

Oracle® Enterprise Manager

Advanced Management and Monitoring for Engineered Systems



13c Release 5

F93993-01

July 2024

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2024, Oracle and/or its affiliates.

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, MySQL, and NetSuite 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	v
Documentation Accessibility	vi
Related Resources	vi
Conventions	vi

1 Introduction to the Engineered Systems Packs

Features of the Engineered Systems Packs	1-2
--	-----

2 Configure Access to the Engineered Systems Packs

Configure Access to the Exadata Management Pack	2-1
Configure Access to the Zero Data Loss Recovery Appliance Management Pack	2-2

3 Using the Engineered Systems Packs

Features Common in Both Packs	3-1
Exadata Fleet Dashboards	3-1
Exadata Fleet Configuration	3-5
Exadata Fleet Aggregated Capacity	3-7
Exadata Fleet Capacity	3-8
Exadata Fleet Component Capacity	3-14
Engineered System Autonomous Health Framework Administration	3-19
Monitoring AHF and Exachk for the Engineered System Fleet	3-22
Managing Engineered System AHF Installations	3-26
Exadata Management Pack Features	3-31
Database Impact Advisor	3-31
AHF Configuration for Database Impact Advisor	3-36
I/O Resource Management Advisor	3-38
Exadata Cloud Advanced Monitoring	3-45
Visualizing Database Topology on Exadata VM Clusters	3-46
Visualizing Resource Usage for Exadata Infrastructure	3-49
Exadata Exascale Advanced Monitoring	3-51

Exadata Infrastructure Patching Using Fleet Maintenance	3-53
Set up Infrastructure for Exadata Patching	3-54
Image Maintenance	3-73
Subscribe Exadata Targets	3-79
Exadata Fleet Maintenance Patching Operations	3-79
Rollback Exadata Patches	3-83
Zero Data Loss Recovery Appliance Management Pack Features	3-85
Database Fleet Backup Configuration for Recovery Appliance	3-86
Database Fleet Backup Scheduling for Recovery Appliance	3-94
Touchless Automated End-to-End Database Fleet Backup Management	3-97
Recovery Appliance Protected Database Fleet Archival Backup Management	3-101
Database Fleet Advanced Backup Configuration for Recovery Appliance Using EMCLI	3-104

4 Troubleshooting

Issues with Database Impact Advisor	4-1
-------------------------------------	-----

Preface

Oracle Enterprise Manager Cloud Control now offers advanced management and monitoring of Engineered Systems through the following packs:

- *Exadata Management Pack* provides advanced monitoring, management, and diagnostics features for hybrid Exadata fleets spanning on-premises, Cloud@Customer, Oracle Cloud Infrastructure, and Multicloud, as well as the databases running on these systems.
- *Zero Data Loss Recovery Appliance Management Pack* provides advanced features for managing database fleet backup to one or more Recovery Appliances, and for monitoring and managing a fleet of Recovery Appliances.

Note that some of the features are common to both packs.

This book describes how you can set up and use these packs in Oracle Enterprise Manager Cloud Control to access rich UI and EMCLI management and monitoring features for the respective Engineered Systems. With these packs, you can access the following functionality:

- Rich, multi-tiered dashboards to analyze Exadata and Recovery Appliance capacity, inventory and performance.
- Comprehensive UI console and tooling to monitor and administer Autonomous Health Framework across an Exadata and Recovery Appliance fleet.
- UI tooling and reports to analyze database CPU usage within an Exadata system.
- UI components to visualize the database topology and inventory for an Exadata Cloud system.
- UI components to visualize the Exascale storage configuration for an Oracle Exadata Database Machine.
- UI and EMCLI workflows to streamline configuration and management of database fleet backup to multiple Recovery Appliances, including advanced backup topologies.
- UI workflow to manage database fleet archival backups from Recovery Appliance to the tape.

In the context of this book, **Engineered Systems** can include one or more of *Oracle Exadata Database Machine*, *Zero Data Loss Recovery Appliance*, and *Exadata Cloud* which includes *Oracle Exadata Database Service on Dedicated Infrastructure* and *Oracle Exadata Database Service on Cloud@Customer*.

Audience

This guide is intended for Exadata and Recovery Appliance administrators, and also for database administrators responsible for managing database backups to Recovery Appliance destinations.

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

Related Resources

The information about discovery of the Engineered Systems and their monitoring are available here:

- Introduction to Monitoring Exadata Cloud Service
- Introduction to the Plug-in in *Oracle Exadata Database Machine Getting Started Guide*
- [Get started with Cloud Control for Recovery Appliance](#)
- Oracle Engineered Systems Metric Reference Manual
- [Other Engineered Systems Resources](#)

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 the Engineered Systems Packs

The Exadata and Zero Data Loss Recovery Appliance packs contain comprehensive, value-added fleet-level management and monitoring functionality intended for advanced use cases.

Advanced Management and Monitoring Solution

The advanced management and monitoring functionality offered by these packs for hybrid Engineered System fleets focuses on the following value additions to the existing core management and monitoring functionality contained in the Exadata, ZDLRA, and Database plug-ins:

- Automation of fleet-level, multi-step, complex procedures that cannot be easily performed outside of Oracle Enterprise Manager Cloud Control
- Automation that leverages the inherent Oracle Enterprise Manager Cloud Control capabilities like procedure scheduling and execution, credentials management, and the Software Library to streamline Engineered System configuration and maintenance
- Features at both the individual system and fleet levels that integrate extensive Oracle Enterprise Manager Cloud Control metric data with new UIs and dashboards to provide unique, deep analysis and visualization tools

Core Management and Monitoring Functionality

The following core management and monitoring functionality is standard in the Exadata and ZDLRA plug-ins:

- Support for new Exadata, ZDLRA, and database software and hardware versions
- Monitoring, including metrics for core functions
- Per-system management workflows for core Exadata and ZDLRA administration operations

For more information about the plug-ins, see [Related Resources](#).

Advantages of the Exadata and Zero Data Loss Recovery Appliance Management Packs

- **Broad feature support across the major Engineered Systems on-premises and Cloud-based variants:** Oracle Exadata Database Machine, Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, Zero Data Loss Recovery Appliance
- **User Interfaces, EMCLI verbs, workflows, jobs, and deployment procedures** that address advanced Engineered System use cases beyond the functionality currently available in the Exadata, ZDLRA, and Database plug-ins
- **Fleet-scale deployment procedures** that automate complex, multi-step processes across multiple Exadata systems, Recovery Appliances, databases, and hosts
- **Dashboards and consoles** that operate at the Engineered System fleet level to provide large scale capacity, inventory and performance analysis, compliance management, and patching management
- **Advisors** that provide deep, granular database performance analysis capabilities for an individual Exadata

- **Workflows and procedures** that operate at the database fleet level to provide streamlined end-to-end backup configuration and scheduling for Recovery Appliance destinations
- **New Exadata and Recovery Appliance target consoles** from which all the new advanced Engineered System fleet management features can be accessed

Benefits of the Exadata Management Pack

Exadata administrators who need to perform the following activities across a large fleet of on-premises and/or Cloud-based Exadata systems:

- Analyze performance, capacity, and inventory across the Exadata fleet.
- Manage and monitor Exachk compliance check policies, schedules, and invocations across the fleet.
- Perform Exadata hardware component patching across the Exadata fleet.
- Perform noisy-neighbor and CPU utilization analysis across all databases running on an Exadata.
- Visualize Exascale storage usage across all databases running on an Exadata.
- Visualize a detailed inventory of databases running on a Cloud-based Exadata (Oracle Exadata Database Service on Cloud@Customer and Oracle Exadata Database Service on Dedicated Infrastructure).

Benefits of the Zero Data Loss Recovery Appliance Management Pack

Recovery Appliance and database administrators who need to perform the following activities across a large database fleet combined with one or more Recovery Appliances:

- Analyze performance, capacity, and inventory across the Recovery Appliance fleet.
- Monitor Exachk compliance check policies, schedules, and invocations across the Recovery Appliance fleet.
- Configure multiple databases to send backups and redo to one or more Recovery Appliances, including optional configuration of alternate Recovery Appliance destinations.
- Configure databases in one or more Data Guard configurations to send backups and redo from the primary and/or standby databases to one or more Recovery Appliances, including optional configuration of alternate Recovery Appliance destinations.
- Schedule automatic, optimized archival backups for multiple databases from a Recovery Appliance to Oracle Cloud storage or tape.

Features of the Engineered Systems Packs

The following features are included in the Exadata Management Pack and Zero Data Loss Recovery Appliance Pack.

