories are CPU, Memory, and Network.

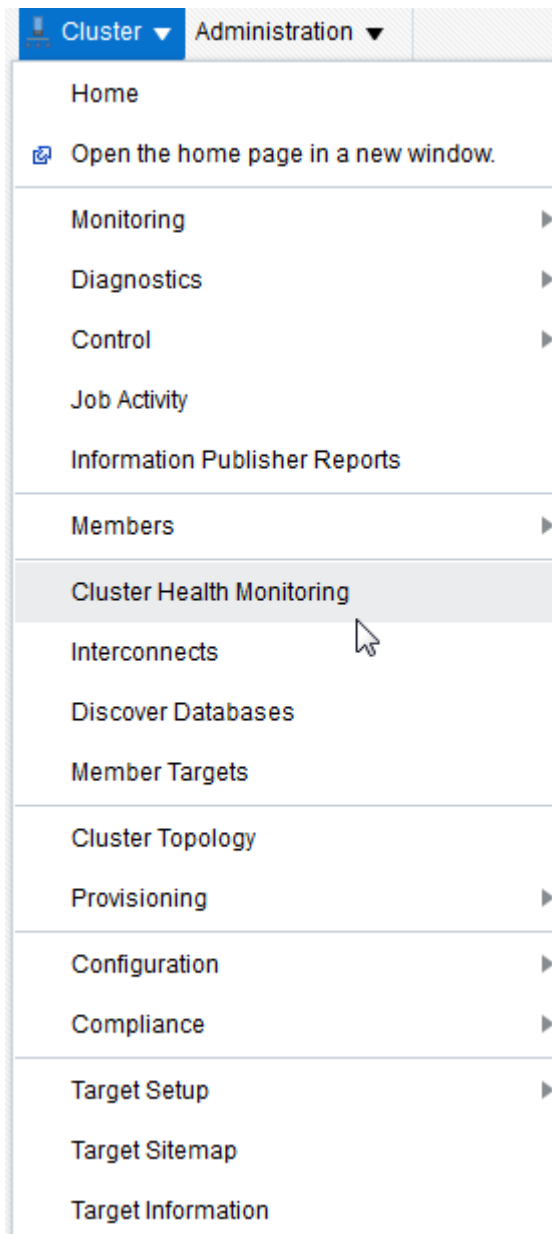
Each category is able to be separately display in greater detail showing more metrics. For example, selecting CPU results in cluster graphs detailing CPU System Usage, CPU User Usage, and CPU Queue Length. From any cluster view, you can select individual node views to more closely examine performance of a single server. As in the case of CPU, the performance of each core is displayed. Move your cursor along the graph to see a tool-tip displaying the numerical values and time stamp of that point.

Besides examining the performance of the current day, you can also review historical data. The amount of historical data is governed by the retention time configured in the Cluster Health Monitor repository in the Grid Infrastructure Management Repository and defaults to 72 hours. This view is selectable at any time by using the **View Mode** drop-down menu and selecting **Historical**. A previous date can then be entered or selected from a pop-up calendar that has dates where data is available bolded. Selecting **Show Chart** then displays the associated metrics graphs.

To view Cluster Health Monitor data:

1. Log in to Enterprise Manager Cloud Control.
2. Select the Cluster Target you want to view.
3. From the **Cluster** drop-down list, select the **Cluster Health Monitoring** option.

Figure 3-1 EMCC - Cluster Health Monitoring



4. Enter Cluster Health Monitor login credentials.
5. From the **View Mode** drop-down list, select the **Real Time** option to view the current data.

By default, EMCC displays the **Overview** of resource utilization. You can filter by **CPU**, **Memory**, and **Network** by selecting an appropriate option from the **Select Chart Type** drop-down list.

While viewing CPU and Network metric graphs, click a node name on the legend to view more details.

Figure 3-2 Cluster Health Monitoring - Real Time Data




6. From the **View Mode** drop-down list, select the **Historical** option to view data for the last 24 hours.
 - To filter historical data by date, select a day on the **Select Date** calendar control and then click **Show Chart**.

By default, EMCC displays the **Overview** of resource utilization. You can filter by **CPU**, **Memory**, and **Network** by selecting an appropriate option from the **Select Chart Type** drop-down list.

While viewing CPU and Network metric graphs, click a node name on the legend to view more details.

Figure 3-3 Cluster Health Monitoring - Historical Data



4

Monitoring System Metrics for Cluster Nodes

This chapter explains the methods to monitor Oracle Clusterware.

Oracle recommends that you use Oracle Enterprise Manager to monitor everyday operations of Oracle Clusterware.

Cluster Health Monitor monitors the complete technology stack, including the operating system, ensuring smooth cluster operations. Both the components are enabled, by default, for any Oracle cluster. Oracle strongly recommends that you use both the components. Also, monitor Oracle Clusterware-managed resources using the Clusterware resource activity log.

- [Monitoring Oracle Clusterware with Oracle Enterprise Manager](#)
Use Oracle Enterprise Manager to monitor the Oracle Clusterware environment.
- [Monitoring Oracle Clusterware with Cluster Health Monitor](#)
You can use the OCLUMON command-line tool to interact with Cluster Health Monitor.
- [Using the Cluster Resource Activity Log to Monitor Cluster Resource Failures](#)
The cluster resource activity log provides precise and specific information about a resource failure, separate from diagnostic logs.

4.1 Monitoring Oracle Clusterware with Oracle Enterprise Manager

Use Oracle Enterprise Manager to monitor the Oracle Clusterware environment.

When you log in to Oracle Enterprise Manager using a client browser, the **Cluster Database Home** page appears where you can monitor the status of both Oracle Database and Oracle Clusterware environments. Oracle Clusterware monitoring includes the following details:

- Current and historical Cluster Health Monitor data in Oracle Enterprise Manager on the cluster target
- Notifications if there are any VIP relocations
- Status of the Oracle Clusterware on each node of the cluster using information obtained through the Cluster Verification Utility (CVU)
- Notifications if node applications (`nodeapps`) start or stop
- Notification of issues in the Oracle Clusterware alert log for the Oracle Cluster Registry, voting file issues (if any), and node evictions

The **Cluster Database Home** page is similar to a single-instance Database Home page. However, on the Cluster Database Home page, Oracle Enterprise Manager displays the system state and availability. The system state and availability includes a summary about alert messages and job activity, and links to all the database and Oracle Automatic Storage Management (Oracle ASM) instances. For example, track problems with services on the cluster including when a service is not running on all the preferred instances or when a service response time threshold is not being met.

Use the Oracle Enterprise Manager **Interconnects** page to monitor the Oracle Clusterware environment. The Interconnects page displays the following details:

- Public and private interfaces on the cluster
- Overall throughput on the private interconnect
- Individual throughput on each of the network interfaces
- Error rates (if any)
- Load contributed by database instances on the interconnect
- Notifications if a database instance is using public interface due to misconfiguration
- Throughput contributed by individual instances on the interconnect

All the information listed earlier is also available as collections that have a historic view. The historic view is useful with cluster cache coherency, such as when diagnosing problems related to cluster wait events. Access the Interconnects page by clicking the **Interconnect** tab on the Cluster Database home page.

Also, the Oracle Enterprise Manager **Cluster Database Performance** page provides a quick glimpse of the performance statistics for a database. Statistics are rolled up across all the instances in the cluster database in charts. Using the links next to the charts, you can get more specific information and perform any of the following tasks:

- Identify the causes of performance issues
- Decide whether resources must be added or redistributed
- Tune your SQL plan and schema for better optimization
- Resolve performance issues

The charts on the Cluster Database Performance page include the following:

- **Chart for Cluster Host Load Average:** The **Cluster Host Load Average** chart in the Cluster Database Performance page shows potential problems that are outside the database. The chart shows maximum, average, and minimum load values for available nodes in the cluster for the previous hour.
- **Chart for Global Cache Block Access Latency:** Each cluster database instance has its own buffer cache in its System Global Area (SGA). Using Cache Fusion, Oracle RAC environments logically combine buffer cache of each instance to enable the database instances to process data as if the data resided on a logically combined, single cache.
- **Chart for Average Active Sessions:** The **Average Active Sessions** chart in the Cluster Database Performance page shows potential problems inside the database. Categories, called wait classes, show how much of the database is using a resource, such as CPU or disk I/O. Comparing CPU time to wait time helps to determine how much of the response time is consumed with useful work rather than waiting for resources that are potentially held by other processes.
- **Chart for Database Throughput:** The **Database Throughput** charts summarize any resource contention that appears in the Average Active Sessions chart, and also show how much work the database is performing on behalf of the users or applications. The **Per Second** view shows the number of transactions compared to the number of logons, and the amount of physical reads compared to the redo size for each second. The **Per Transaction** view shows the amount of physical

reads compared to the redo size for each transaction. Logons is the number of users that are logged on to the database.

In addition, the **Top Activity** drop-down menu on the **Cluster Database Performance** page enables you to see the activity by wait events, services, and instances. In addition, you can see the details about SQL/sessions by going to a prior point in time by moving the slider on the chart.

4.2 Monitoring Oracle Clusterware with Cluster Health Monitor

You can use the OCLUMON command-line tool to interact with Cluster Health Monitor.

OCLUMON is included with Cluster Health Monitor. You can use it to query the Cluster Health Monitor repository to display node-specific metrics for a specified time period. You can also use OCLUMON to perform miscellaneous administrative tasks, such as the following:

- Changing the debug levels with the `oclumon debug` command
- Querying the version of Cluster Health Monitor with the `oclumon version` command
- Viewing the collected information in the form of a node view using the `oclumon dumpnodeview` command
- Changing the metrics database size using the `oclumon manage` command

Related Topics

- [OCLUMON Command Reference](#)
Use the command-line tool to query the Cluster Health Monitor repository to display node-specific metrics for a specific time period.

4.3 Using the Cluster Resource Activity Log to Monitor Cluster Resource Failures

The cluster resource activity log provides precise and specific information about a resource failure, separate from diagnostic logs.

If an Oracle Clusterware-managed resource fails, then Oracle Clusterware logs messages about the failure in the **cluster resource activity log** located in the Grid Infrastructure Management Repository. Failures can occur as a result of a problem with a resource, a hosting node, or the network. The cluster resource activity log provides a unified view of the cause of resource failure.

Writes to the cluster resource activity log are tagged with an activity ID and any related data gets the same parent activity ID, and is nested under the parent data. For example, if Oracle Clusterware is running and you run the `crsctl stop clusterware -all` command, then all activities get activity IDs, and related activities are tagged with the same parent activity ID. On each node, the command creates sub-IDs under the parent IDs, and tags each of the respective activities with their corresponding activity ID. Further, each resource on the individual nodes creates sub-IDs based on the parent ID, creating a hierarchy of activity IDs. The hierarchy of activity IDs enables you to analyze the data to find specific activities.

For example, you may have many resources with complicated dependencies among each other, and with a database service. On Friday, you see that all of the resources are running on one node but when you return on Monday, every resource is on a different node, and you want to know why. Using the `crsctl query calog` command, you can query the cluster

resource activity log for all activities involving those resources and the database service. The output provides a complete flow and you can query each sub-ID within the parent service failover ID, and see, specifically, what happened and why.

You can query any number of fields in the cluster resource activity log using filters. For example, you can query all the activities written by specific operating system users such as `root`. The output produced by the `crsctl query calog` command can be displayed in either a tabular format or in XML format.

The cluster resource activity log is an adjunct to current Oracle Clusterware logging and alert log messages.



Note:

Oracle Clusterware does not write messages that contain security-related information, such as log-in credentials, to the cluster activity log.

Use the following commands to manage and view the contents of the cluster resource activity log:

Part III

Automatic Problem Solving

Some situations can be automatically resolved with tools in the Autonomous Health Framework.

- [Resolving Database and Database Instance Delays](#)
Blocker Resolver preserves the database performance by resolving delays and keeping the resources available.

5

Resolving Database and Database Instance Delays

Blocker Resolver preserves the database performance by resolving delays and keeping the resources available.

- [Blocker Resolver Architecture](#)
Blocker Resolver autonomously runs as a `DIA0` task within the database.
- [Optional Configuration for Blocker Resolver](#)
You can adjust the sensitivity, and control the size and number of the log files used by Blocker Resolver.
- [Blocker Resolver Diagnostics and Logging](#)
Blocker Resolver autonomously resolves delays and continuously logs the resolutions in the database alert logs and the diagnostics in the trace files.