Topics:

- [Features Common to Both Packs](#)
- [Features Specific to Exadata Management Pack](#)
- [Features Specific to Zero Data Loss Recovery Appliance Management Pack \(ZDLRA\)](#)

Features, Packs Supporting Them, and the Plug-ins Required:

Feature	Supporting Pack	Required Plug-in
Engineered System Fleet Console	Not licensed	Exadata
Exadata Fleet Dashboards	<ul style="list-style-type: none"> Exadata Management Pack Zero Data Loss Recovery Appliance Management Pack 	Exadata
Engineered System Autonomous Health Framework Administration	<ul style="list-style-type: none"> Exadata Management Pack Zero Data Loss Recovery Appliance Management Pack 	Exadata
Database Impact Advisor	Exadata Management Pack	Exadata
I/O Resource Management Advisor	Exadata Management Pack	Exadata
Exadata Cloud Advanced Monitoring	Exadata Management Pack	Exadata
Exadata Exascale Advanced Monitoring	Exadata Management Pack	Exadata
Exadata Infrastructure Patching Using Fleet Maintenance	Exadata Management Pack	Exadata
Database Fleet Backup Configuration and Scheduling for Recovery Appliance	Zero Data Loss Recovery Appliance Management Pack	<ul style="list-style-type: none"> Database ZDLRA
Recovery Appliance Protected Database Fleet Archival Backup Management	Zero Data Loss Recovery Appliance Management Pack	ZDLRA
Database Fleet Advanced Backup Configuration for Recovery Appliance Using EMCLI	Zero Data Loss Recovery Appliance Management Pack	<ul style="list-style-type: none"> Database ZDLRA

Features Common to Both Packs

- Engineered System Fleet Console**
 Although it's a standard Exadata and ZDLRA plug-in feature and not part of a pack, the Engineered System Fleet Console is the gateway to many of the features in Exadata Management Pack and Zero Data Loss Recovery Appliance Management Pack. This page is the new destination for the **Exadata** and **Recovery Appliance** menu items under **Targets**, with a separate tab for each target type. It encompasses all functionality of the previous fleet pages for these target types, plus the following new features:
 - Top-level charts summarizing fleet target status, member status, and pack enablement status
 - Table columns indicating whether a specific Exadata or Recovery Appliance target has the respective pack enabled
 - Launch menu where most fleet-level pack features can be accessed
- Exadata Fleet Dashboards**
 The Exadata Fleet Dashboards distill Exadata component hardware and software metric data for individual Exadata (on-premises and Cloud-based) and Recovery Appliance systems into integrated, detailed, fleet-level visualizations that can be highly customized through an extensive set of filters across different dimensions. The available dashboards are as follows:
 - **Exadata Fleet Configuration:** This dashboard shows the current Exadata hardware and software configuration across the Exadata and Recovery Appliance fleet or filtered subset of the fleet, and includes the following data:

- * *Engineered System Overview*: Type, hardware generation, deployment type, service type, database deployment option, and switch type
- * *Hardware Configuration*: Exadata generation, database server model, storage server model, storage server media type, and switch model
- * *Software Configuration*: Host operating system version, database server ILOM version, storage server Exadata system software version, database server Exadata system software version, and switch version
- * *Database Configuration*: Container type, database version, pluggable database version, ASM disk group redundancy, and Grid Infrastructure version
- **Exadata Fleet Aggregated Capacity**: This dashboard summarizes Exadata capacity for the Exadata and Recovery Appliance fleet or filtered subset of the fleet. Data is averaged over the last 31 days. The dashboard includes the following data:
 - * *Database Server Capacity and Performance*: CPU allocation and utilization, memory allocation and utilization
 - * *Storage Capacity and Performance*: Storage allocation and utilization
 - * *Storage Server I/O*: Read latency/throughput/IOPS, write latency/throughput/IOPS, IO utilization/load, total read throughput, and total write throughput
- **Exadata Fleet Capacity**: This dashboard shows individual Exadata capacity across the Exadata and Recovery Appliance fleet or filtered subset of the fleet, allowing comparison by system. Data is averaged over the last 31 days. The dashboard includes the following data:
 - * *ASM Usage Growth*: Data disk group usage growth, recovery disk group usage growth, and database storage usage growth
 - * *Storage Server I/O*: Disk and flash read and write latency, utilization, IOPS, and throughput
 - * *Database Server Performance Details*: CPU utilization and memory utilization
 - * *ASM Performance Details*: Disk group space utilization and database tablespace utilization
 - * *Storage Server I/O Details*: Capacity, disk read latency/utilization/IOPS/throughput and flash latency/utilization/IOPS/throughput
- **Exadata Fleet Component Capacity**: This dashboard shows individual Exadata component capacity across the Exadata and Recovery Appliance fleet or filtered subset of the fleet, allowing identification and comparison of components that may be over or under-utilized. Data is averaged over the last 31 days. The information categories are similar to the Exadata Fleet Capacity dashboard, but the data is shown on a per-component basis instead of a per-system basis.

For details about the feature, see [Exadata Fleet Dashboards](#).

- **Engineered System Autonomous Health Framework Administration**

The Engineered System Autonomous Health Framework (AHF) Administration console provides comprehensive monitoring and management of AHF and Exachk across the Engineered System fleet, including the following features:

- Monitoring of AHF software versions, upgrade status, service status, Exachk status, and Exachk runs across the Exadata and Recovery Appliance fleet
- Management of AHF software automatic and manual upgrades across the Exadata fleet
- Management of Exachk automatic and manual invocations across the Exadata fleet

- Invocation of AHF commands across the Exadata fleet

For details about the feature, see [Engineered System Autonomous Health Framework Administration](#).

Features Specific to Exadata Management Pack

- **Database Impact Advisor**

The Database Impact Advisor can be run against an individual Exadata system to perform system-wide database CPU usage *noisy-neighbor* analysis in order to identify databases whose performance is potentially impacted by other databases or other operating system processes. The analysis applies algorithms established by the Autonomous Health Framework (AHF) Balance feature to Enterprise Manager historical metric data from past 30 days. Further, the Database Impact Advisor integrates directly with AHF Balance to generate recommendations for optimizing Database Resource Manager (DBRM) settings across all databases to minimize any CPU-based performance impacts that were found.

The Database Impact Advisor offers the following features:

- Summary of potentially impacted databases
- Summary of impacts across clusters, hosts, databases and instances
- Detailed CPU impact breakdown by Exadata, cluster, host, database, and instance:
 - * Impact status (PASS/WARNING/FAIL) at each level
 - * Number of impacted and exposed hours for each database and instance
- In-context per-instance historical charts showing:
 - * CPU usage in comparison with the DBRM guarantee and limit
 - * CPU usage in comparison with other instances and process on the host
 - * Impacted and exposed time ranges
 - * Occurrences of notable database and Exadata incidents
- Overall CPU allocation breakdown across database instances, grouped by cluster and host, and filtered by impact status
- Generation and display of AHF Balance reports at the Exadata, cluster, and database levels
- Integrated management of a specific AHF installation for use with Enterprise Manager analytics

For details about the feature, see [Database Impact Advisor](#).

- **I/O Resource Management Advisor**

The I/O Resource Management Advisor facilitates effective usage of Exadata I/O Resource Management (IORM) by providing tools to analyze flash I/O utilization across all databases running on an Exadata system and to identify potential measures to optimize system I/O resource usage. It provides the following capabilities:

- **Comprehensive System-Level I/O Overview:** Different visualizations at multiple levels like treemaps, tables and metric charts provide a comprehensive overview of IORM settings and a summary of I/O utilization across the Exadata system. These visualizations allow identification of any container database (CDB), pluggable database (PDB), or non-container database that is experiencing flash I/O throttling due to system limits or I/O usage by neighboring databases.
- **Detailed Database-Level I/O Analysis:** Once a potentially impacted database is identified, additional drill-down visualizations allow analysis of the I/O performance history of a database and its neighbors across multiple dimensions, including I/O

utilization, throughput, and latency. Analysis at this level can assist in determining whether a database is actually experiencing throttling and associated performance degradation, and whether its IORM plan needs to be updated.

For details about the feature, see [I/O Resource Management Advisor](#).

- **Exadata Cloud Advanced Monitoring**

The **Exadata Infrastructure** target (used for *Oracle Exadata Database Service on Cloud@Customer* and *Oracle Exadata Database Service on Dedicated Infrastructure*) home page provides advanced database inventory monitoring through the **Database Inventory** tab. This tab provides a Sankey-based topology diagram and a tabular listing of database inventory on the system, tracing the topology starting from the database major version down through the minor version (patch release), Oracle home, number of container/non-container databases, number of pluggable databases, ending with the specific pluggable (or non-container) databases.

For details about the feature, see [Exadata Cloud Advanced Monitoring](#).

- **Exadata Exascale Advanced Monitoring**

The Exascale target home page provides advanced IOPS monitoring through the IOPS tab using an interactive Sankey-based illustration of IOPS allocations across the Vaults, Storage Pools, and Databases. You can use advanced filtering options to visually identify IOPS bottlenecks in the Exascale topology and view these in context with historical Exascale IOPS utilization in order to triage performance issues.

For details about the feature, see [Exadata Exascale Advanced Monitoring](#).

- **Exadata Infrastructure Patching Using Fleet Maintenance**

Exadata Fleet Maintenance is a solution designed to simplify and streamline the patching process for your entire Exadata infrastructure, ensuring optimal performance, security, and uptime.

For details about the feature, see [Exadata Infrastructure Patching Using Fleet Maintenance](#).