Related Topics

- [Introduction to Blocker Resolver](#)
Blocker Resolver is an Oracle Real Application Clusters (Oracle RAC) environment feature that autonomously resolves delays and keeps the resources available.

5.1 Blocker Resolver Architecture

Blocker Resolver autonomously runs as a `DIA0` task within the database.

Blocker Resolver works in the following three phases:

- **Detect:** In this phase, Blocker Resolver collects the data on all the nodes and detects the sessions that are waiting for the resources held by another session.
- **Analyze:** In this phase, Blocker Resolver analyzes the sessions detected in the **Detect** phase to determine if the sessions are part of a potential delay. If the sessions are suspected as delayed, Blocker Resolver then waits for a certain threshold time period to ensure that the sessions are delayed.
- **Verify:** In this phase, after the threshold time period is up, Blocker Resolver verifies that the sessions are delayed and selects a session that's causing the delay.

After selecting the session that's causing the delay, Blocker Resolver applies resolution methods on that session. If the chain of sessions or the delay resolves automatically, then Blocker Resolver does not apply delay resolution methods. However, if the delay does not resolve by itself, then Blocker Resolver resolves the delay by terminating the session that's causing the delay. If terminating the session fails, then Blocker Resolver terminates the process of the session. This entire process is autonomous and does not block resources for a long period and does not affect the performance.

For example, if a high rank session is included in the chain of delayed sessions, then Blocker Resolver expedites the termination of the session that's causing the delay. Termination of the session that's causing the delay prevents the high rank session from waiting too long and helps to maintain performance objective of the high rank session.

5.2 Optional Configuration for Blocker Resolver

You can adjust the sensitivity, and control the size and number of the log files used by Blocker Resolver.

Note:

The `DBMS_HANG_MANAGER` package is deprecated in Oracle Database 23c. Use `DBMS_BLOCKER_RESOLVER` instead. The `DBMS_HANG_MANAGER` package provides a method of changing some configuration parameters and constraints to address session issues. This package is being replaced with `DBMS_BLOCKER_RESOLVER`. `DBMS_HANG_MANAGER` can be removed in a future release.

Sensitivity

If Blocker Resolver detects a delay, then Blocker Resolver waits for a certain threshold time period to ensure that the sessions are delayed. Change threshold time period by using `DBMS_BLOCKER_RESOLVER` to set the `sensitivity` parameter to either `Normal` or `High`. If the `sensitivity` parameter is set to `Normal`, then Blocker Resolver waits for the default time period. However, if the sensitivity is set to `High`, then the time period is reduced by 50%.

By default, the `sensitivity` parameter is set to `Normal`. To set Blocker Resolver sensitivity, run the following commands in SQL*Plus as SYS user:

- To set the `sensitivity` parameter to `Normal`:

```
exec dbms_blocker_resolver.set(dbms_blocker_resolver.sensitivity,  
dbms_blocker_resolver.sensitivity_normal);
```

- To set the `sensitivity` parameter to `High`:

```
exec dbms_blocker_resolver.set(dbms_blocker_resolver.sensitivity,  
dbms_blocker_resolver.sensitivity_high);
```

Size of the Trace Log File

The Blocker Resolver logs detailed diagnostics of the delays in the trace files with `_base_` in the file name. Change the size of the trace files in bytes with the `base_file_size_limit` parameter. Run the following command in SQL*Plus, for example, to set the trace file size limit to 100 MB:

```
exec  
dbms_blocker_resolver.set(dbms_blocker_resolver.base_file_size_limit,  
104857600);
```

Number of Trace Log Files

The base Blocker Resolver trace files are part of a trace file set. Change the number of trace files in trace file set with the `base_file_set_count` parameter. Run the following command in SQL*Plus, for example, to set the number of trace files in trace file set to 6:

```
exec dbms_blocker_resolver.set(dbms_blocker_resolver.base_file_set_count,6);
```

By default, `base_file_set_count` parameter is set to 5.

5.3 Blocker Resolver Diagnostics and Logging

Blocker Resolver autonomously resolves delays and continuously logs the resolutions in the database alert logs and the diagnostics in the trace files.

Blocker Resolver logs the resolutions in the database alert logs as Automatic Diagnostic Repository (ADR) incidents with incident code `ORA-32701`.

You also get detailed diagnostics about the delay detection in the trace files. Trace files and alert logs have file names starting with `database_instance_dia0_`.

- The trace files are stored in the `$ ADR_BASE/diag/rdbms/database name/database instance/incident/incdir_XXXXXX` directory
- The alert logs are stored in the `$ ADR_BASE/diag/rdbms/database name/database instance/trace` directory

Example 5-1 Blocker Resolver Trace File for a Local Instance

This example shows an example of the output you see for Blocker Resolver for the local database instance

```
Trace Log File .../oracle/log/diag/rdbms/hm1/hm11/incident/incdir_111/
hm11_dia0_11111_i111.trc
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production
...
*** 2016-07-16T12:39:02.715475-07:00
HM: Hang Statistics - only statistics with non-zero values are listed

          current number of active sessions 3
          current number of hung sessions 1
instance health (in terms of hung sessions) 66.67%
          number of cluster-wide active sessions 9
          number of cluster-wide hung sessions 5
          cluster health (in terms of hung sessions) 44.45%

*** 2016-07-16T12:39:02.715681-07:00
Resolvable Hangs in the System
```

Hang ID	Hang Type	Hang Status	Root Inst Num	Chain Root Sess	Total #hung Sess	Hang Conf	Hang Span	Resolution Action
1	HANG	RSLNPEND	3	44	3	5	HIGH GLOBAL	Terminate Process

Hang Resolution Reason: Although hangs of this root type are typically self-resolving, the previously ignored hang was automatically resolved.

Example 5-2 Error Message in the Alert Log Indicating a Delayed Session

This example shows an example of a Blocker Resolver alert log on the primary instance

```
2016-07-16T12:39:02.616573-07:00
Errors in file .../oracle/log/diag/rdbms/hml/hml/trace/
hml_dia0_i1111.trc (incident=1111):
ORA-32701: Possible hangs up to hang ID=1 detected
Incident details in: .../oracle/log/diag/rdbms/hml/hml/incident/
inmdir_1111/hml_dia0_11111_i1111.trc
2016-07-16T12:39:02.674061-07:00
DIA0 requesting termination of session sid:44 with serial # 23456
(ospid:34569) on instance 3
    due to a GLOBAL, HIGH confidence hang with ID=1.
    Hang Resolution Reason: Although hangs of this root type are
    typically
    self-resolving, the previously ignored hang was automatically
    resolved.
DIA0: Examine the alert log on instance 3 for session termination
status of hang with ID=1.
```

Example 5-3 Error Message in the Alert Log Showing a Session Delay Resolved by Blocker Resolver

This example shows an example of a Blocker Resolver alert log on the local instance for resolved delays

```
2016-07-16T12:39:02.707822-07:00
Errors in file .../oracle/log/diag/rdbms/hml/hml1/trace/
hml1_dia0_11111.trc (incident=169):
ORA-32701: Possible hangs up to hang ID=1 detected
Incident details in: .../oracle/log/diag/rdbms/hml/hml1/incident/
inmdir_169/hml1_dia0_30676_i169.trc
2016-07-16T12:39:05.086593-07:00
DIA0 terminating blocker (ospid: 30872 sid: 44 ser#: 23456) of hang
with ID = 1
    requested by master DIA0 process on instance 1
    Hang Resolution Reason: Although hangs of this root type are
    typically
    self-resolving, the previously ignored hang was automatically
    resolved.
    by terminating session sid:44 with serial # 23456 (ospid:34569)
...
DIA0 successfully terminated session sid:44 with serial # 23456
(ospid:34569) with status 0.
```

Part IV

Appendixes

- [OCLUMON Command Reference](#)
Use the command-line tool to query the Cluster Health Monitor repository to display node-specific metrics for a specific time period.
- [Managing the Cluster Resource Activity Log](#)
Oracle Clusterware stores logs about resource failures in the cluster resource activity log, which is located in the Grid Infrastructure Management Repository.
- [chactl Command Reference](#)
The Oracle Cluster Health Advisor commands enable the Oracle Grid Infrastructure user to administer basic monitoring functionality on the targets.
- [Behavior Changes, Deprecated and Desupported Features](#)
Review information about changes, deprecations, and desupports.

A

OCLUMON Command Reference

Use the command-line tool to query the Cluster Health Monitor repository to display node-specific metrics for a specific time period.

Use OCLUMON to perform miscellaneous administrative tasks, such as changing the debug levels, querying the version of Cluster Health Monitor, and changing the metrics database size.

- [oclumon debug](#)
Use the `oclumon debug` command to set the log level for the Cluster Health Monitor services.
- [oclumon dumpnodeview](#)
Use the `oclumon dumpnodeview` command to view log information from the system monitor service in the form of a node view.
- [oclumon dumpnodeview local](#)
Use the `oclumon dumpnodeview local` command to view log information from the system monitor service in the form of a node view.
- [oclumon manage](#)
Use the `oclumon manage` command to view and change configuration information from the system monitor service.
- [oclumon version](#)
Use the `oclumon version` command to obtain the version of Cluster Health Monitor that you are using.
- [oclumon chmdiag description](#)
Use the `oclumon chmdiag description` command to get a detailed description of all the supported events and actions.
- [oclumon chmdiag query](#)
Use the `oclumon chmdiag query` command to query CHMDiag events/actions sent by various components and generate an HTML or a text report.
- [oclumon chmdiag collect](#)
Use the `oclumon chmdiag collect` command to collect all events/actions data generated by CHMDiag into the specified output directory location. This command will primarily be used by Oracle Trace File Analyzer to collect all events/actions that fall within a problematic window.

A.1 oclumon debug

Use the `oclumon debug` command to set the log level for the Cluster Health Monitor services.

Syntax

```
oclumon debug [log daemon module:log_level] [version]
```

Parameters

Table A-1 oclumon debug Command Parameters

Parameter	Description
<code>log daemon</code> <code>module:log_level</code>	Use this option change the log level of daemons and daemon modules. Supported daemons are: osysmond ologgerd client all Supported daemon modules are: osysmond: CRFMOND, CRFM, and allcomp ologgerd: CRFLOGD, CRFLDREP, CRFM, and allcomp client: OCLUMON, CRFM, and allcomp all: allcomp Supported <code>log_level</code> values are 0, 1, 2, and 3.
<code>version</code>	Use this option to display the versions of the daemons.

Example A-1 oclumon debug

The following example sets the log level of the system monitor service (`osysmond`):

```
$ oclumon debug log osysmond CRFMOND:3
```

The following example displays the versions of the daemons:

```
$ oclumon debug version
```

```
Cluster Health Monitor (OS), Release 20.0.0.0.0
Version          : 20.3.0.0.0
NODEVIEW Version : 19.03
Label Date       : 200116
```

A.2 oclumon dumpnodeview

Use the `oclumon dumpnodeview` command to view log information from the system monitor service in the form of a node view.

Usage Notes**Note:**

The `oclumon dumpnodeview` commands work only if GIMR or MGMTDB is configured and is in ONLINE state.

A node view is a collection of all metrics collected by Cluster Health Monitor for a node at a point in time. Cluster Health Monitor attempts to collect metrics every five seconds on every node. Some metrics are static while other metrics are dynamic.

A node view consists of eight views when you display verbose output:

- **SYSTEM:** Lists system metrics such as CPU COUNT, CPU USAGE, and MEM USAGE
- **TOP CONSUMERS:** Lists the top consuming processes in the following format:

```
metric_name: 'process_name(process_identifier) utilization'
```

- **CPUS:** Lists statistics for each CPU
- **PROCESSES:** Lists process metrics such as PID, name, number of threads, memory usage, and number of file descriptors
- **DEVICES:** Lists device metrics such as disk read and write rates, queue length, and wait time per I/O
- **NICS:** Lists network interface card metrics such as network receive and send rates, effective bandwidth, and error rates
- **FILESYSTEMS:** Lists file system metrics, such as total, used, and available space
- **PROTOCOL ERRORS:** Lists any protocol errors

Generate a summary report that only contains the SYSTEM and TOP CONSUMERS views.