Features Specific to Zero Data Loss Recovery Appliance Management Pack (ZDLRA)

The Zero Data Loss Recovery Appliance Management Pack offers an integrated suite of tools like UIs, EMCLI verbs, REST APIs, and associated deployment procedures to manage the end-to-end backup lifecycle for a database fleet. For database administrators managing fleet backups, the fundamental initial task is to configure the database fleet to send backups and redo to one or more Recovery Appliances. Next, the Oracle-suggested incremental-forever backup strategy must be scheduled, followed by ongoing fleet configuration updates and maintenance. All aspects of these tasks can be managed by the following tools:

- **Database Fleet Backup Configuration and Scheduling for Recovery Appliance**

- **Configure Fleet Backup to Recovery Appliance**

This wizard-style UI provides the ability to configure a hybrid fleet of databases (versions, types, roles) to send backups and redo to a Recovery Appliance, that is, configure them for Recovery Appliance protection. It submits a deployment procedure that fully automates all aspects of configuring the databases, including database, Grid Infrastructure, SQLNET.ORA file, wallet, and recovery catalog settings. The workflow and associated deployment procedure offer the following features:

- * Selection of multiple individual databases or a group target containing the databases to be configured
- * Support for all database versions, platforms, and types (single-instance and RAC)
- * Support for configuring backup from Data Guard primary and standby databases

- * Support for Recovery Appliance Backup and Recovery Continuity (designation of an alternate Recovery Appliance for backups and redo in the event of a failure)
- * Optional auto-configure mode for groups, whereby databases that subsequently join the group will be automatically configured without user intervention
- * UI functionality also accessible via an Enterprise Manager REST API and an EMCLI verb

For details about the feature, see [Database Fleet Backup Configuration for Recovery Appliance](#).

– **Schedule Fleet Backup to Recovery Appliance**

This wizard-style UI is associated with the Configure wizard above. It provides the ability to schedule the standardized Oracle-suggested Recovery Appliance backup strategy for a fleet of databases that have already been configured for Recovery Appliance protection. It schedules a recurring deployment procedure that invokes individual database backups in parallel. The workflow and associated deployment procedure offer the following features:

- * Selection of multiple individual databases or a group target containing the databases to be backed-up (can use the same group targets as Configure wizard)
- * Implementation of Oracle-suggested incremental-forever, daily backup strategy
- * New databases that join the group are automatically picked up in subsequent scheduled backup executions
- * All backups start on same schedule, and run in parallel
- * Individual databases can backup to different Recovery Appliances (depending on which Recovery Appliance they were configured to send backups to)
- * UI functionality also accessible with the Enterprise Manager REST API and EMCLI verb

For details about the feature, see [Database Fleet Backup Scheduling for Recovery Appliance](#).

For a discussion on end-to-end use case for implementing configuration and scheduling, see [Touchless Automated End-to-End Database Fleet Backup Management](#).

• **Recovery Appliance Protected Database Fleet Archival Backup Management**

A Recovery Appliance administrator may need to perform and manage archival backups of certain databases in the fleet to meet regulatory requirements. This can be accomplished by using the **Create Archival Backup** workflow, which is accessible from the Recovery Appliance target **Archival Backups** page. The Create Archival Backup workflow submits a deployment procedure that, for each database, polls the Recovery Appliance for the availability of an indexed backup correlating to the desired archival backup point in time, and then invokes the archival backup. The workflow and associated deployment procedure offer the following features:

- Selection of an individual protected database or all databases within a Recovery Appliance protection policy
- The ability to perform a one-time archival backup or schedule repeating archival backups
- Creation of the archival backup based on a point in time for individual and multiple databases, or a system change number (SCN) or restore point for individual databases
- Specification of the backup retention time
- Specification of compression and encryption options

- Special verification measures to ensure creation of optimized archival backups that contain the minimum possible number of archived log backups

For details about the feature, see [Recovery Appliance Protected Database Fleet Archival Backup Management](#).

- **Database Fleet Advanced Backup Configuration for Recovery Appliance Using EMCLI**

Support for configuration of advanced backup topologies beyond the scenarios covered by the above UI tooling is available via EMCLI. In addition to all the UI capabilities, the EMCLI `configure_db_ha` command also allows configuration of the databases in a Data Guard configuration to send backup and redo to multiple Recovery Appliances, in accordance with several possible Maximum Availability Architecture (MAA) recommended database protection configurations. Further, even when submitting a fleet operation, use of EMCLI allows highly granular per-database customization of backup configuration and scheduling options. The EMCLI `configure_db_ha` command submits the same deployment procedure as the above UI workflow, ensuring consistent results regardless of what interface is used. The EMCLI `configure_db_ha` command offers the following advanced features:

- **Configuration of the following Oracle MAA backup configurations** (involves Data Guard databases and two Recovery Appliances):
 - * **MAA Gold:** Backups scheduled from both primary and standby databases to different Recovery Appliances that are not configured for backup replication
 - * **MAA Gold + ZDLRA Backup Anywhere Replication:** Backups scheduled from either the primary or standby database to their respective local Recovery Appliance, with Backup Anywhere replication configured between the Recovery Appliances
- **Configuration of ZDLRA Backup and Recovery Continuity with the above MAA Gold configurations** (involves Data Guard databases and four Recovery Appliances):
 - * Specification of a preferred Recovery Appliance and alternate Recovery Appliance for both the primary and standby database (preferred and alternate Recovery Appliances must be configured for *Backup Anywhere* replication)
 - * Backups go to the preferred Recovery Appliance by default and automatically failover to the alternate Recovery Appliance if needed.
- **Configuration of ZDLRA Backup and Recovery Continuity for a Non-Data Guard Database:**
 - * Specification of a preferred Recovery Appliance and alternate Recovery Appliance for the database
 - * Backups go to the Preferred Recovery Appliance and automatically failover to the alternate Recovery Appliance ensuring Backup and Recovery Continuity.

For details about the feature, see [Database Fleet Advanced Backup Configuration for Recovery Appliance Using EMCLI](#).

2

Configure Access to the Engineered Systems Packs

Topics:

- [Configure Access to the Exadata Management Pack](#)
- [Configure Access to the Zero Data Loss Recovery Appliance Management Pack](#)

Also, see Displaying Management Pack Information in Enterprise Manager in *Licensing Information User Manual*.

To modify the default settings for the new targets discovered in Enterprise Manager Cloud Control, see the section **Default Settings** in Enabling and Disabling Management Packs in Enterprise Manager in *Licensing Information User Manual*.

Configure Access to the Exadata Management Pack

Ensure to complete the prerequisite tasks before you access the features available in the pack. Verify that the pack is enabled for your targets.

The Exadata Management Pack is enabled by default. Hence, you will be able to leverage the features of the pack automatically.

However, to disable or to re-enable the access of the pack at a later point if you had disabled it, follow these steps:

1. Log in to your Enterprise Manager Cloud Control account.
2. Click **Settings**, click **Management Packs**, and click **Management Pack Access**. The Management Packs page opens. Here, all the targets and the applicable packs are listed. The **Target based** pack access is enabled by default.

3. From the **Search** menu, select from the targets for which Exadata Management Pack is applicable. Select **Exadata Infrastructure** or **Oracle Exadata Database Machine**.

The specified type of Exadata targets are displayed in the table.

4. Enable or disable the **Exadata Management Pack** check box corresponding to your targets. Click **Apply**.

Alternatively, if you want to enable the pack for all your Exadata targets, follow these steps:

1. Log in to your Enterprise Manager Cloud Control account.
2. Click **Settings**, click **Management Packs**, and click **Management Pack Access**. The Management Packs page opens. Here, all the targets and the applicable packs are listed. The **Target Based** pack access is enabled by default.
3. Change the Pack access to **Patch Based Batch Update**. Now you can select from the available packs to apply on your targets.
4. Under **Type**, select the type of target for which you want to select the pack. Select **Exadata Infrastructure** or **Oracle Exadata Database Machine**. Now the packs applicable for your type of target are listed.

5. Select **Exadata Management Pack** and click **Move**. Now the pack is listed in the **Selected Packs** section.
6. Click the **Enable** button for Access. Click **Apply**.

Configure Access to the Zero Data Loss Recovery Appliance Management Pack

Ensure to complete the prerequisite tasks before you access the features available in the pack. Verify that the pack is enabled for your targets.

The Zero Data Loss Recovery Appliance Management Pack is enabled by default. Hence, you will be able to leverage the features of the pack automatically.

However, to disable or to re-enable the access of the pack at a later point if you had disabled it, follow these steps:

1. Log in to your Enterprise Manager Cloud Control account.
2. Click **Settings**, click **Management Packs**, and click **Management Pack Access**. The Management Packs page opens. Here, all the targets and the applicable packs are listed. The **Target based** pack access is enabled by default.
3. From the **Search** menu, select from the targets for which Zero Data Loss Recovery Appliance Management Pack is applicable. Select **Recovery Appliance**.

The specified type of Zero Data Loss Recovery Appliance targets are displayed in the table.

4. Enable or disable the **Zero Data Loss Recovery Appliance Management Pack** check box corresponding to your targets. Click **Apply**.

Alternatively, if you want to enable the pack for all your Zero Data Loss Recovery Appliance targets, follow these steps:

1. Log in to your Enterprise Manager Cloud Control account.
2. Click **Settings**, click **Management Packs**, and click **Management Pack Access**. The Management Packs page opens. Here, all the targets and the applicable packs are listed. The **Target Based** pack access is enabled by default.
3. Change the Pack access to **Patch Based Batch Update**. Now you can select from the available packs to apply on your targets.
4. Under **Type**, select the type of target for which you want to select the pack. Select **Recovery Appliance**. Now the packs applicable for your type of target are listed.
5. Select **Zero Data Loss Recovery Appliance Management Pack** and click **Move**. Now the pack is listed in the **Selected Packs** section.
6. Click the **Enable** button for Access. Click **Apply**.

3

Using the Engineered Systems Packs

Topics:

- [Features Common in Both Packs](#)
- [Exadata Management Pack Features](#)
- [Zero Data Loss Recovery Appliance Management Pack Features](#)

Features Common in Both Packs

Topics:

- [Exadata Fleet Dashboards](#)
- [Engineered System Autonomous Health Framework Administration](#)

Exadata Fleet Dashboards

The Exadata Dashboards are a set of Oracle-defined dashboards that distill Exadata component hardware and software metric data for individual Exadata (on-premises and Cloud-based) and Recovery Appliance systems into integrated, detailed fleet-level visualizations.

These dashboards are part of the Exadata Management Pack and Zero Data Loss Recovery Appliance Management Pack, and will include data only for Database Machine, Exascale Infrastructure, and Recovery Appliance targets for which the corresponding pack is enabled. Some of the dashboards incorporate database target information for the databases running on the Engineered Systems. Since such data will only be shown for databases monitored by Oracle Enterprise Manager, it is recommended that all Oracle Enterprise Manager databases be discovered in Oracle Enterprise Manager in order to leverage the full capability of the dashboards.

Topics:

- [Example Scenario to Analyze Capacity and Resource Usage With the Dashboards](#)
- [Filters for Dashboards](#)
- [Exadata Fleet Configuration](#)
 - [Overview](#)
 - [Hardware](#)
 - [Software](#)
 - [Database](#)
- [Exadata Fleet Aggregated Capacity](#)
 - [Database Server CPU and Memory, and Storage](#)
 - [Storage Server I/O](#)

- [Exadata Fleet Capacity](#)
 - [Database Server CPU and Memory](#)
 - [Storage](#)
 - [Storage Server I/O](#)
 - [Database Server CPU and Memory Details](#)
 - [Storage Details](#)
 - [Storage Server I/O Details](#)
- [Exadata Fleet Component Capacity](#)
 - [Database Server CPU and Memory](#)
 - [Storage](#)
 - [Storage Server I/O](#)
 - [Database Server CPU and Memory Details](#)
 - [Storage Details](#)
 - [Storage Server I/O Details](#)

Example Scenario to Analyze Capacity and Resource Usage With the Dashboards

- **Get a high-level overview of resource usage across the fleet:** Start with the **Exadata Fleet Aggregated Capacity** dashboard, using the filters to narrow down the results to a subset of the fleet, say all *Engineered System Type = Exadata* with *Lifecycle Status = Production* and *Location = Chicago*. If a trend of high storage utilization is noted, then the next step would be to consult the **Exadata Fleet Capacity** dashboard to view the bubble charts representing individual Engineered Systems and compare their metrics.
- **Identify which Exadata systems need attention:** After navigating to the **Exadata Fleet Capacity** dashboard, the filter settings set on the previous dashboard are automatically carried over and the dashboard displays each Exadata system as a bubble on the Storage charts. This makes it easy to identify which Exadata systems are contributing towards the high storage utilization. If filtering with *Group By = Department* shows that the Exadata targets in the sales department have higher storage utilization than the others, then filtering by *Target Property: Department = Sales* would focus attention only on the systems with higher storage utilization.
- **Analyze usage for components within an Exadata system:** The next step is to consult the **Exadata Fleet Component Capacity** dashboard. With the same filters automatically applied, the bubble charts would show the usage patterns of ASM Data and Recovery disk groups for the production Exadata systems in Chicago serving the sales department. The growth for these disk groups could be analyzed further in the adjacent bubble charts and tables.
- **Analyze metric trends:** If the analysis shows that the high storage utilization is isolated to one Database Machine (*DBM1*), this can be further isolated by going back to the **Exadata Fleet Aggregated Capacity** dashboard to view the storage utilization metric history for this Exadata by specifying filter *Engineered System Target Name = DBM1*.

Filters for Dashboards

A common filter section is available across all Exadata Fleet Dashboards to customize the data displayed, allowing for a more focused and manageable data set in each dashboard. These filters are *sticky*, in that they retain their values when navigating across different Exadata Fleet Dashboards, as applicable.

Engineered System Target Name = Select or enter a va...	Engineered System Type All	Engineered System Deployment Type All	Exadata Generation All	Host Operating System Major Version All	Exadata Software Major Version All
Target Property: Contact No values to select.	Target Property: Cost Center No values to select.	Target Property: CSI No values to select.	Target Property: Comment No values to select.	Target Property: Department No values to select.	Target Property: Downtime Contact No values to select.
Target Property: Lifecycle Status No values to select.	Target Property: Line of Business No values to select.	Target Property: Location No values to select.	Target Property: Site No values to select.	Group By LifeCycle Status	

The *Engineered System Target Name*, *Engineered System Type*, and *Engineered System Deployment Type* filters are applied at the Engineered System target level. When applied, they refine the dashboard content, displaying results from the matched Engineered Systems as well as their hierarchy. The *Group By* filter is applicable only to the Exadata Fleet Capacity and Exadata Component Capacity dashboards.

The remaining filters are applied at the individual component level. When the filter criteria are met for at least one component of an Engineered System, the entire Engineered System and all its components are considered a match and rendered along with all other matched systems on the dashboard.

The following supported filters are common across all Exadata Fleet Dashboards:

Filter	Supported Operators	Supported Values
<i>Engineered System Target Name</i>	<ul style="list-style-type: none"> Equal to Not equal to Contains Does not contain Starts with Does not start with Ends with Does not end with 	Any valid Oracle Enterprise Manager target name pattern.
<i>Engineered System Type</i>	Equal to	<ul style="list-style-type: none"> Exadata Recovery Appliance
<i>Engineered System Deployment Type</i>	Equal to	<ul style="list-style-type: none"> Cloud@Customer On-Premises Oracle Cloud Infrastructure
<i>Exadata Generation</i>	Equal to	<p>The Exadata hardware rack generation; for example, X8-2, X7-2.</p> <p>The Exadata generation is derived from the storage server make/model or the database server rack identifier; for example:</p> <ul style="list-style-type: none"> If the storage server make/model value is <i>Oracle Corporation ORACLE SERVER X8-2L High Capacity</i>, then the generation is inferred to be X8-2. If the database server rack identifier is <i>Exadata Database Machine X8-2 AK00831126</i>, then the generation is inferred to be X8-2.


Filter	Supported Operators	Supported Values
<i>Host Operating System Major Version</i>	Equal to	<p>The host operating system version in the format X.Y; for example, 8.9, 7.9.</p> <p>The major version is derived from the full version by extracting the first two numbers in the dotted notation. Usually, the host OS version contains five numbers separated by dots; for example, 8.9.0.0.0 would yield a major version of 8.9.</p>
<i>Exadata Software Major Version</i>	Equal to	<p>The Exadata software version in the format X.Y; for example, 21.0, 23.0.</p> <p>The major version is derived from the full version by extracting the first 2 numbers in the dotted notation; for example, if the full version of the storage server is 21.0.0.0.0, the major version is 21.0.</p>
<p><i>Target Properties:</i></p> <ul style="list-style-type: none"> • <i>Contact</i> • <i>Cost Center</i> • <i>CSI</i> • <i>Comment</i> • <i>Department</i> • <i>Downtime Contact</i> • <i>Lifecycle Status</i> • <i>Line of Business</i> • <i>Location</i> • <i>Site</i> 	Equal to	<p>Filters will be populated based on the values set for the target properties.</p> <p>Targets within the hierarchy of the Exadata or Recovery Appliance topology can have different properties set at the target level. If any of the targets in the hierarchy match the filter criteria, the dashboard will include results from the parent Exadata or Recovery Appliance.</p>
<i>Group By</i>	Equal to	<p>This option is available only for the Exadata Fleet Capacity and Exadata Component Capacity dashboards. It helps to categorize results based on the selection, and renders them with different colors. For instance, if <i>Group By = Department</i> is set and the results contain three departments (HR, Finance, Sales), then the bubbles on the Exadata Fleet Capacity dashboard is rendered in three colors, each corresponding to one department.</p>

- All filters described above are applied using the AND condition. For example, specifying the filters as *Engineered System Target Name = abc* with *Target Property:Location = Austin* customizes the dashboard content with results that match both these criteria.
- The *Host Operating System Major Version* filter and the *Exadata Software Major Version* filter are applied to the appropriate child component targets under the parent Engineered System target. An Engineered System and all its components are considered a match if

any one of the version filter criteria is met. For example, a filter on *Host Operating System Major Version = 8.7* and *Exadata Software Major Version = 24.1* selects Engineered Systems and all their components if at least one of their hosts has operating system major version 8.7 or if at least one of their storage servers has Exadata software version 24.1.

Exadata Fleet Configuration

The Exadata Fleet Configuration dashboard shows the Exadata hardware and software configurations across the Exadata and Recovery Appliance fleet, based on the last known metric collection. You can apply the filters to limit the display to a subset of systems.