Syntax

```
oculumon dumpnodeview [-allnodes | -n node1...]
[-last duration | -s timestamp -e timestamp]
[-i interval]
[-v | [-system [-v2]]]
[-process]
[-procag]
[-device]
[-filesystem]
[-nic]
[-advm]
[-protocols]
[-cpu]
[-topconsumer]
[-asminst_db]
[-nfs]]
[-format format type]
[-dir directory [-append]]
```

Parameters

Table A-2 oclumon dumpnodeview Command Parameters

Parameter	Description
-allnodes	Use this option to dump the node views of all the nodes in the cluster.
-n node1 node2	Specify one node or several nodes in a space-delimited list for which you want to dump the node view.

Table A-2 (Cont.) oclumon dumpnodeview Command Parameters

Parameter	Description
-last " <i>duration</i> "	Use this option to specify a time, given in HH24:MM:SS format surrounded by double quotation marks (""), to retrieve the last metrics. For example: "23:05:00"
-s " <i>time_stamp</i> " -e " <i>time_stamp</i> "	Use the -s option to specify a time stamp from which to start a range of queries and use the -e option to specify a time stamp to end the range of queries. Specify time in YYYY-MM-DD HH24:MM:SS format surrounded by double quotation marks (""). For example: "2011-05-10 23:05:00"
	Note: Specify these two options together to obtain a range.
-i <i>interval</i>	Specify a collection interval, in five-second increments.
-v	Displays verbose node view output.
-system[-v2], -topconsumer, -process, -cpu, -procag, -device, -filesystem, -nic, -protocols, -advm, -asminst_db, -nfs	Dumps each specified node view parts.
-v2	Dumps nodeview part output with version2 outlook. Currently available with the "system" part.
-format " <i>format type</i> "	Specify the output format. " <i>format type</i> " can be legacy, tabular, json, or csv. The default format is mostly tabular with legacy for node view parts with only one row.
-dir <i>directory</i>	Dumps the node view to the files in the directory that you specify. Specify the -append option to append the files of the current to the existing files. If you do not specify -append, then the command overwrites the existing files, if present. For example, the command oclumon dumpnodeview -dir <i>dir_name</i> dumps the data in the specified directory. If this command is run twice, it overwrites the data dumped by the previous run. Running the command with -append, for example, oclumon dumpnodeview -dir <i>dir_name</i> -append, appends the data of the current run with the previous one in the specified directory.

Table A-2 (Cont.) ocumon dumpnodeview Command Parameters

Parameter	Description
-procag	Outputs the process of the node view, aggregated by category: <ul style="list-style-type: none"> • DBBG (DB backgrounds) • DBFG (DB foregrounds) • CLUST (Cluster) • OTHER (other processes) Note: -procag is currently available only on Linux, Solaris, and AIX. It is not supported on Microsoft Windows systems.
-h	Displays online help for the ocumon dumpnodeview command.

Usage Notes

- In certain circumstances, data can be delayed for some time before the command replays the data.
For example, the `crsctl stop cluster -all` command can cause data delay. After running `crsctl start cluster -all`, it may take several minutes before `ocumon dumpnodeview` shows any data collected during the interval.
- The default is to continuously dump node views. To stop continuous display, use Ctrl+C on Linux and Microsoft Windows.
- Both the local system monitor service (`osysmond`) and the cluster logger service (`ologgerd`) must be running to obtain node view dumps.
- The `ocumon dumpnodeview` command displays only 127 CPUs of the CPU core, omitting a CPU at random from the list.

Metric Descriptions

This section includes descriptions of the metrics in each of the seven views that comprise a node view listed in the following tables.

Table A-3 ocumon dumpnodeview SYSTEM View Metric Descriptions

Metric	Description
#pcpus	Number of physical CPUs.
#cores	Number of CPU cores in the system.
#vcpus	Number of logical compute units.
cpuht	CPU hyperthreading enabled (Y) or disabled (N).
chipname	Name of the CPU vendor.
cpu	Average CPU utilization per processing unit within the current sample interval (%).
cpuusage	Total CPU usage = <code>cpusystem + cpuuser + cpunice</code> Percentage of over all CPU cores. 100% indicates that all cores are spent for that metric.
cpusystem	CPU used by processes in kernel mode.
cpuuser	CPU used by normal processes in user mode.

Table A-3 (Cont.) ocumon dumpnodeview SYSTEM View Metric Descriptions

Metric	Description
cpunice	CPU used by "niced" processes (low priority).
cpuioawait	CPU waiting for I/O.
cpusteal	Virtual CPU waiting for physical CPU to be freed by other VM.
cpuq	Number of processes waiting in the run queue within the current sample interval.
physmemfree	Amount of free RAM (KB).
physmemtotal	Amount of total usable RAM (KB).
shmem	Shared memory.
mcache	Amount of physical RAM used for file buffers plus the amount of physical RAM used as cache memory (KB). On Microsoft Windows systems, this is the number of bytes currently being used by the file system cache. Note: This metric is not available on Oracle Solaris.
swapfree	Amount of swap memory free (KB)
swaptotal	Total amount of physical swap memory (KB)
hugepagetotal	Total size of huge in KB Note: This metric is not available on Solaris or Microsoft Windows systems.
hugepagefree	Free size of huge page in KB Note: This metric is not available on Solaris or Microsoft Windows systems.
hugepagesize	Smallest unit size of huge page Note: This metric is not available on Solaris or Microsoft Windows systems.
ior	Average total disk read rate within the current sample interval (KB per second).
iow	Average total disk write rate within the current sample interval (KB per second).
ios	Average disk I/O operation rate within the current sample interval (I/O operations per second).
swpin	Average swap in rate within the current sample interval (KB per second). Note: This metric is not available on Microsoft Windows systems.
swpout	Average swap out rate within the current sample interval (KB per second). Note: This metric is not available on Microsoft Windows systems.
pgin	Average page in rate within the current sample interval (pages per second).
pgout	Average page out rate within the current sample interval (pages per second).
netr	Average total network receive rate within the current sample interval (KB per second).
netw	Average total network send rate within the current sample interval (KB per second).

Table A-3 (Cont.) ocumon dumpnodeview SYSTEM View Metric Descriptions

Metric	Description
procs	Number of processes.
procsoncpu	The current number of processes running on the CPU.
#procs_blocked	Number of processes currently blocked waiting for I/O.
rtprocs	Number of real-time processes.
rtprocsoncpu	The current number of real-time processes running on the CPU.
#fds	Number of open file descriptors. <i>or</i> Number of open handles on Microsoft Windows.
#sysfdlimit	System limit on the number of file descriptors. Note: This metric is not available on either Solaris or Microsoft Windows systems.
#disks	Number of disks.
#nics	Number of network interface cards.
nicErrors	Average total network error rate within the current sample interval (errors per second).
#nfs	Number of network file system.
loadavg1 loadavg5 loadavg15	Load average (average number of jobs in the run queue or waiting for disk I/O) of the last 1, 5, 15 minutes.

Table A-4 ocumon dumpnodeview PROCESSES View Metric Descriptions

Metric	Description
name	The name of the process executable.
pid	The process identifier assigned by the operating system.
ppid	PID of the parent process. For example, if process 1 spawns process 2, then ppid of process 2 is pid of process 1.
cumulative_cpu	The total amount of CPU time this process is scheduled to run since it started. The total amount of CPU time spent for this process so far is measured in micro seconds.
#procfdlimit	Limit on number of file descriptors for this process. Note: This metric is not available on Microsoft Windows, AIX, and HP-UX systems.
cpuusage	Process CPU utilization (%). Note: The utilization value can be up to 100 times the number of processing units.
vmem	Process virtual memory usage (KB).
privmem	Process private memory usage (KB).
shmem, shm, and sharedmem	Process shared memory usage (KB). Note: This metric is not available on Microsoft Windows, Solaris, and AIX systems. It is supported only on Linux systems.

Table A-4 (Cont.) oclumon dumpnodeview PROCESSES View Metric Descriptions

Metric	Description
workingset	Working set of a program (KB) Note: This metric is only available on Microsoft Windows.
#fd	Number of file descriptors open by this process. <i>or</i> Number of open handles by this process on Microsoft Windows.
#threads	Number of threads created by this process.
priority	The process priority.
nice	The nice value of the process. Note: This metric is not applicable to Microsoft Windows systems.
state	The state of the process. Note: This metric is not applicable to Microsoft Windows systems.

Table A-5 oclumon dumpnodeview DEVICES View Metric Descriptions

Metric	Description
ior	Average disk read rate within the current sample interval (KB per second).
iow	Average disk write rate within the current sample interval (KB per second).
ios	Average disk I/O operation rate within the current sample interval (I/O operations per second)
qlen	Number of I/O requests in <code>WAIT</code> state within the current sample interval.
wait	Average wait time per I/O within the current sample interval (msec).
type	If applicable, identifies what the device is used for. Possible values are: <ul style="list-style-type: none"> • SWAP • SYS • OCR • ASM • VOTING

Table A-6 oclumon dumpnodeview NICS View Metric Descriptions

Metric	Description
netrr	Average network receive rate within the current sample interval (KB per second).
netwr	Average network sent rate within the current sample interval (KB per second).
neteff	Average effective bandwidth within the current sample interval (KB per second)

Table A-6 (Cont.) oclumon dumpnodeview NICS View Metric Descriptions

Metric	Description
nicerrors	Average error rate within the current sample interval (errors per second).
pktsin	Average incoming packet rate within the current sample interval (packets per second).
pktsout	Average outgoing packet rate within the current sample interval (packets per second).
errsin	Average error rate for incoming packets within the current sample interval (errors per second).
errout	Average error rate for outgoing packets within the current sample interval (errors per second).
indiscarded	Average drop rate for incoming packets within the current sample interval (packets per second).
outdiscarded	Average drop rate for outgoing packets within the current sample interval (packets per second).
inunicast	Average packet receive rate for unicast within the current sample interval (packets per second).
type	Whether PUBLIC or PRIVATE.
inonunicast	Average packet receive rate for multi-cast (packets per second).
latency	Estimated latency for this network interface card (msec).

Table A-7 oclumon dumpnodeview FILESYSTEMS View Metric Descriptions

Metric	Description
total	Total amount of space (KB).
mount	Mount point.
type	File system type, whether local file system, NFS, or other.
used	Amount of used space (KB).
available	Amount of available space (KB).
used%	Percentage of used space (%)
ifree%	Percentage of free file nodes (%).
	Note: This metric is not available on Microsoft Windows systems.

Table A-8 oclumon dumpnodeview PROTOCOL ERRORS View Metric Descriptions

Metric	Description
IPHdrErr	Number of input datagrams discarded due to errors in the IPv4 headers of the datagrams.
IPAddrErr	Number of input datagrams discarded because the IPv4 address in their IPv4 header's destination field was not a valid address to be received at this entity.
IPUnkProto	Number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.

Table A-8 (Cont.) oclumon dumpnodeview PROTOCOL ERRORS View Metric Descriptions

Metric	Description
IPReasFail	Number of failures detected by the IPv4 reassembly algorithm.
IPFragFail	Number of IPv4 discarded datagrams due to fragmentation failures.
TCPFailedConn	Number of times that TCP connections have made a direct transition to the CLOSED state from either the SYN-SENT state or the SYN-RCVD state, plus the number of times that TCP connections have made a direct transition to the LISTEN state from the SYN-RCVD state.
TCPEstRst	Number of times that TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state.
TCPRetraSeg	Total number of TCP segments retransmitted.
UDPUnkPort	Total number of received UDP datagrams for which there was no application at the destination port.
UDPRcvErr	Number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port.

Table A-9 oclumon dumpnodeview CPUS View Metric Descriptions

Metric	Description
<i>cpuid</i>	Virtual CPU.
<i>sys-usage</i>	CPU usage in system space.
<i>user-usage</i>	CPU usage in user space.
<i>nice</i>	Value of NIC for a specific CPU.
<i>usage</i>	CPU usage for a specific CPU.
<i>iowait</i>	CPU wait time for I/O operations.

Example A-2 dumpnodeview -n

The following example dumps node views from `node1`, `node2`, and `node3` collected over the last 12 hours:

```
$ oclumon dumpnodeview -n node1 node2 node3 -last "12:00:00"
```

The following example displays node views from all nodes collected over the last 15 minutes at a 30-second interval:

```
$ oclumon dumpnodeview -allnodes -last "00:15:00" -i 30
```

Example A-3 dumphnodeview -format csv

The following example shows how to use the option `-format csv` to output content in comma-separated values file format:

```
# oclumon dumphnodeview -format csv

dumphnodeview: Node name not given. Querying for the local host

-----
Node: node1 Clock: '2016-09-02 11.18.00-0700' SerialNo:310668
-----

SYSTEM:
"#pcpus", "#cores", "#vcpus", "cpuht", "chipname", "cpuusage[%]", "cpusys[%]", "cpuuser[%]",
"cpunice[%]", "cpuiowait[%]", "cpusteal[%]", "cpuq", "physmemfree[KB]", "physmemtotal[KB]",
"mcache[KB]", "swapfree[KB]", "swaptotal[KB]", "hugepagetotal", "hugepagefree", "hugepagesize",
"ior[KB/S]", "iow[KB/S]", "ios[#/S]", "swpin[KB/S]", "swpout[KB/S]", "pgin[#/S]", "pgout[#/S]",
"netr[KB/S]", "netw[KB/S]", "#procs", "#procscpu", "#procs_blocked", "#rtprocs", "#rtprocscpu",
"#fds", "#sysfdlimit", "#disks", "#nics", "loadavg1", "loadavg5", "loadavg15", "#nic
Errors"
2,12,24,Y,"Intel(R) Xeon(R) CPU X5670 @
2.93GHz",68.66,5.40,63.26,0.00,0.00,0.00,0,820240,
73959636,61520568,4191424,4194300,0,0,
2048,143,525,64,0,0,0,279,600.888,437.070,951,24,0,58,N/
A,33120,6815744,13,5,19.25,17.67,16.09,0

TOPCONSUMERS:
"topcpu", "topprivmem", "topshm", "topfd", "topthread"
"java(25047) 225.44", "java(24667) 1008360", "ora_lms1_prod_1(28913)
4985464", "polkit-gnome-au(20730) 1038", "java(2734) 209"
```

Example A-4 dumphnodeview -procag

The following example shows how to output node views, aggregated by category: DBBG (DB backgrounds), DBFG (DB foregrounds), CLUST (Cluster), and OTHER (other processes).

```
# oclumon dumphnodeview -procag

-----
Node: node1 Clock: '2016-09-02 11.14.15-0700' SerialNo:310623
-----

PROCESS AGGREGATE:
cpuusage[%]  privatemem[KB]  maxshm[KB]  #threads  #fd
#processes  category  sid
      0.62      45791348      4985200      187      10250
183      DBBG      prod_1
      0.52      29544192      3322648      191      10463
187      DBBG      webdb_1
      17.81      8451288      967924      22      511
```


22	DBFG	webdb_1			
	75.94		34930368	1644492	64
1067		64	DBFG	prod_1	
	3.42		3139208	120256	480
3556		25	CLUST		
	1.66		1989424	16568	1110
4040		471	OTHER		

Example A-5 Node View Output

```
-----
Node: rwsak10 Clock: '2016-05-08 02.11.25-0800' SerialNo:155631
-----
```

SYSTEM:

```
#pcpus: 2 #vcpus: 24 cpuht: Y chipname: Intel(R) cpu: 1.23 cpuq: 0
physmemfree: 8889492 physmemtotal: 74369536 mcache: 55081824 swapfree:
18480404
swaptotal: 18480408 hugepagetotal: 0 hugepagefree: 0 hugepagesize:
2048 ior: 132
iow: 236 ios: 23 swpin: 0 swpout: 0 pgin: 131 pgout: 235 netr: 72.404
netw: 97.511 procs: 969 procsoncpu: 6 rtprocs: 62 rtprocsoncpu N/A
#fds: 32640
#sysfdlimit: 6815744 #disks: 9 #nics: 5 nicErrors: 0
```

TOP CONSUMERS:

```
topcpu: 'osysmond.bin(30981) 2.40' topprivmem: 'oraagent.bin(14599)
682496'
topshm: 'ora_dbw2_oss_3(7049) 2156136' topfd: 'ocssd.bin(29986) 274'
topthread: 'java(32255) 53'
```

CPUS:

```
cpu18: sys-2.93 user-2.15 nice-0.0 usage-5.8 iowait-0.0 steal-0.0
.
.
.
```

PROCESSES:

```
name: 'osysmond.bin' pid: 30891 #procfdlimit: 65536 cpuusage: 2.40
privmem: 35808
shm: 81964 #fd: 119 #threads: 13 priority: -100 nice: 0 state: S
.
.
.
```

DEVICES:

```
sdi ior: 0.000 iow: 0.000 ios: 0 qlen: 0 wait: 0 type: SYS
sdal ior: 0.000 iow: 61.495 ios: 629 qlen: 0 wait: 0 type: SYS
.
.
.
```

NICS:

```
lo netrr: 39.935 netwr: 39.935 neteff: 79.869 nicerrors: 0 pktsin: 25
pktsout: 25 errsin: 0 errsout: 0 indiscarded: 0 outdiscarded: 0
inunicast: 25 innonunicast: 0 type: PUBLIC
eth0 netrr: 1.412 netwr: 0.527 neteff: 1.939 nicerrors: 0 pktsin: 15
pktsout: 4 errsin: 0 errsout: 0 indiscarded: 0 outdiscarded: 0
inunicast: 15 innonunicast: 0 type: PUBLIC latency: <1
```

FILESYSTEMS:

```
mount: / type: rootfs total: 563657948 used: 78592012 available: 455971824
used%: 14 ifree%: 99 GRID_HOME
.
.
.
```

PROTOCOL ERRORS:

```
IPHdrErr: 0 IPAddrErr: 0 IPUnkProto: 0 IPReasFail: 0 IPFragFail: 0
TCPFailedConn: 5197 TCPEstRst: 717163 TCPRetraSeg: 592 UDPUnkPort: 103306
UDPRcvErr: 70
```

A.3 oclumon dumpnodeview local

Use the `oclumon dumpnodeview local` command to view log information from the system monitor service in the form of a node view.

Usage Notes



Note:

The `oclumon dumpnodeview local` command has no dependency on GIMR or MGMTDB and can return Cluster Health Monitor data irrespective of GIMR being configured.

Syntax

```
dumpnodeview local [[((-system | -protocols | -v) |
  [(-cpu | -process | -procagg | -device | -nic | -filesystem | -nfs)
  [-detail] [-all] [-sort <metric_name>] [-filter <string>] [-head
  <rows_count>] [-i <seconds>]])
  [[(-s <start_time> -e <end_time>] | -last <duration>))] |
  [-inputDataDir <absolute_path> -logDir <absolute_path>]
  [-h]]
```

Parameters**Table A-10 oclumon dumpnodeview local Command Parameters**

Parameter	Description
local	Cluster Health Monitor local dump. Dumps metrics from node local Cluster Health Monitor repository.
-system	Dumps system metrics. For example: oclumon dumpnodeview local -system .
-cpu	Dumps CPU metrics. For example: oclumon dumpnodeview local -cpu .
-process	Dumps process metrics. For example: oclumon dumpnodeview local -process .
-procagg	Dumps process aggregate metrics. For example: oclumon dumpnodeview local -procagg .
-device	Dumps disk metrics. For example: oclumon dumpnodeview local -device .
-nic	Dumps network interface metrics. For example: oclumon dumpnodeview local -nic .
-filesystem	Dumps filesystem metrics. For example: oclumon dumpnodeview local -filesystem .

Table A-10 (Cont.) oclumon dumpnodeview local Command Parameters

Parameter	Description
-nfs	Dumps NFS metrics. For example: <pre>oclumon dumpnodeview local -nfs</pre>
-protocols	Dumps network protocol metrics, cumulative values from system start. For example: <pre>oclumon dumpnodeview local -protocols</pre>
-v	Displays verbose node view output. For example: <pre>oclumon dumpnodeview local -v</pre>
-h, --help	Displays the command-line help and exits.

Table A-11 oclumon dumpnodeview local Command Flags

Flag	Description
-detail	Use this option to dump detailed metrics. Applicable to the <code>-process</code> and <code>-nic</code> options. For example: <pre>oclumon dumpnodeview local -process -detail</pre>
-all	Use this option to dump the node views of all entries. Applicable to the <code>-process</code> option. For example: <pre>oclumon dumpnodeview local -process -all</pre>

Table A-11 (Cont.) oclumon dumpnodeview local Command Flags

Flag	Description
<code>-head <i>rows_count</i></code>	<p>Use this option to dump the node view of the specified number of metrics rows in the result. Applicable to the <code>-process</code> option. Default is set to 5.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -process -head 7</pre>
<code>-sort <i>metric_name</i></code>	<p>Use this option to sort based on the specified metric name, supported with the <code>-process</code>, <code>-device</code>, <code>-nic</code>, <code>-cpu</code>, <code>-procagg</code>, <code>-filesystem</code>, <code>-nfs</code> options.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -device -sort "ioR"</pre>
<code>-i <i>seconds</i></code>	<p>Display data separated by the specified interval in seconds. Must be a multiple of 5. Applicable to continuous mode query.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -device -i 5</pre>
<code>-filter <i>string</i></code>	<p>Use this option to search for a filter string in the Name column of the respective metric.</p> <p>For example, <code>-process -filter "ora"</code> will display the process metrics, which contain "ora" substring in their name.</p> <p>Supported with the <code>-process</code>, <code>-device</code>, <code>-nic</code>, <code>-cpu</code>, <code>-procagg</code>, <code>-filesystem</code>, <code>-nfs</code> options.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -process -filter "ora"</pre>

Table A-11 (Cont.) oclumon dumpnodeview local Command Flags

Flag	Description
<code>-show_all_sample_with_filter</code>	<p>All samples where filter doesn't matches will also show in the output . Can be used only with <code>-filter</code> option.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -filter filter_criteria - show_all_sample_with_filter</pre>

Table A-12 oclumon dumpnodeview local Command Log File Directories

Directory	Description
<code>-inputDataDir <i>absolute_dir_path</i></code>	<p>Specifies absolute path of the directory that contains JSON logs files.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -cpu - inputDataDir <i>absolute_path</i></pre>
<code>-logDir <i>absolute_log_dir_path</i></code>	<p>Specifies absolute path of the directory, which will contain the script run logs.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -cpu - inputDataDir <i>absolute_path</i> -logDir <i>absolute_log_dir_path</i></pre>

Table A-13 oclumon dumpnodeview local Command Historical Query Options

Flag	Description
-s <i>start_time</i> -e <i>end_time</i>	<p>Use the <code>-s</code> option to specify a time stamp from which to start a range of queries and use the <code>-e</code> option to specify a time stamp to end the range of queries.</p> <p>Specify time in the <code>YYYY-MM-DD HH24:MM:SS</code> format surrounded by double quotation marks (<code>"</code>).</p> <p>Specify these two options together to obtain a range.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -cpu -s "2019-07-10 03:40:25" -e "2019-07-10 03:45:25"</pre>
-last <i>duration</i>	<p>Use this option to specify a time, given in <code>HH24:MM:SS</code> format surrounded by double quotation marks (<code>"</code>), to retrieve the last metrics.</p> <p>Specifying <code>"00:45:00"</code> will dump metrics for the last 45 minutes.</p> <p>For example:</p> <pre>oclumon dumpnodeview local -nic -last "00:45:00"</pre>

A.4 oclumon manage

Use the `oclumon manage` command to view and change configuration information from the system monitor service.

Syntax

```
oclumon manage -repos {{changeretentiontime time} | {changerepossize memory_size}} | -get {key1 [key2 ...] | alllogger [-details] | mylogger [-details]}
```

Parameters

Table A-14 oclumon manage Command Parameters

Parameter	Description
<code>-repos {{changeretentiontime time} {changerepossize memory_size}}</code>	<p>The <code>-repos</code> flag is required to specify the following Cluster Health Monitor repository-related options:</p> <ul style="list-style-type: none"> <code>changeretentiontime time</code>: Use this option to confirm that there is sufficient tablespace to hold the amount of Cluster Health Monitor data that can be accumulated in a specific amount of time. Note: This option <i>does not</i> change retention time. <code>changerepossize memory_size</code>: Use this option to change the Cluster Health Monitor repository space limit to a specified number of MB Caution: If you decrease the space limit of the Cluster Health Monitor repository, then all data collected before the resizing operation is permanently deleted.
<code>-get key1 [key2 ...]</code>	<p>Use this option to obtain Cluster Health Monitor repository information using the following keywords:</p> <ul style="list-style-type: none"> <code>resize</code>: Size of the Cluster Health Monitor repository, in seconds <code>reppath</code>: Directory path to the Cluster Health Monitor repository <code>master</code>: Name of the parent node <code>allogger</code>: Special key to obtain a list of all nodes running Cluster Logger Service <code>mylogger</code>: Special key to obtain the node running the Cluster Logger Service which is serving the current node <code>-details</code>: Use this option with <code>allogger</code> and <code>mylogger</code> for listing nodes served by the Cluster Logger Service <p>You can specify any number of keywords in a space-delimited list following the <code>-get</code> flag.</p>
<code>-h</code>	Displays online help for the <code>oclumon manage</code> command

Usage Notes

- The local system monitor service must be running to change the retention time of the Cluster Health Monitor repository.
- The Cluster Logger Service must be running to change the retention time of the Cluster Health Monitor repository.