Click the numbers in the center of the ring charts or a ring chart section to see a detailed, in-depth breakdown of the dataset. The image below is an example of the details displayed when the number at the center of the *Engineered System Type* ring chart is clicked:

The expanded view shows a table with the following columns: Engineered System, Engineered System Target Type, Engineered System Deployment Type, Deployment Option, Engineered System Type, Rack Generation, and Switch Type. The table lists various system configurations, including DB Machine, MAA-Manageability-ExaXBM, and ZDLRA Boston systems, with their respective target types, deployment options, and rack generations.

Engineered System	Engineered System Target Type	Engineered System Deployment Type	Deployment Option	Engineered System Type	Rack Generation	Switch Type
DB Machine	Oracle Database Machine	On-Premises	Bare Metal	Exadata	X8M	ROCE
DB Machine	Oracle Database Machine	On-Premises	Virtual Machine	Exadata	X10M	InfiniBand
DB Machine	Oracle Database Machine	On-Premises	Virtual Machine	Exadata	X10M	InfiniBand
DB Machine	Oracle Database Machine	On-Premises	Virtual Machine	Exadata	X10M	InfiniBand
DB Machine	Oracle Database Machine	On-Premises	Virtual Machine	Exadata	X10M	InfiniBand
DB Machine	Oracle Database Machine	On-Premises	Virtual Machine	Exadata	X10M	InfiniBand
DB Machine	Oracle Database Machine	On-Premises	Virtual Machine	Exadata	X10M	InfiniBand
MAA-Manageability-ExaXBM_1	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
MAA-Manageability-ExaXBM_2	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
MAA-Manageability-ExaXBM_3	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
MAA-Manageability-ExaXBM_4	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
MAA-Manageability-ExaXBM_5	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
MAA-Manageability-ExaXBM_6	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
MAA-Manageability-ExaXBM_7	Exadata Infrastructure	Oracle Cloud Infrastructure	Virtual Machine	Exadata	X10M	ROCE
ZDLRA Boston_1	Recovery Appliance	On-Premises	Bare Metal	Recovery Appliance	X4	InfiniBand
ZDLRA Boston_2	Recovery Appliance	On-Premises	Bare Metal	Recovery Appliance	X4	InfiniBand
ZDLRA Boston_3	Recovery Appliance	On-Premises	Bare Metal	Recovery Appliance	X4	InfiniBand
ZDLRA Boston_4	Recovery Appliance	On-Premises	Bare Metal	Recovery Appliance	X4	InfiniBand
ZDLRA Boston_5	Recovery Appliance	On-Premises	Bare Metal	Recovery Appliance	X4	InfiniBand
ZDLRA Boston_6	Recovery Appliance	On-Premises	Bare Metal	Recovery Appliance	X4	InfiniBand

The Exadata Fleet Configuration dashboard has four vertical sections:

- [Overview](#)
- [Hardware](#)
- [Software](#)
- [Database](#)

Overview

The ring charts in this section summarize the basic characteristics of the Engineered System fleet:

- **Engineered System Type:** The distribution of system types. Possible values are *Exadata* or *Recovery Appliance*. *Exadata* includes Database Machine and Exadata Infrastructure targets.
- **Engineered System Deployment Type:** The distribution of deployment types. Possible values are *Cloud@Customer*, *On-Premises*, and *Oracle Cloud Infrastructure*.
- **Engineered System Service Type:** The distribution of cloud service types across Exadata VM Clusters and Autonomous Exadata VM Clusters. Possible values are *Autonomous Database* and *Exadata Database Service*.
- **Database Deployment Option:** The distribution of Exadata deployment options. Possible values are *Virtual Machine* and *Bare Metal*.
- **Switch Type:** The distribution of network and management switch types. Possible values are *Infiniband* and *RoCE*.

Hardware

The ring charts in this section summarize key hardware characteristics of the Engineered System fleet:

- **Exadata Generation:** The distribution of Exadata hardware rack generations; for example, X8-2, X7-2. Database Machines composed of multiple racks of different generations are categorized as *Mixed*.
- **Database Server Model:** The distribution of hardware models of the database servers; for example, *Oracle Corporation ORACLE SERVER X7-2L High Capacity*, *Oracle Corporation ORACLE SERVER X8-2L High Capacity*.
- **Storage Server Model:** The distribution of hardware models of the storage servers; for example, X7, X8.
- **Storage Server Media Type:** The distribution of storage server media types. Possible values are *Extreme Flash* and *High Capacity*.
- **Switch Model:** The distribution of network switch hardware models; for example, *N9K-C9348GC-FXP*, *Sun Datacenter InfiniBand Switch 36*.

Software

The ring charts in this segment summarize key characteristics of the software running on components of the Engineered System fleet:

- **Host Operating System Major Version:** The distribution of host operating system major versions; for example, 8.9, 7.9.
- **Database Server ILOM Version:** The distribution of ILOM server major software versions; for example, 5.1, 6.1.
- **Storage Server Exadata Software Version:** The distribution of storage server major software versions; for example, 26.1, 23.1.

- **Database Server Exadata System Software Version:** The distribution of database server Exadata software major versions; for example, *24.1*, *23.1*.
- **Switch Version:** The distribution of network switch software versions; for example, *10.3*, *10.4*.

Database

The ring charts in this segment summarize key characteristics of databases deployed on the Engineered System fleet:

- **Database Container Type:** The distribution of container and non-container databases. Possible values are *CDB* and *Non-CDB*.
- **Database Version:** The distribution of database release update versions for CDBs and non-CDBs; for example, *19.19*, *23.5*.
- **Pluggable Database Version:** The distribution of database release update versions for pluggable databases; for example, *19.19*, *23.5*.
- **ASM Disk Group Redundancy Type:** The distribution of ASM disk group redundancy levels. Possible values are *High*, *Low*, and *Normal*.
- **Grid Infrastructure Version:** The distribution of grid infrastructure release update versions; for example, *19.19*, *23.5*.

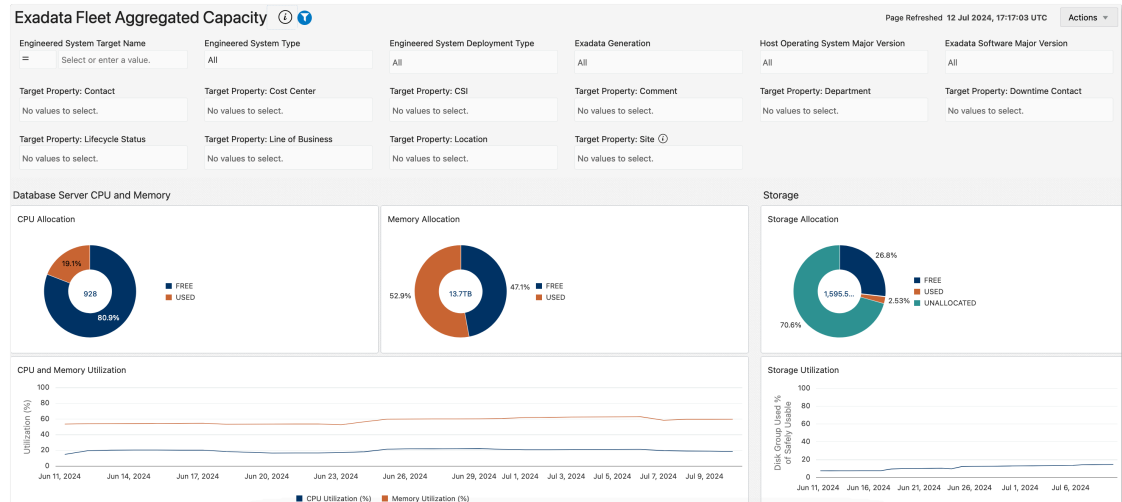
Exadata Fleet Aggregated Capacity

The Exadata Fleet Aggregated Capacity dashboard summarizes the capacity averaged over the last 31 days for the Exadata and Recovery Appliance fleet or filtered subset of the fleet. It provides a high-level summary of fleet resource usage indicating overall average utilization.

This dashboard is organized into two sections:

- [Database Server CPU and Memory, and Storage](#)
- [Storage Server I/O](#)

Database Server CPU and Memory, and Storage



The ring charts in these sections summarize the average CPU, memory and storage resource allocation across the Engineered System fleet over the last 31 days.

- **CPU Allocation:** Shows total, free and used CPU on the database servers
- **Memory Allocation:** Shows total, free, and used memory on the on the database servers
- **Storage Allocation:** Shows total, free and used storage space on the disk groups across all the ASM clusters

The line charts below the ring charts illustrate the utilization trend of these resources for the last 31 days, aggregated across all the Engineered Systems. Analyze these charts in conjunction with each other to understand the resource usage. For instance, if the ring charts show low storage utilization but the corresponding line chart shows peaks of high storage utilization over the last 31 days, it may warrant deeper investigation into whether storage has recently been cleaned up. Further insight as to which systems have contributed to the high usage could be obtained by going to the Exadata Fleet Capacity and Exadata Fleet Component Capacity dashboards. Conversely, if the ring charts show low storage utilization and the line charts also show low average storage utilization over the last 31 days, it may broadly indicate that this subset of Engineered Systems can accommodate more data storage.

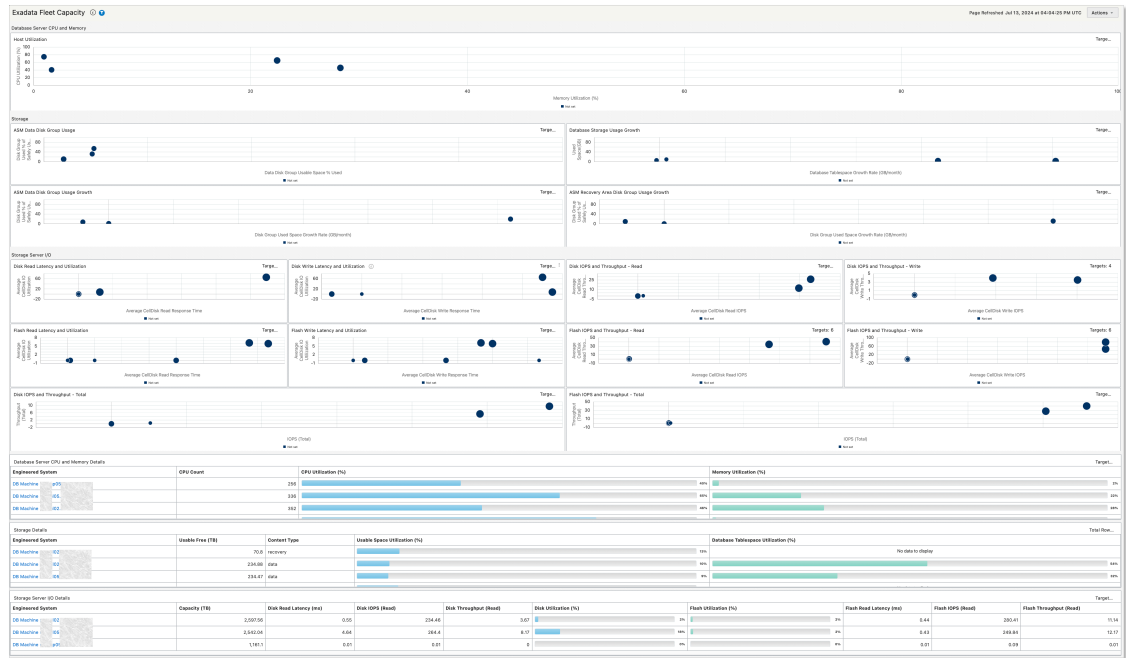
Storage Server I/O

This section consists of line charts that help in visually correlating key storage server performance indicator metrics like *disk read/write latency, throughput and IOPS, I/O load, and I/O utilization*. All the charts in this section have congruous timelines. The metric values on all the charts for a specific time can be viewed together by moving the cursor to a position on any of the charts. Correlating these metrics can be especially helpful in performance diagnostics; for example, to determine whether time periods with high latencies can be correlated to time periods of high I/O load.



Exadata Fleet Capacity

The Exadata Fleet Capacity dashboard shows individual Exadata capacity averaged over the last 31 days across the Exadata and Recovery Appliance fleet or filtered subset of the fleet, allowing comparison by system.




The dashboard includes multiple sections that provide detailed insights into specific dimensions of fleet resource capacity like *host utilization*, *storage utilization*, *I/O performance*, *CPU*, *memory*, *disk usage*, and *IOPS* (input/output operations per second). Detailed descriptions for each section are provided below:

- [Database Server CPU and Memory](#)
- [Storage](#)
- [Storage Server I/O](#)
- [Database Server CPU and Memory Details](#)
- [Storage Details](#)
- [Storage Server I/O Details](#)

Database Server CPU and Memory

The *Host Utilization* bubble chart in this section visualizes the average utilization of CPU and memory across all hosts (database servers) in an Engineered Systems over the past month.

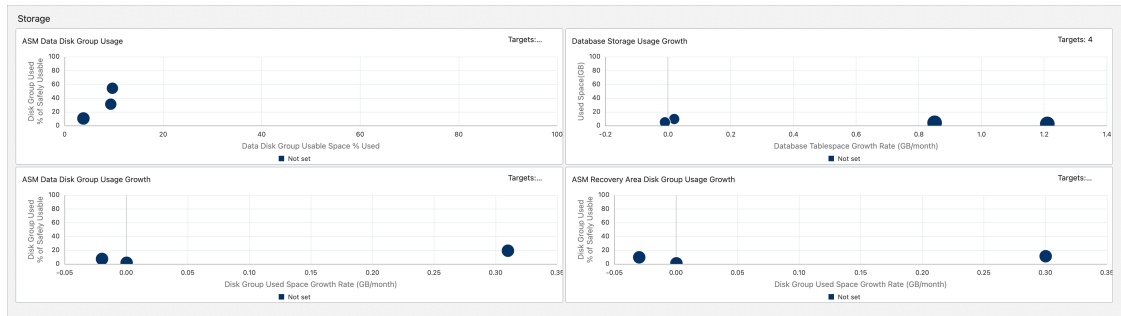


- **X-axis (Memory Utilization %):** Represents the memory utilization percentage, ranging from 0% to 100%
- **Y-axis (CPU Utilization %):** Represents the CPU utilization percentage, ranging from 0% to 100%

- **Bubbles:** Each bubble corresponds to a different Engineered System.
 - The size of the bubble indicates the total CPU count of the database servers in an Engineered System, with larger bubbles representing Engineered Systems with a higher number of CPUs.
 - The color of the bubbles is determined by the *Group By* filter setting.
- **Tooltip:** A tooltip provides detailed information for a specific Engineered System when hovered over. For example, the Database Machine target shown in the example above has a CPU utilization of 65.34%, memory utilization of 22.46%, and a CPU count of 336. The lifecycle status for this target is not set. In all the bubble charts, the tooltip displays the values for the target property as specified in the *Group By* filter. The bubble color is determined by the value of the specified target property.

Storage

This section presents an overview of the storage utilization through four scatter plots, offering insights into disk space usage and growth.



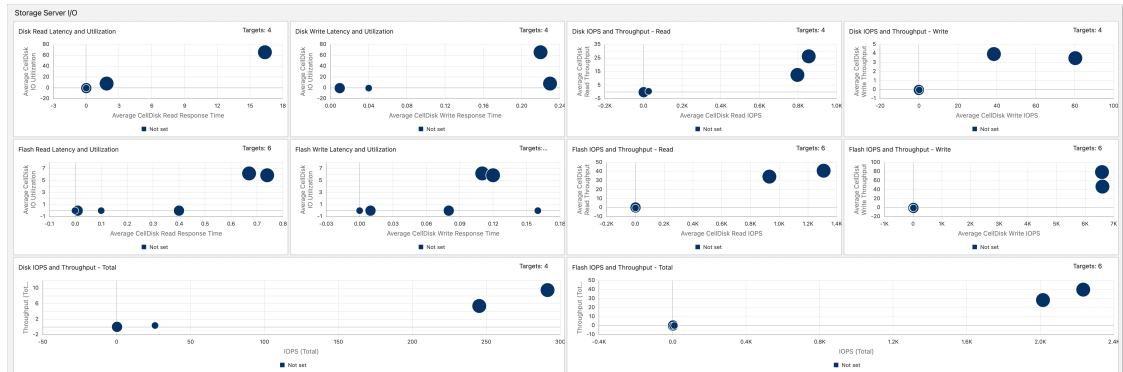
- **ASM Data Disk Group Usage:** The chart shows the average usage across all ASM disk groups used for storing datafiles across all the clusters in an Engineered System over the past month.
 - **Y-axis:** Space used by disk groups as a percentage of safely usable space
 - **X-axis:** Disk group percentage of usable space used
 - **Bubbles:** Each bubble represents an Engineered System. The size of the bubble is proportional to the total capacity of all disk groups used for storing datafiles across all the ASM clusters in the Engineered System.
- **Database Storage Usage Growth:** The chart shows the average database storage usage growth of Engineered Systems in the fleet over the past month.
 - **Y-axis:** Used space in GB
 - **X-axis:** Database tablespace growth rate in GB/month
 - **Bubbles:** Each bubble represents an Engineered System. The size of the bubble is proportional to the allocated storage capacity of all the databases on the Engineered System.
- **ASM Data Disk Group Usage Growth Chart:** The chart shows the average usage growth across all ASM disk groups used for storing datafiles across all the clusters in an Engineered System over the past month.
 - **Y-axis:** Space used by disk groups as a percentage of safely usable space

- **X-axis:** Disk group used space growth rate in GB/month
- **Bubbles:** Each bubble represents an Engineered System. The size of the bubble is proportional to the total capacity of all disk groups used for storing datafiles across all the ASM clusters in the Engineered System.
- **ASM Recovery Area Disk Group Usage Growth:** The chart shows the average usage across all ASM disk groups used for storing recovery files across all the clusters in an Engineered System over the past month.
 - **Y-axis:** Space used by disk groups as a percentage of safely usable space
 - **X-axis:** Disk group used space growth rate in GB/month
 - **Bubbles:** Each bubble represents an Engineered System. The size of the bubble is proportional to the total capacity of all disk groups used for storing recovery files across all the ASM clusters in the Engineered System.

The disk group content type for these charts is determined by examining the value of the ASM content.type attribute: *data* indicates datafiles, *recovery* indicates recovery files.