Example A-6 oclumon manage

The following examples show commands and sample output:

```
$ oclumon manage -get MASTER
Master = node1
```

```
$ oclumon manage -get alllogger -details
Logger = node1
Nodes = node1,node2
```

```
$ oclumon manage -repos changeretentiontime 86400
```

```
$ oclumon manage -repos changerepossize 6000
```

A.5 oclumon version

Use the `oclumon version` command to obtain the version of Cluster Health Monitor that you are using.

Syntax

```
oclumon version
```

Example A-7 oclumon version

This command produces output similar to the following:

```
Cluster Health Monitor (OS), Release 20.0.0.0.0
Version : 20.3.0.0.0
```

A.6 oclumon chmdiag description

Use the `oclumon chmdiag description` command to get a detailed description of all the supported events and actions.

Syntax

```
oclumon chmdiag description
[-h]
[-f {text,html} | --format {text,html}]
[--outdir OUTDIR]
```

Parameters

Table A-15 oclumon chmdiag description Command Parameters

Parameter	Description
<code>-f {text,html} --format {text,html}</code>	Specify to generate a report in the Text or HTML format. Note that you need to specify the <code>--outdir</code> if you need the report in the HTML format. HTML report is an interactive <code>html</code> page, which has support for easy navigation and filtering of data. It also has built-in help, which describes the events/actions.
<code>--outdir OUTDIR</code>	Specify the location of the directory where the description data (text/html) will be written.

A.7 oclumon chmdiag query

Use the `oclumon chmdiag query` command to query CHMDiag events/actions sent by various components and generate an HTML or a text report.

Syntax

```
oclumon chmdiag query
[-h]
[-f {text,html} | --format {text,html}]
[--outdir OUTDIR]
[--summary]
[-s S]
[-e E]
[--last LAST]
[--evtids EVTIDS [EVTIDS,...]]
[--actids ACTIDS [ACTIDS,...]]
[--compids COMPIDS [COMPIDS,...]]
[--inputdir INPUTDIR]
[--logdir LOGDIR]
```

Parameters

Table A-16 oclumon chmdiag query Command Parameters

Parameter	Description
<code>-f {text,html} --format {text,html}</code>	Specify to generate a report in the Text or HTML format. Note that you need to specify the <code>--outdir</code> if you need the report in the HTML format. HTML report is an interactive <code>html</code> page, which has support for easy navigation and filtering of data. It also has built-in help, which describes the events/actions.
<code>--outdir OUTDIR</code>	Specify the location of the directory where the queried data report (text/html) will be written.
<code>--summary</code>	Specify to display only the summary of the query result.

Table A-16 (Cont.) oclumon chmdiag query Command Parameters

Parameter	Description
-s <i>S</i>	Specify the start timestamp for the query in the YYYY-MM-DDTHH24:MI:SS-z format.
-e <i>E</i>	Specify the end timestamp for the query in the YYYY-MM-DDTHH24:MI:SS-z format.
--last <i>LAST</i>	Specify to query logs for a specific duration. The duration can be specified as nd h m s. Where: <ul style="list-style-type: none"> • d stands for Days • h stands for Hours • m stands for Minutes • s stands for Seconds • n can be a floating number For example, --last 1d2h3m1s.
--evtids <i>EVTIDS</i>	Specify to query only the events with the specified event IDs.
--actids <i>ACTIDS</i>	Specify to query only the actions with the specified action IDs.
--compids <i>COMPIDS</i>	Specify to query all events/actions generated by the specified set of components.
--inputdir <i>INPUTDIR</i>	Specify the absolute path of the directory that contains chmdiag logs files.
--logdir <i>LOGDIR</i>	Specify the absolute path of the directory, which will contain the script run logs.

Action Status can be one of the following:

- **Succeeded:** Action completed with success return code.
- **Failed:** Action returned a non-success return code
- **NotRunExceededMaxLimit:** Action has not been run as the maximum number of allowed concurrent runs for this action has been exceeded.
- **Purged:** Action results have been purged, that is, the results have been recycled based on the set retention policy for this Action ID.
- **ExceededMaxRunTime:** This action has not completed within the configured maximum run time and hence the action has been killed. Whatever the output that gets generated by the command will be captured in the `cmd.out` file.

Example A-8 oclumon chmdiag query

```
oclumon chmdiag query -f html --outdir /outdir/path/
```

```
oclumon chmdiag query --last 1h30m
```

```
oclumon chmdiag query -s 2018-06-20T10:30:00 -e 2018-06-20T11:30:00
```

A.8 oclumon chmdiag collect

Use the `oclumon chmdiag collect` command to collect all events/actions data generated by CHMDiag into the specified output directory location. This command will primarily be used by Oracle Trace File Analyzer to collect all events/actions that fall within a problematic window.

Syntax

```
oclumon chmdiag collect
[-h]
[--outdir OUTDIR]
[--gzip]
[-s S]
[-e E]
[--last LAST]
[--evtids EVTIDS [EVTIDS,...]]
[--actids ACTIDS [ACTIDS,...]]
[--compids COMPIDS [COMPIDS,...]]
[--inputdir INPUTDIR]
[--logdir LOGDIR]
```

Parameters

Table A-17 oclumon chmdiag collect Command Parameters

Parameter	Description
<code>--outdir <i>OUTDIR</i></code>	Specify the location of the directory where the collected data will be written.
<code>-gzip</code>	Specify to compress the output directory data into a <code>gzip</code> file.
<code>-s <i>S</i></code>	Specify the start timestamp for the collection in the <code>YYYY-MM-DDTHH24:MI:SS-z</code> format.
<code>-e <i>E</i></code>	Specify the end timestamp for the collection in the <code>YYYY-MM-DDTHH24:MI:SS-z</code> format.
<code>--last <i>LAST</i></code>	Specify to collect logs for a specific duration. The duration can be specified as <code>nd h m s</code> . Where: <ul style="list-style-type: none"> <code>d</code> stands for Days <code>h</code> stands for Hours <code>m</code> stands for Minutes <code>s</code> stands for Seconds <code>n</code> can be a floating number For example, <code>--last 1d2h3m1s</code> .
<code>--evtids <i>EVTIDS</i></code>	Specify to collect only the events with the specified event IDs.
<code>--actids <i>ACTIDS</i></code>	Specify to collect only the actions with the specified action IDs.
<code>--compids <i>COMPIDS</i></code>	Specify to collect all events/actions generated by the specified set of components.
<code>--inputdir <i>INPUTDIR</i></code>	Specify the absolute path of the directory that contains <code>chmdiag</code> logs files.

Table A-17 (Cont.) ocumon chmdiag collect Command Parameters

Parameter	Description
<code>--logdir LOGDIR</code>	Specify the absolute path of the directory, which will contain the script run logs.

Example A-9 ocumon chmdiag collect

```
ocumon chmdiag collect --outdir /outdir/path/
```

```
ocumon chmdiag collect --last 1.5h --outdir /outdir/path/
```

B

Managing the Cluster Resource Activity Log

Oracle Clusterware stores logs about resource failures in the cluster resource activity log, which is located in the Grid Infrastructure Management Repository.

Failures can occur as a result of a problem with a resource, a hosting node, or the network.

The cluster resource activity log provides precise and specific information about a resource failure, separate from diagnostic logs. The cluster resource activity log also provides a unified view of the cause of resource failure.

Use the following commands to manage and view the contents of the cluster resource activity log:

- [crsctl query calog](#)
Query the cluster resource activity logs matching specific criteria.
- [crsctl get calog maxsize](#)
To store Oracle Clusterware-managed resource activity information, query the maximum space allotted to the cluster resource activity log.
- [crsctl get calog retentiontime](#)
Query the retention time of the cluster resource activity log.
- [crsctl set calog maxsize](#)
Configure the maximum amount of space allotted to store Oracle Clusterware-managed resource activity information.
- [crsctl set calog retentiontime](#)
Configure the retention time of the cluster resource activity log.

B.1 crsctl query calog

Query the cluster resource activity logs matching specific criteria.

Syntax

```
crsctl query calog [-aftertime "timestamp"] [-beforetime "timestamp"]  
  [-duration "time_interval" | -follow] [-filter "filter_expression"]  
  [-fullfmt | -xmlfmt]
```

Parameters

Table B-1 crsctl query calog Command Parameters

Parameter	Description
-aftertime "timestamp"	<p>Displays the activities logged after a specific time.</p> <p>Specify the timestamp in the YYYY-MM-DD HH24:MI:SS [.FF] [TZH:TZM] or YYYY-MM-DD or HH24:MI:SS [.FF] [TZH:TZM] format.</p> <p>TZH and TZM stands for time zone hour and minute, and FF stands for microseconds.</p> <p>If you specify [TZH:TZM], then the crsctl command assumes UTC as time zone. If you do not specify [TZH:TZM], then the crsctl command assumes the local time zone of the cluster node from where the crsctl command is run.</p> <p>Use this parameter with -beforetime to query the activities logged at a specific time interval.</p>
-beforetime "timestamp"	<p>Displays the activities logged before a specific time.</p> <p>Specify the timestamp in the YYYY-MM-DD HH24:MI:SS [.FF] [TZH:TZM] or YYYY-MM-DD or HH24:MI:SS [.FF] [TZH:TZM] format.</p> <p>TZH and TZM stands for time zone hour and minute, and FF stands for microseconds.</p> <p>If you specify [TZH:TZM], then the crsctl command assumes UTC as time zone. If you do not specify [TZH:TZM], then the crsctl command assumes the local time zone of the cluster node from where the crsctl command is run.</p> <p>Use this parameter with -aftertime to query the activities logged at a specific time interval.</p>
-duration "time_interval" - follow	<p>Use -duration to specify a time interval that you want to query when you use the -aftertime parameter.</p> <p>Specify the timestamp in the DD HH:MM:SS format.</p> <p>Use -follow to display a continuous stream of activities as they occur.</p>
-filter "filter_expression"	<p>Query any number of fields in the cluster resource activity log using the -filter parameter.</p> <p>To specify multiple filters, use a comma-delimited list of filter expressions surrounded by double quotation marks ("").</p>
-fullfmt -xmlfmt	To display cluster resource activity log data, choose full or XML format.

Cluster Resource Activity Log Fields

Query any number of fields in the cluster resource activity log using the -filter parameter.

Table B-2 Cluster Resource Activity Log Fields

Field	Description	Use Case
timestamp	The time when the cluster resource activities were logged.	Use this filter to query all the activities logged at a specific time. This is an alternative to - aftertime, -beforetime, and - duration command parameters.
writer_process_id	The ID of the process that is writing to the cluster resource activity log.	Query only the activities spawned by a specific process.
writer_process_name	The name of the process that is writing to the cluster resource activity log.	When you query a specific process, CRSD returns all the activities for a specific process.
writer_user	The name of the user who is writing to the cluster resource activity log.	Query all the activities written by a specific user.
writer_group	The name of the group to which a user belongs who is writing to the cluster resource activity log.	Query all the activities written by users belonging to a specific user group.
writer_hostname	The name of the host on which the cluster resource activity log is written.	Query all the activities written by a specific host.
writer_clustername	The name of the cluster on which the cluster resource activity log is written.	Query all the activities written by a specific cluster.
nls_product	The product of the NLS message, for example, CRS, ORA, or srvm.	Query all the activities that have a specific product name.
nls_facility	The facility of the NLS message, for example, CRS or PROC.	Query all the activities that have a specific facility name.
nls_id	The ID of the NLS message, for example 42008.	Query all the activities that have a specific message ID.
nls_field_count	The number of fields in the NLS message.	Query all the activities that correspond to NLS messages with more than, less than, or equal to nls_field_count command parameters.
nls_field1	The first field of the NLS message.	Query all the activities that match the first parameter of an NLS message.
nls_field1_type	The type of the first field in the NLS message.	Query all the activities that match a specific type of the first parameter of an NLS message.
nls_format	The format of the NLS message, for example, Resource '%s' has been modified.	Query all the activities that match a specific format of an NLS message.
nls_message	The entire NLS message that was written to the cluster resource activity log, for example, Resource 'ora.cvu' has been modified.	Query all the activities that match a specific NLS message.