Storage Server I/O

This section provides an overview of the various performance metrics related to I/O serviced by flash and hard disks across all Engineered System storage servers. Each chart provides detailed insights into different aspects of storage performance, including latency, utilization, IOPS, and throughput. Each bubble on the charts represents an Engineered System. The size of the bubble is proportional to the number of storage servers on the Engineered System.



- **Disk Read Latency and Utilization:** The chart shows the average read latency and utilization values of the storage server disks over the past month.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Read Response Time
- **Disk Write Latency and Utilization:** The chart shows the average write latency and utilization values of the disk storage over the past month.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Write Response Time
- **Disk IOPS and Throughput - Read:** The chart shows the average IOPS and throughput values of read operations on the disk storage over the past month.
 - **Y-axis:** Average Disk Read Throughput
 - **X-axis:** Average Disk Read IOPS

- **Disk IOPS and Throughput - Write:** The chart shows the average IOPS and throughput values of write operations on the disk storage over the past month.
 - **Y-axis:** Average Disk Write Throughput
 - **X-axis:** Average Disk Write IOPS
- **Flash Read Latency and Utilization:** The chart shows the average read latency and utilization values of the flash storage over the past month.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Read Response Time
- **Flash Write Latency and Utilization:** The chart shows the average read latency and utilization values of the flash storage over the past month.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Write Response Time
- **Flash IOPS and Throughput - Read:** The chart shows the average IOPS and throughput values of read operations on the flash storage over the past month.
 - **Y-axis:** Average Disk Read Throughput
 - **X-axis:** Average Disk Read IOPS
- **Flash IOPS and Throughput - Write:** The chart shows the average IOPS and throughput values of write operations on the flash storage over the past month.
 - **Y-axis:** Average Disk Write Throughput
 - **X-axis:** Average Disk Write IOPS
- **Disk IOPS and Throughput - Total:** The chart shows the average IOPS and throughput values for both read and write operations on the disk storage over the past month.
 - **Y-axis:** Throughput (Total)
 - **X-axis:** IOPS (Total)
- **Flash IOPS and Throughput - Total:** The chart shows the average IOPS and throughput values for both read and write operations on the flash storage over the past month.
 - **Y-axis:** Throughput (Total)
 - **X-axis:** IOPS (Total)

Database Server CPU and Memory Details

The **Database Server CPU and Memory Details** table summarizes CPU and memory utilization for Engineered Systems. This table allows comparison of resource usage across different systems, which can be useful for performance monitoring, capacity planning, and identifying potential bottlenecks in CPU and memory utilization.

Database Server CPU and Memory Details				Targets...
Engineered System	CPU Count	CPU Utilization (%)	Memory Utilization (%)	
DB Machine g05	256	<div style="width: 40%;"></div> 40%	<div style="width: 20%;"></div> 20%	
DB Machine g06	336	<div style="width: 60%;"></div> 60%	<div style="width: 22%;"></div> 22%	
DB Machine g02	312	<div style="width: 48%;"></div> 48%	<div style="width: 28%;"></div> 28%	
DB Machine g07	256	<div style="width: 72%;"></div> 72%	<div style="width: 18%;"></div> 18%	

The table contains the following columns:

- **Engineered System:** Name of the Engineered System. Click the name of the target to navigate to its home page.

- **CPU Count:** Number of CPUs for each system
- **CPU Utilization (%):** Average CPU utilization usage, averaged over the last 31 days
- **Memory Utilization (%):** Memory utilization percentage, averaged over the last 31 days

Storage Details

The **Storage Details** table provides an overview of the storage details for Engineered Systems. This table allows comparison of storage utilization across different systems, which can provide insights into free space availability, content distribution, and the efficiency of space utilization in both general storage and specific database tablespaces.

Engineered System	Usable Free (TB)	Content Type	Usable Space Utilization (%)	Database Tablespace Utilization (%)
DB Machine :02	70.8	recovery	13%	No data to display
DB Machine :02	234.88	data	10%	64%
DB Machine :05	234.47	data	9%	92%
DB Machine :05	70.35	recovery	12%	No data to display

The table includes the following columns:

- **Engineered System:** Name of the Engineered System. Click the name of the target to navigate to its home page.
- **Usable Free (TB):** Amount of usable free space available in TB
- **Content Type:** Type of content stored; for example, *data* or *recovery*.
- **Usable Space Utilization (%):** Percentage of usable space utilization
- **Database Tablespace Utilization (%):** Percentage of database tablespace utilization

Storage Server I/O Details

The **Storage Server I/O Details** table provides a detailed overview of I/O performance metrics for various Engineered Systems in tabular format. This table allows comparison of I/O performance and utilization across different systems, which can provide insights into storage efficiency and performance bottlenecks in both disk and flash storage.

Engineered System	Capacity (TB)	Disk Read Latency (ms)	Disk IOPS (Read)	Disk Throughput (Read)	Disk Utilization (%)	Flash Utilization (%)	Flash Read Latency (ms)	Flash IOPS (Read)	Flash Throughput (Read)	Targets...
DB Machine :02	2,597.56	0.55	234.46	3.67	2%	2%	0.44	280.41	11.14	
DB Machine :05	2,542.04	0.64	264.4	8.17	19%	19%	0.43	249.84	12.77	
DB Machine :p05	1,761.1	0.07	0.07	0	0%	0%	0.07	0.09	0.07	

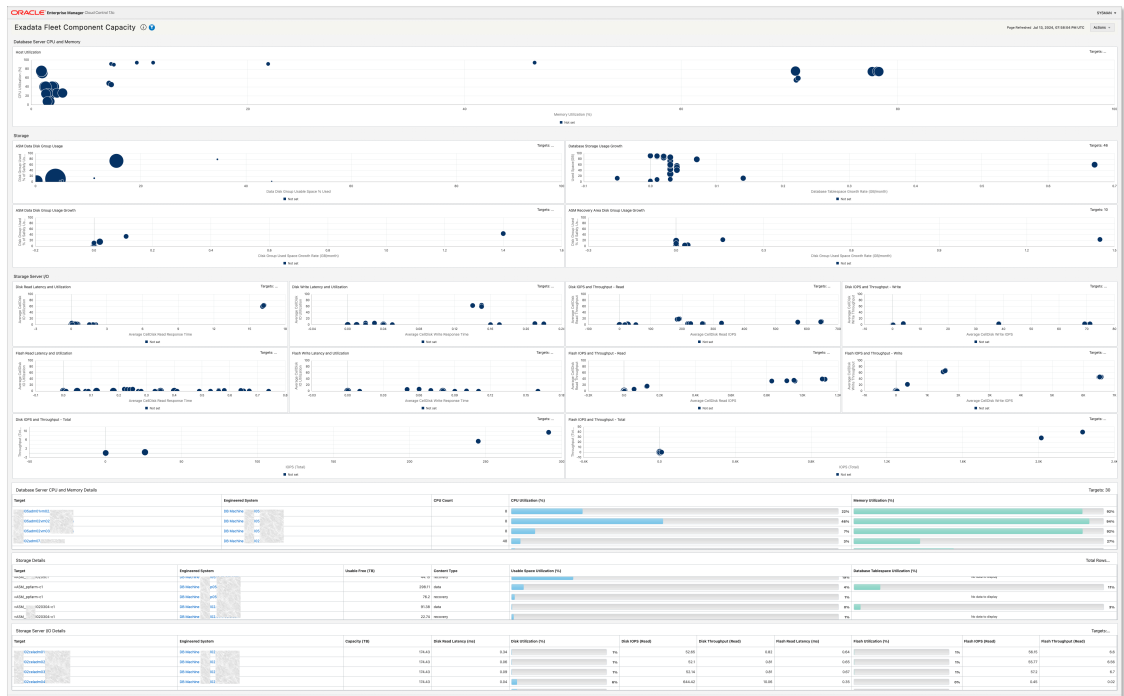
The table includes the following columns:

- **Engineered System:** Name of the Engineered System. Click the name of the target to navigate to its home page.
- **Capacity (TB):** Storage capacity in TB
- **Disk Read Latency (ms):** Latency for disk reads in milliseconds
- **Disk IOPS (Read):** Number of input/output operations per second for disk reads
- **Disk Throughput (Read):** Read throughput for disk storage
- **Disk Utilization (%):** Percentage of disk utilization
- **Flash Read Latency (ms):** Latency for flash reads in milliseconds
- **Flash IOPS (Read):** Number of input/output operations per second for flash reads

- **Flash Throughput (Read):** Read throughput for flash storage
- **Flash Utilization (%):** Percentage of flash utilization

Exadata Fleet Component Capacity

In contrast to the Exadata Fleet Capacity dashboard which focuses on capacities and performance metrics at the Engineered System level, the **Exadata Fleet Component Capacity** dashboard provides a holistic view of fleet capacity and performance metrics at the individual component level. It allows for the identification and comparison of over or under-utilized components, with data averaged over the last 31 days.




The dashboard is organized into sections for database servers, storage, and storage I/O, each focusing on performance aspects of different component types. Detailed descriptions for each section are provided below:

- [Database Server CPU and Memory](#)
- [Storage](#)
- [Storage Server I/O](#)
- [Database Server CPU and Memory Details](#)
- [Storage Details](#)
- [Storage Server I/O Details](#)

Database Server CPU and Memory


This chart shows the average utilization of CPU and memory on the hosts over the past 31 days.