Table B-2 (Cont.) Cluster Resource Activity Log Fields

Field	Description	Use Case
actid	The unique activity ID of every cluster activity log.	Query all the activities that match a specific ID. Also, specify only partial actid and list all activities where the actid is a subset of the activity ID.
is_planned	Confirms if the activity is planned or not. For example, if a user issues the command <code>crsctl stop crs</code> on a node, then the stack stops and resources bounce. Running the <code>crsctl stop crs</code> command generates activities and logged in the calog. Since this is a planned action, the <code>is_planned</code> field is set to true (1). Otherwise, the <code>is_planned</code> field is set to false (0).	Query all the planned or unplanned activities.
onbehalfof_user	The name of the user on behalf of whom the cluster activity log is written.	Query all the activities written on behalf of a specific user.
entity_isoraentity	Confirms if the entity for which the calog activities are being logged is an oracle entity or not. If a resource, such as <code>ora.***</code> , is started or stopped, for example, then all those activities are logged in the cluster resource activity log. Since <code>ora.***</code> is an Oracle entity, the <code>entity_isoraentity</code> field is set to true (1). Otherwise the <code>entity_isoraentity</code> field is set to false (0).	Query all the activities logged by Oracle or non-Oracle entities.
entity_type	The type of the entity, such as <code>server</code> , for which the cluster activity log is written.	Query all the activities that match a specific entity.
entity_name	The name of the entity, for example, <code>foo</code> for which the cluster activity log is written.	Query all the cluster activities that match a specific entity name.
entity_hostname	The name of the host, for example, <code>node1</code> , associated with the entity for which the cluster activity log is written.	Query all the cluster activities that match a specific host name.
entity_clustername	The name of the cluster, for example, <code>cluster1</code> associated with the entity for which the cluster activity log is written.	Query all the cluster activities that match a specific cluster name.

Usage Notes

Combine simple filters into expressions called expression filters using Boolean operators.

Enclose timestamps and time intervals in double quotation marks ("").

Enclose the filter expressions in double quotation marks ("").

Enclose the values that contain parentheses or spaces in single quotation marks (").

If no matching records are found, then the Oracle Clusterware Control (CRSCTL) utility displays the following message:

```
CRS-40002: No activities match the query.
```

Examples

Examples of filters include:

- "writer_user==root": Limits the display to only root user.
- "customer_data=='GEN_RESTART@SERVERNAME(rwsbi08)=StartCompleted~'" : Limits the display to customer_data that has the specified value
GEN_RESTART@SERVERNAME(node1)=StartCompleted~.

To query all the resource activities and display the output in full format:

```
$ crsctl query calog -fullfmt
```

```

----ACTIVITY START----
timestamp           : 2016-09-27 17:55:43.152000
writer_process_id   : 6538
writer_process_name : crsd.bin
writer_user         : root
writer_group        : root
writer_hostname     : node1
writer_clustername  : cluster1-mb1
customer_data       : CHECK_RESULTS=-408040060~
nls_product         : CRS
nls_facility        : CRS
nls_id              : 2938
nls_field_count     : 1
nls_field1          : ora.cvu
nls_field1_type     : 25
nls_field1_len      : 0
nls_format          : Resource '%s' has been modified.
nls_message         : Resource 'ora.cvu' has been modified.
actid               : 14732093665106538/1816699/1
is_planned          : 1
onbehalfof_user     : grid
onbehalfof_hostname : node1
entity_isoraentity  : 1
entity_type         : resource
entity_name         : ora.cvu
entity_hostname     : node1
entity_clustername  : cluster1-mb1
----ACTIVITY END----

```

To query all the resource activities and display the output in XML format:

```
$ crsctl query calog -xmlfmt
```

```
<?xml version="1.0" encoding="UTF-8"?>
<activities>
  <activity>
    <timestamp>2016-09-27 17:55:43.152000</timestamp>
    <writer_process_id>6538</writer_process_id>
    <writer_process_name>crsd.bin</writer_process_name>
    <writer_user>root</writer_user>
    <writer_group>root</writer_group>
    <writer_hostname>node1</writer_hostname>
    <writer_clustername>cluster1-mbl</writer_clustername>
    <customer_data>CHECK_RESULTS=-408040060~</customer_data>
    <nls_product>CRS</nls_product>
    <nls_facility>CRS</nls_facility>
    <nls_id>2938</nls_id>
    <nls_field_count>1</nls_field_count>
    <nls_field1>ora.cvu</nls_field1>
    <nls_field1_type>25</nls_field1_type>
    <nls_field1_len>0</nls_field1_len>
    <nls_format>Resource '%s' has been modified.</nls_format>
    <nls_message>Resource 'ora.cvu' has been modified.</nls_message>
    <actid>14732093665106538/1816699/1</actid>
    <is_planned>1</is_planned>
    <onbehalfof_user>grid</onbehalfof_user>
    <onbehalfof_hostname>node1</onbehalfof_hostname>
    <entity_isoraentity>1</entity_isoraentity>
    <entity_type>resource</entity_type>
    <entity_name>ora.cvu</entity_name>
    <entity_hostname>node1</entity_hostname>
    <entity_clustername>cluster1-mbl</entity_clustername>
  </activity>
</activities>
```

To query resource activities for a two-hour interval after a specific time and display the output in XML format:

```
$ crsctl query calog -aftertime "2016-09-28 17:55:43" -duration "02:00:00" -xmlfmt
```

```
<?xml version="1.0" encoding="UTF-8"?>
<activities>
  <activity>
    <timestamp>2016-09-28 17:55:45.992000</timestamp>
    <writer_process_id>6538</writer_process_id>
    <writer_process_name>crsd.bin</writer_process_name>
    <writer_user>root</writer_user>
    <writer_group>root</writer_group>
    <writer_hostname>node1</writer_hostname>
    <writer_clustername>cluster1-mbl</writer_clustername>
    <customer_data>CHECK_RESULTS=1718139884~</customer_data>
    <nls_product>CRS</nls_product>
    <nls_facility>CRS</nls_facility>
```

```

<nls_id>2938</nls_id>
<nls_field_count>1</nls_field_count>
<nls_field1>ora.cvu</nls_field1>
<nls_field1_type>25</nls_field1_type>
<nls_field1_len>0</nls_field1_len>
<nls_format>Resource '%s' has been modified.</nls_format>
<nls_message>Resource 'ora.cvu' has been modified.</nls_message>
<actid>14732093665106538/1942009/1</actid>
<is_planned>1</is_planned>
<onbehalfof_user>grid</onbehalfof_user>
<onbehalfof_hostname>node1</onbehalfof_hostname>
<entity_isoraentity>1</entity_isoraentity>
<entity_type>resource</entity_type>
<entity_name>ora.cvu</entity_name>
<entity_hostname>node1</entity_hostname>
<entity_clustername>cluster1-mbl</entity_clustername>
</activity>
</activities>

```

To query resource activities at a specific time:

```
$ crsctl query calog -filter "timestamp=='2016-09-28 17:55:45.992000'"
```

```
2016-09-28 17:55:45.992000 : Resource 'ora.cvu' has been modified. :
14732093665106538/1942009/1 :
```

To query resource activities using filters `writer_user` and `customer_data`:

```
$ crsctl query calog -filter "writer_user==root AND customer_data==
'GEN_RESTART@SERVERNAME (node1)=StartCompleted~'" -fullfmt
```

or

```
$ crsctl query calog -filter "(writer_user==root) AND (customer_data==
'GEN_RESTART@SERVERNAME (node1)=StartCompleted~')" -fullfmt
```

```

----ACTIVITY START----
timestamp                : 2016-09-15 17:42:57.517000
writer_process_id        : 6538
writer_process_name      : crsd.bin
writer_user               : root
writer_group              : root
writer_hostname           : node1
writer_clustername        : cluster1-mbl
customer_data              : GEN_RESTART@SERVERNAME (rwsbi08)=StartCompleted~
nls_product               : CRS
nls_facility               : CRS
nls_id                    : 2938
nls_field_count           : 1
nls_field1                : ora.testdb.db
nls_field1_type           : 25
nls_field1_len            : 0

```

```
nls_format          : Resource '%s' has been modified.
nls_message         : Resource 'ora.devdb.db' has been modified.
actid               : 14732093665106538/659678/1
is_planned          : 1
onbehalfof_user     : oracle
onbehalfof_hostname : node1
entity_isoraentity  : 1
entity_type         : resource
entity_name         : ora.testdb.db
entity_hostname     : node1
entity_clustername  : cluster1-mb1
----ACTIVITY END----
```

To query all the calogs that were generated after UTC+08:00 time "2016-11-15 22:53:08":

```
$ crsctl query calog -aftertime "2016-11-15 22:53:08+08:00"
```

To query all the calogs that were generated after UTC-08:00 time "2016-11-15 22:53:08":

```
$ crsctl query calog -aftertime "2016-11-15 22:53:08-08:00"
```

To query all the calogs by specifying the timestamp with microseconds:

```
$ crsctl query calog -aftertime "2016-11-16 01:07:53.063000"
```

```
2016-11-16 01:07:53.558000 : Resource 'ora.cvu' has been modified. :
14792791129816600/2580/7 :
2016-11-16 01:07:53.562000 : Clean of 'ora.cvu' on 'rwsam02'
succeeded : 14792791129816600/2580/8 :
```

B.2 crsctl get calog maxsize

To store Oracle Clusterware-managed resource activity information, query the maximum space allotted to the cluster resource activity log.

Syntax

```
crsctl get calog maxsize
```

Parameters

The `crsctl get calog maxsize` command has no parameters.

Example

The following example returns the maximum space allotted to the cluster resource activity log to store activities:

```
$ crsctl get calog maxsize
```

CRS-6760: The maximum size of the Oracle cluster activity log is 1024 MB.

B.3 crsctl get calog retentiontime

Query the retention time of the cluster resource activity log.

Syntax

```
crsctl get calog retentiontime
```

Parameters

The `crsctl get calog retentiontime` command has no parameters.

Examples

The following example returns the retention time of the cluster activity log, in number of hours:

```
$ crsctl get calog retentiontime
```

CRS-6781: The retention time of the cluster activity log is 73 hours.

B.4 crsctl set calog maxsize

Configure the maximum amount of space allotted to store Oracle Clusterware-managed resource activity information.

Syntax

```
crsctl set calog maxsize maximum_size
```

Usage Notes

Specify a value, in MB, for the maximum size of the storage space that you want to allot to the cluster resource activity log.



Note:

If you reduce the amount of storage space, then the contents of the storage are lost.

Example

The following example sets maximum amount of space, to store Oracle Clusterware-managed resource activity information, to 1024 MB:

```
$ crsctl set calog maxsize 1024
```

B.5 crsctl set calog retentiontime

Configure the retention time of the cluster resource activity log.

Syntax

```
crsctl set calog retentiontime hours
```

Parameters

The `crsctl set calog retentiontime` command takes a number of hours as a parameter.

Usage Notes

Specify a value, in hours, for the retention time of the cluster resource activity log.

Examples

The following example sets the retention time of the cluster resource activity log to 72 hours:

```
$ crsctl set calog retentiontime 72
```

C

chactl Command Reference

The Oracle Cluster Health Advisor commands enable the Oracle Grid Infrastructure user to administer basic monitoring functionality on the targets.

- [chactl monitor](#)
Use the `chactl monitor` command to start monitoring all the instances of a specific Oracle Real Application Clusters (Oracle RAC) database using the current set model.
- [chactl unmonitor](#)
Use the `chactl unmonitor` command to stop monitoring all the instances of a specific database.
- [chactl status](#)
Use the `chactl status` command to check monitoring status of the running targets.
- [chactl config](#)
Use the `chactl config` command to list all the targets being monitored, along with the current model of each target.
- [chactl calibrate](#)
Use the `chactl calibrate` command to create a new model that has greater sensitivity and accuracy.
- [chactl query diagnosis](#)
Use the `chactl query diagnosis` command to return problems and diagnosis, and suggested corrective actions associated with the problem for specific cluster nodes or Oracle Real Application Clusters (Oracle RAC) databases.
- [chactl query model](#)
Use the `chactl query model` command to list all Oracle Cluster Health Advisor models or to view detailed information about a specific Oracle Cluster Health Advisor model.
- [chactl query repository](#)
Use the `chactl query repository` command to view the maximum retention time, number of targets, and the size of the Oracle Cluster Health Advisor repository.
- [chactl query calibration](#)
Use the `chactl query calibration` command to view detailed information about the calibration data of a specific target.
- [chactl remove model](#)
Use the `chactl remove model` command to delete an Oracle Cluster Health Advisor model along with the calibration data and metadata of the model from the Oracle Cluster Health Advisor repository.
- [chactl rename model](#)
Use the `chactl rename model` command to rename an Oracle Cluster Health Advisor model in the Oracle Cluster Health Advisor repository.
- [chactl export model](#)
Use the `chactl export model` command to export Oracle Cluster Health Advisor models.
- [chactl import model](#)
Use the `chactl import model` command to import Oracle Cluster Health Advisor models.

- [chactl set maxretention](#)
Use the `chactl set maxretention` command to set the maximum retention time for the diagnostic data.
- [chactl resize repository](#)
Use the `chactl resize repository` command to resize the tablespace of the Oracle Cluster Health Advisor repository based on the current retention time and the number of targets.

C.1 chactl monitor

Use the `chactl monitor` command to start monitoring all the instances of a specific Oracle Real Application Clusters (Oracle RAC) database using the current set model.

Oracle Cluster Health Advisor monitors all instances of this database using the same model assigned to the database.

Oracle Cluster Health Advisor uses Oracle-supplied gold model when you start monitoring a target for the first time. Oracle Cluster Health Advisor stores monitoring status of the target in the internal store. Oracle Cluster Health Advisor starts monitoring any new database instance when Oracle Cluster Health Advisor detects or redetects the new instance.

Syntax

```
chactl monitor database -db db_unique_name [-model model_name [-force]]
[-help]
```

```
chactl monitor cluster [-model model_name [-force]]
```

Parameters

Table C-1 chactl monitor Command Parameters

Parameter	Description
<code>db_unique_name</code>	Specify the name of the database.
<code>model_name</code>	Specify the name of the model.
<code>force</code>	Use the <code>-force</code> option to monitor with the specified model without stopping monitoring the target. Without the <code>-force</code> option, run <code>chactl unmonitor</code> first, and then <code>chactl monitor</code> with the model name.

Examples

- To monitor the *SalesDB* database using the *BlkFridayShopping* default model:

```
$ chactl monitor database -db SalesDB -model BlkFridayShopping
```

- To monitor the *InventoryDB* database using the *Nov2014* model:

```
$ chactl monitor database -db InventoryDB -model Nov2014
```

If you specify the *model_name*, then Oracle Cluster Health Advisor starts monitoring with the specified model and stores the model in the Oracle Cluster Health Advisor internal store.

If you use both the *-model* and *-force* options, then Oracle Cluster Health Advisor stops monitoring and restarts monitoring with the specified model.

- To monitor the *SalesDB* database using the *Dec2014* model:

```
$ chactl monitor database -db SalesDB -model Dec2014
```

- To monitor the *InventoryDB* database using the *Dec2014* model and the *-force* option:

```
$ chactl monitor database -db InventoryDB -model Dec2014 -force
```

Error Messages

Error: no CHA resource is running in the cluster.

Description: Returns when there is no hub or leaf node running the Oracle Cluster Health Advisor service.

Error: the database is not configured.

Description: Returns when the database is not found in either the Oracle Cluster Health Advisor configuration repository or as a CRS resource.

Error: input string "xc#? %" is invalid.

Description: Returns when the command-line cannot be parsed. Also displays the top-level help text.

Error: CHA is already monitoring target <dbname>.

Description: Returns when the database is already monitored.

C.2 chactl unmonitor

Use the `chactl unmonitor` command to stop monitoring all the instances of a specific database.

Syntax

```
chactl unmonitor database -db db_unique_name [-help]
```

Examples

To stop monitoring the *SalesDB* database:

```
$ chactl unmonitor database -db SalesDB
Database SalesDB is not monitored
```

C.3 chactl status

Use the `chactl status` command to check monitoring status of the running targets.

If you do not specify any parameters, then the `chactl status` command returns the status of all running targets.

The monitoring status of an Oracle Cluster Health Advisor target can be either `Monitoring` or `Not Monitoring`. The `chactl status` command shows four types of results and depends on whether you specify a target and `-verbose` option.

The `-verbose` option of the command also displays the monitoring status of targets contained within the specified target and the names of executing models of each printed target. The `chactl status` command displays targets with positive monitoring status only. The `chactl status` command displays negative monitoring status only when the corresponding target is explicitly specified on the command-line.

Syntax

```
chactl status {cluster|database [-db db_unique_name]} [-verbose][-help]
```

Examples

- To display the list of cluster nodes and databases being monitored:

```
#chactl status
Monitoring nodes rac1Node1, rac1Node2
Monitoring databases SalesDB, HRdb
```

Note:

A database is displayed with **Monitoring** status, if Oracle Cluster Health Advisor is monitoring one or more of the instances of the database, even if some of the instances of the database are not running.

- To display the status of Oracle Cluster Health Advisor:

```
$ chactl status
Cluster Health Advisor service is offline.
```

No target or the `-verbose` option is specified on the command-line. Oracle Cluster Health Advisor is not running on any node of the cluster.

- To display various Oracle Cluster Health Advisor monitoring states for cluster nodes and databases:

```
$ chactl status database -db SalesDB
Monitoring database SalesDB
```

```
$ chactl status database -db bogusDB
Not Monitoring database bogusDB
```

```
$ chactl status cluster
Monitoring nodes rac1,rac2
Not Monitoring node rac3
```

or

```
$ chactl status cluster
Cluster Health Advisor is offline
```

- To display the detailed Oracle Cluster Health Advisor monitoring status for the entire cluster:

```
$ chactl status -verbose
Monitoring node(s) racNd1, racNd2, racNd3, racNd4 using model MidSparc
```

```
Monitoring database HRdb2, Instances HRdb2I1, HRdb2I2 in server pool
SilverPool using model M6
Monitoring database HRdb, Instances HRdbI4, HRdbI6 in server pool
SilverPool using model M23
Monitoring database testHR, Instances inst3 on node racN7 using model
TestM13
Monitoring database testHR, Instances inst4 on node racN8 using model
TestM14
```

When the target is not specified and the `-verbose` option is specified, the `chactl status` command displays the status of the database instances and names of the models.

C.4 chactl config

Use the `chactl config` command to list all the targets being monitored, along with the current model of each target.

If the specified target is a multitenant container database (CDB) or a cluster, then the `chactl config` command also displays the configuration data status.

Syntax

```
chactl config {cluster|database -db db_unique_name }[-help]
```

Examples

To display the monitor configuration and the specified model of each target:

```
$ chactl config
Databases monitored: prodDB, hrDB

$ chactl config database -db prodDB
Monitor: Enabled
Model: GoldDB

$ chactl config cluster
Monitor: Enabled
Model: DEFAULT_CLUSTER
```

C.5 chactl calibrate

Use the `chactl calibrate` command to create a new model that has greater sensitivity and accuracy.

The user-generated models are effective for Oracle Real Application Clusters (Oracle RAC) monitored systems in your operating environment as the user-generated models use calibration data from the target. Oracle Cluster Health Advisor adds the user-generated model to the list of available models and stores the new model in the Oracle Cluster Health Advisor repository.

If a model with the same name exists, then overwrite the old model with the new one by using the `-force` option.

Key Performance and Workload Indicators

A set of metrics or Key Performance Indicators describe high-level constraints to the training data selected for calibration. This set consists of relevant metrics to describe performance goals and resource utilization bandwidth, for example, response times or CPU utilization.

The Key Performance Indicators are also operating system and database signals which are monitored, estimated, and associated with fault detection logic. Most of these Key Performance Indicators are also either predictors, that is, their state is correlated with the state of other signals, or predicted by other signals. The fact that the Key Performance Indicators correlate with other signals makes them useful as filters for the training or calibration data.

The Key Performance Indicators ranges are used in the `query calibrate` and `calibrate` commands to filter out data points.

The following Key Performance Indicators are supported for database:

- `CPUPERCENT` - CPU utilization - Percent
- `IOREAD` - Disk read - Mbyte/sec
- `DBTIMEPERCALL` - Database time per user call - usec/call

- IOWRITE - Disk write - Mbyte/sec
- IOTHROUGHPUT - Disk throughput - IO/sec

The following Key Performance Indicators are supported for cluster:

- CPUPERCENT - CPU utilization - Percent
- IOREAD - Disk read - Mbyte/sec
- IOWRITE - Disk write - Mbyte/sec
- IOTHROUGHPUT - Disk throughput - IO/sec

Syntax

```
chactl calibrate {cluster|database -db db_unique_name} -model model_name
[-force] [-timeranges 'start=time_stamp,end=time_stamp,...']
[-kpiiset 'name=kpi_name min=val max=val,...' ][-help]
```

Specify timestamp in the YYYY-MM-DD HH24:MI:SS format.

Examples

```
chactl calibrate database -db oracle -model weekday
-timeranges 'start=start=2016-09-09 16:00:00,end=2016-09-09 23:00:00'
```

```
chactl calibrate database -db oracle -model weekday
-timeranges 'start=start=2016-09-09 16:00:00,end=2016-09-09 23:00:00'
-kpiiset 'name=CPUPERCENT min=10 max=60'
```

Error Messages

Error: input string "xc#? %" is misconstructured

Description: Confirm if the given model name exists with warning: *model_name* already exists, please use [-force] message.

Error: *start_time* and/or *end_time* are misconstructured

Description: Input time specifiers are badly constructed.

Error: no sufficient calibration data exists for the specified period, please reselect another period

Description: Evaluator couldn't find enough calibration data.

C.6 chactl query diagnosis

Use the `chactl query diagnosis` command to return problems and diagnosis, and suggested corrective actions associated with the problem for specific cluster nodes or Oracle Real Application Clusters (Oracle RAC) databases.

Syntax

```
chactl query diagnosis [-cluster|-db db_unique_name] [-start time -end time] [-htmlfile file_name] [-help]
```

Specify date and time in the YYYY-MM-DD HH24:MI:SS format.

In the preceding syntax, you must consider the following points:

- If you do not provide any options, then the `chactl query diagnosis` command returns the current state of all monitored nodes and databases. The `chactl query diagnosis` command reports general state of the targets, for example, **ABNORMAL** by showing their diagnostic identifier, for example, Storage Bandwidth Saturation. This is a quick way to check for any **ABNORMAL** state in a database or cluster.
- If you provide a time option after the target name, then the `chactl query diagnosis` command returns the state of the specified target restricted to the conditions in the time interval specified. The compressed time series lists the identifiers of the causes for distinct incidents which occurred in the time interval, its start and end time.
- If an incident and cause recur in a specific time interval, then the problem is reported only once. The start time is the start time of the first occurrence of the incident and the end time is the end time of the last occurrence of the incident in the particular time interval.
- If you specify the `-db` option without a database name, then the `chactl query diagnosis` command displays diagnostic information for all databases. However, if a database name is specified, then the `chactl query diagnosis` command displays diagnostic information for all instances of the database that are being monitored.
- If you specify the `-cluster` option without a host name, then the `chactl query diagnosis` command displays diagnostic information for all hosts in that cluster.
- If you do not specify a time interval, then the `chactl query diagnosis` command displays only the current issues for all or the specified targets. The `chactl query diagnosis` command does not display the frequency statistics explicitly. However, you can count the number of normal and abnormal events that occurred in a target in the last 24 hours.
- If no incidents have occurred during the specified time interval, then the `chactl query diagnosis` command returns a text message, for example, Database/host is operating NORMALLY, or no incidents were found.
- If the state of a target is **NORMAL**, the command does not report it. The `chactl query diagnosis` command reports only the targets with **ABNORMAL** state for the specified time interval.

Output parameters:

- Incident start Time
- Incident end time (only for the default database and/or host, non-verbose output)
- Target (for example, database, host)
- Problem

Description: Detailed description of the problem

Cause: Root cause of the problem and contributing factors

- Action: an action that corrects the abnormal state covered in the diagnosis

Reporting Format: The diagnostic information is displayed in a time compressed or time series order, grouped by components.

Examples

To display diagnostic information of a database for a specific time interval:

```
$ chactl query diagnosis -db oltpacdb -start "2016-02-01 02:52:50.0" -end
"2016-02-01 03:19:15.0"
2016-02-01 01:47:10.0 Database oltpacdb DB Control File IO Performance
(oltpacdb_1) [detected]
2016-02-01 01:47:10.0 Database oltpacdb DB Control File IO Performance
(oltpacdb_2) [detected]
2016-02-01 02:52:15.0 Database oltpacdb DB CPU Utilization (oltpacdb_2)
[detected]
2016-02-01 02:52:50.0 Database oltpacdb DB CPU Utilization (oltpacdb_1)
[detected]
2016-02-01 02:59:35.0 Database oltpacdb DB Log File Switch (oltpacdb_1)
[detected]
2016-02-01 02:59:45.0 Database oltpacdb DB Log File Switch (oltpacdb_2)
[detected]
```

Problem: DB Control File IO Performance

Description: CHA has detected that reads or writes to the control files are slower than expected.

Cause: The Cluster Health Advisor (CHA) detected that reads or writes to the control files were slow

because of an increase in disk IO.

The slow control file reads and writes may have an impact on checkpoint and Log Writer (LGWR) performance.

Action: Separate the control files from other database files and move them to faster disks or Solid State Devices.

Problem: DB CPU Utilization

Description: CHA detected larger than expected CPU utilization for this database.

Cause: The Cluster Health Advisor (CHA) detected an increase in database CPU utilization

because of an increase in the database workload.

Action: Identify the CPU intensive queries by using the Automatic Diagnostic and Defect Manager (ADDM)

and follow the recommendations given there. Limit the number of CPU intensive queries

or relocate sessions to less busymachines. Add CPUs if the CPU capacity is insufficient to support the load

without a performance degradation or effects on other databases.

Problem: DB Log File Switch

Description: CHA detected that database sessions are waiting longer than expected for log switch completions.

Cause: The Cluster Health Advisor (CHA) detected high contention during log

switches
because the redo log files were small and the redo logs switched frequently.
Action: Increase the size of the redo logs.

Error Message

Message: *Target* is operating normally

Description: No incidents are found on the target.

Message: No data was found for active *Target*

Description: No data was found, but the target was operating or active at the time of the query.

Message: Target is not active or was not being monitored.

Description: No data was found because the target was not monitored at the time of the query.

C.7 chactl query model

Use the `chactl query model` command to list all Oracle Cluster Health Advisor models or to view detailed information about a specific Oracle Cluster Health Advisor model.

Syntax

```
chactl query model [-name model_name [-verbose]][-help]
```

Examples

- To list all base Oracle Cluster Health Advisor models:

```
$ chactl query model
Models: MOD1, MOD2, MOD3, MOD4, MOD5, MOD6, MOD7
```

```
$ chactl query model -name weekday
Model: weekday
Target Type: DATABASE
Version: 12.2.0.1_0
OS Calibrated on: Linux amd64
Calibration Target Name: prod
Calibration Date: 2016-09-10 12:59:49
Calibration Time Ranges: start=2016-09-09 16:00:00,end=2016-09-09
23:00:00
Calibration KPIs: not specified
```

- To view detailed information, including calibration metadata, about the specific Oracle Cluster Health Advisor model:

```
$ chactl query model -name MOD5 -verbose
Model: MOD5
```

```

CREATION_DATE:          Jan 10,2016 10:10
VALIDATION_STATUS:     Validated
DATA_FROM_TARGET :     inst72, inst75
USED_IN_TARGET :       inst76, inst75, prodDB, evalDB-evalSP
CAL_DATA_FROM_DATE:    Jan 05,2016 10:00
CAL_DATA_TO_DATE:      Jan 07,2016 13:00
CAL_DATA_FROM_TARGETS  inst73, inst75
...

```

C.8 chactl query repository

Use the `chactl query repository` command to view the maximum retention time, number of targets, and the size of the Oracle Cluster Health Advisor repository.

Syntax

```
chactl query repository [-help]
```

Examples

To view information about the Oracle Cluster Health Advisor repository:

```

$ chactl query repository
specified max retention time(hrs) : 72
available retention time(hrs)     : 212
available number of entities      : 2
allocated number of entities      : 0
total repository size(gb)        : 2.00
allocated repository size(gb)     : 0.07

```

C.9 chactl query calibration

Use the `chactl query calibration` command to view detailed information about the calibration data of a specific target.

Syntax

```
chactl query calibration {-cluster|-db db_unique_name} [-timeranges
'start=time_stamp,end=time_stamp,...'] [-kpiset 'name=kpi_name min=val
max=val,...'] [-interval val][-help]
```

Specify the interval in hours.

Specify date and time in the YYYY-MM-DD HH24:MI:SS format.

Note:

If you do not specify a time interval, then the `chactl query calibration` command displays all the calibration data collected for a specific target.

The following Key Performance Indicators are supported for database:

- CPUPERCENT - CPU utilization - Percent
- IOREAD - Disk read - Mbyte/sec
- DBTIMEPERCALL - Database time per user call - usec/call
- IOWRITE - Disk write - Mbyte/sec
- IOTHROUGHPUT - Disk throughput - IO/sec

The following Key Performance Indicators are supported for cluster:

- CPUPERCENT - CPU utilization - Percent
- IOREAD - Disk read - Mbyte/sec
- IOWRITE - Disk write - Mbyte/sec
- IOTHROUGHPUT - Disk throughput - IO/sec

Examples

To view detailed information about the calibration data of the specified target:

```
$ chactl query calibration -db oltpacdb -timeranges
'start=2016-07-26 01:00:00,end=2016-07-26 02:00:00,start=2016-07-26
03:00:00,end=2016-07-26 04:00:00'
-kpiset 'name=CPUPERCENT min=20 max=40, name=IOTHROUGHPUT min=500
max=9000' -interval 2
```

```
Database name : oltpacdb
Start time : 2016-07-26 01:03:10
End time : 2016-07-26 01:57:25
Total Samples : 120
Percentage of filtered data : 8.32%
The number of data samples may not be sufficient for calibration.
```

1) Disk read (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
4.96	0.20	8.98	0.06	25.68
<25	<50	<75	<100	>=100
97.50%	2.50%	0.00%	0.00%	0.00%

2) Disk write (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
27.73	9.72	31.75	4.16	109.39
<50	<100	<150	<200	>=200
73.33%	22.50%	4.17%	0.00%	0.00%

3) Disk throughput (ASM) (IO/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
2407.50	1500.00	1978.55	700.00	7800.00

<5000	<10000	<15000	<20000	>=20000
83.33%	16.67%	0.00%	0.00%	0.00%

4) CPU utilization (total) (%)

MEAN	MEDIAN	STDDEV	MIN	MAX
21.99	21.75	1.36	20.00	26.80

<20	<40	<60	<80	>=80
0.00%	100.00%	0.00%	0.00%	0.00%

5) Database time per user call (usec/call)

MEAN	MEDIAN	STDDEV	MIN	MAX
267.39	264.87	32.05	205.80	484.57

<10000000	<20000000	<30000000	<40000000	<50000000	<60000000	<70000000	>=70000000
100.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

Database name : oltpacdb
 Start time : 2016-07-26 03:00:00
 End time : 2016-07-26 03:53:30
 Total Samples : 342
 Percentage of filtered data : 23.72%
 The number of data samples may not be sufficient for calibration.

1) Disk read (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
12.18	0.28	16.07	0.05	60.98

<25	<50	<75	<100	>=100
64.33%	34.50%	1.17%	0.00%	0.00%

2) Disk write (ASM) (Mbyte/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
57.57	51.14	34.12	16.10	135.29

<50	<100	<150	<200	>=200
49.12%	38.30%	12.57%	0.00%	0.00%

3) Disk throughput (ASM) (IO/sec)

MEAN	MEDIAN	STDDEV	MIN	MAX
5048.83	4300.00	1730.17	2700.00	9000.00

<5000	<10000	<15000	<20000	>=20000
63.74%	36.26%	0.00%	0.00%	0.00%

4) CPU utilization (total) (%)

MEAN	MEDIAN	STDDEV	MIN	MAX
------	--------	--------	-----	-----

```

23.10      22.80      1.88      20.00      31.40
<20       <40        <60       <80       >=80
0.00%     100.00%   0.00%     0.00%     0.00%

```

5) Database time per user call (usec/call)

```

MEAN      MEDIAN     STDDEV     MIN        MAX
744.39    256.47    2892.71    211.45     45438.35

```

```

<10000000 <20000000 <30000000 <40000000 <50000000 <60000000
<70000000 >=70000000
100.00%    0.00%     0.00%     0.00%     0.00%     0.00%     0.00%
0.00%

```

C.10 chactl remove model

Use the `chactl remove model` command to delete an Oracle Cluster Health Advisor model along with the calibration data and metadata of the model from the Oracle Cluster Health Advisor repository.

Note:

If the model is being used to monitor the targets, then the `chactl remove model` command cannot delete any model.

Syntax

```
chactl remove model -name model_name [-help]
```

Error Message

Error: *model_name* does not exist

Description: The specified Oracle Cluster Health Advisor model does not exist in the Oracle Cluster Health Advisor repository.

C.11 chactl rename model

Use the `chactl rename model` command to rename an Oracle Cluster Health Advisor model in the Oracle Cluster Health Advisor repository.

Assign a descriptive and unique name to the model. Oracle Cluster Health Advisor preserves all the links related to the renamed model.

Syntax

```
chactl rename model -from model_name -to model_name [-help]
```

Error Messages

Error: `model_name` does not exist

Description: The specified model name does not exist in the Oracle Cluster Health Advisor repository.

Error: `dest_name` already exist

Description: The specified model name already exists in the Oracle Cluster Health Advisor repository.

C.12 chactl export model

Use the `chactl export model` command to export Oracle Cluster Health Advisor models.

Syntax

```
chactl export model -name model_name -file output_file [-help]
```

Example

```
$ chactl export model -name weekday -file /tmp//weekday.mod
```

C.13 chactl import model

Use the `chactl import model` command to import Oracle Cluster Health Advisor models.

Syntax

```
chactl import model -name model_name -file model_file [-force] [-help]
```

While importing, if there is an existing model with the same name as the model being imported, then use the `-force` option to overwrite.

Example C-1 Example

```
$ chactl import model -name weekday -file /tmp//weekday.mod
```

C.14 chactl set maxretention

Use the `chactl set maxretention` command to set the maximum retention time for the diagnostic data.

The default and minimum retention time is 72 hours. If the Oracle Cluster Health Advisor repository does not have enough space, then the retention time is decreased for all the targets.

 **Note:**

Oracle Cluster Health Advisor stops monitoring if the retention time is less than 24 hours.

Syntax

```
chactl set maxretention -time retention_time [-help]
```

Specify the retention time in hours.

Examples

To set the maximum retention time to 80 hours:

```
$ chactl set maxretention -time 80  
max retention successfully set to 80 hours
```

Error Message

Error: Specified time is smaller than the allowed minimum

Description: This message is returned if the input value for maximum retention time is smaller than the minimum value.

C.15 chactl resize repository

Use the `chactl resize repository` command to resize the tablespace of the Oracle Cluster Health Advisor repository based on the current retention time and the number of targets.

 **Note:**

The `chactl resize repository` command fails if your system does not have enough free disk space or if the tablespace contains data beyond requested resize value.

Syntax

```
chactl resize repository -entities total number of hosts and database instances [-force | -eval] [-help]
```

Examples

To set the number of targets in the tablespace to 32:

```
chactl resize repository -entities 32  
repository successfully resized for 32 targets
```

D

Behavior Changes, Deprecated and Desupported Features

Review information about changes, deprecations, and desupports.

- [Oracle Database Quality of Service \(QoS\) Management is Deprecated in Release 21c](#)
Starting in Oracle Database release 21c, Oracle Database Quality of Service (QoS) Management is deprecated and will be desupported in a future release.

D.1 Oracle Database Quality of Service (QoS) Management is Deprecated in Release 21c

Starting in Oracle Database release 21c, Oracle Database Quality of Service (QoS) Management is deprecated and will be desupported in a future release.

Oracle Database Quality of Service (QoS) Management automates the workload management for an entire system by adjusting the system configuration based on pre-defined policies to keep applications running at the performance levels needed. Applications and databases are increasingly deployed in systems that provide some of the resource management capabilities of Oracle Database Quality of Service (QoS) Management. At the same time, Oracle's Autonomous Health Framework has been enhanced to adjust and provide recommendations to mitigate events and conditions that impact the health and operational capability of a system and its associated components. For those reasons, Oracle Database Quality of Service (QoS) Management has been deprecated with Oracle Database 21c.