- **Y-axis:** CPU Utilization (%)
- **X-axis:** Memory Utilization (%)
- **Bubbles:** Each bubble represents a host in Engineered System fleet. The size of the bubble is proportional to the number of CPU cores on the host.

Storage

This section presents an overview of the storage utilization through four scatter plots, providing insights into different aspects of disk usage and growth.



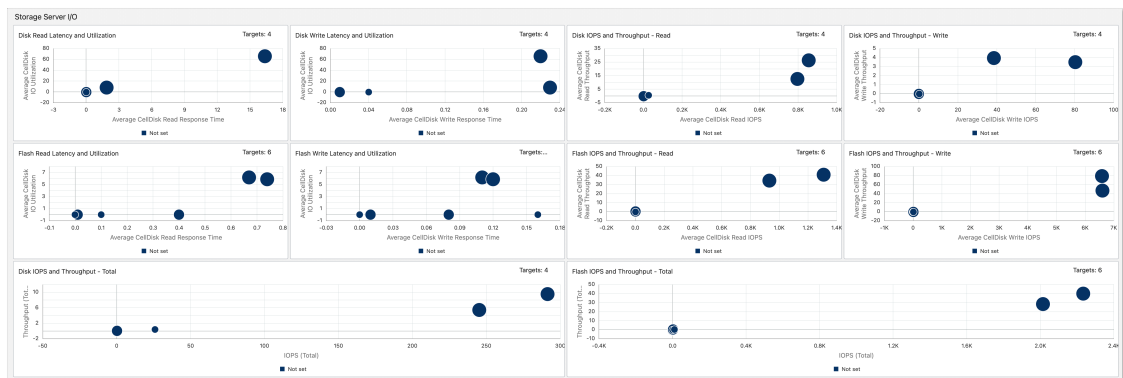
- **ASM Data Disk Group Usage:** The chart shows the average usage across all ASM disk groups used for storing datafiles in the ASM cluster over the past month.
 - **Y-axis:** Space used by disk groups as a percentage of safely usable space
 - **X-axis:** Disk group percentage of usable space used
 - **Bubbles:** Each bubble represents an ASM cluster in the Engineered System fleet. The size of the bubble is proportional to the total capacity of all disk groups used for storing datafiles in the ASM cluster.
- **Database Storage Usage Growth:** The chart shows the average database storage usage growth over the past month.
 - **Y-axis:** Used space in GB
 - **X-axis:** Database tablespace growth rate in GB/month
 - **Bubbles:** Each bubble represents a database in the Engineered Systems fleet. The size of the bubble is proportional to the allocated capacity of the database.
- **ASM Data Disk Group Usage Growth Chart:** The chart shows the average usage growth across all ASM disk groups used for storing datafiles in the ASM cluster over the past month.

- **Y-axis:** Space used by disk groups as a percentage of safely usable space
- **X-axis:** Disk group used space growth rate in GB/month
- **Bubbles:** Each bubble represents an ASM cluster in the Engineered System fleet. The size of the bubble is proportional to the total capacity of all disk groups used for storing datafiles in the ASM cluster.
- **ASM Recovery Area Disk Group Usage Growth:** The chart shows the average usage growth across all ASM disk groups used for storing recovery files in the ASM cluster over the past month.
 - **Y-axis:** Space used by disk groups as a percentage of safely usable space
 - **X-axis:** Disk group used space growth rate in GB/month
 - **Bubbles:** Each bubble represents an ASM cluster in the Engineered System fleet. The size of the bubble is proportional to the total capacity of all disk groups used for storing recovery files in the ASM cluster.

The disk group content type for these charts is determined by examining the value of the ASM content.type attribute: *data* indicates datafiles, *recovery* indicates recovery files.

Storage Server I/O

This section provides an overview of the various performance metrics related to I/O serviced by flash and hard disks of all Engineered System storage servers. Each chart provides insights into different aspects of latency, utilization, IOPS, and throughput. Each bubble represents a storage server in the Engineered System fleet. The size of the bubble is proportional to the storage capacity of the storage server.



- **Disk Read Latency and Utilization:** The chart shows the average read latency and utilization values of the disk storage over the past 31 days.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Read Response Time
- **Disk Write Latency and Utilization:** The chart shows the average write latency and utilization values of the disk storage over the past 31 days.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Write Response Time
- **Disk IOPS and Throughput - Read:** The chart shows the average IOPS and throughput values of read operations on the disk storage over the past 31 days.
 - **Y-axis:** Average Disk Read Throughput

- **X-axis:** Average Disk Read IOPS
- **Disk IOPS and Throughput - Write:** The chart shows the average IOPS and throughput values of write operations on the disk storage over the past month.
 - **Y-axis:** Average Disk Write Throughput
 - **X-axis:** Average Disk Write IOPS
- **Flash Read Latency and Utilization:** The chart shows the average read latency and utilization values of the flash storage over the past 31 days.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Read Response Time
- **Flash Write Latency and Utilization:** The chart shows the average write latency and utilization values of the flash storage over the past 31 days.
 - **Y-axis:** Average Disk I/O Utilization
 - **X-axis:** Average Disk Write Response Time
- **Flash IOPS and Throughput - Read:** The chart shows the average IOPS and throughput values of read operations on the flash storage over the past 31 days.
 - **Y-axis:** Average Disk Read Throughput
 - **X-axis:** Average Disk Read IOPS
- **Flash IOPS and Throughput - Write:** The chart shows the average IOPS and throughput values of write operations on the flash storage over the past 31 days.
 - **Y-axis:** Average Disk Write Throughput
 - **X-axis:** Average Disk Write IOPS
- **Disk IOPS and Throughput - Total:** The chart shows the average IOPS and throughput values of the disk storage over the past 31 days.
 - **Y-axis:** Throughput (Total)
 - **X-axis:** IOPS (Total)
- **Flash IOPS and Throughput - Total:** The chart shows the average IOPS and throughput values of the flash storage over the past 31 days.
 - **Y-axis:** Throughput (Total)
 - **X-axis:** IOPS (Total)

Database Server CPU and Memory Details

The **Database Server CPU and Memory Details** table summarizes CPU and memory utilization of hosts in the Engineered System fleet. This table allows comparison of CPU and memory usage across different hosts, aiding in the identification of systems that may require resource reallocation or optimization. The visual representation of utilization percentages makes it easy to spot potential bottlenecks or under-utilized resources at an individual host level.

Database Server CPU and Memory Details					Targets: 30	
Target	Engineered System	CPU Count	CPU Utilization (%)	Memory Utilization (%)		
I05adm01vm02	DB Machine I05	8	22%	92%		
I05adm02vm02	DB Machine I05	8	46%	94%		
I05adm02vm03	DB Machine I05	8	7%	92%		
I02adm07	DB Machine I02	48	3%	27%		

The table includes the following columns:

- **Target:** Name of the host. Click the name to navigate to the home page of the host.
- **Engineered System:** Name of the Engineered System to which the host belongs. Click the name to navigate to the home page of the target.
- **CPU Count:** Number of CPUs allocated to the host
- **CPU Utilization (%):** Average CPU utilization, averaged over the past 31 days
- **Memory Utilization (%):** Average memory utilization, averaged over the past 31 days

Storage Details

The **Storage Details** table provides an overview of the storage details for the ASM clusters in the Engineered System fleet. This table provides insights into free space availability, content distribution, and the efficiency of space utilization in both general storage and specific database tablespaces at the ASM cluster level. The database targets must be discovered in Oracle Enterprise Manager in order to display details in the *Database Tablespace Utilization* column.

Storage Details ⓘ						Total Rows: 20
Target	Engineered System	Usable Free (TB)	Content Type	Usable Space Utilization (%)	Database Tablespace Utilization (%)	
+ASM_ppfarm-c1	DB Machine p05	298.11	data		4%	
+ASM_ppfarm-c1	DB Machine p05	76.2	recovery		1%	No data to display
+ASM_I020304-c1	DB Machine I02	91.38	data		0%	
+ASM_I020304-c1	DB Machine I02	22.74	recovery		1%	No data to display

The table includes the following columns:

- **Target:** Name of the ASM cluster
- **Engineered System:** Name of the Engineered System to which the cluster belongs. Click the name to navigate to the home page of the target.
- **Usable Free (TB):** Usable free space available in TB
- **Content Type:** Type of content stored. Possible values are *data* and *recovery*.
- **Usable Space Utilization (%):** Percentage of usable space utilization
- **Database Tablespace Utilization (%):** Percentage of database tablespace utilization

Storage Server I/O Details

The **Storage Server I/O Details** table shows the total storage capacity, average latency, average IOPS and average throughput for read operations on the storage servers in the Engineered System fleet over the past 31 days. This table allows comparison of I/O performance and utilization across different Engineered Systems at the individual storage server level, which can provide insights into storage efficiency and performance bottlenecks in both disk and flash storage.

Storage Server I/O Details											Targets: 40
Target	Engineered System	Capacity (TB)	Disk Read Latency (ms)	Disk Utilization (%)	Disk IOPS (Read)	Disk Throughput (Read)	Flash Read Latency (ms)	Flash Utilization (%)	Flash IOPS (Read)	Flash Throughput (Read)	
I02celadm01	DB Machine I02	174.43	0.34		52.65	0.82	0.64		56.15	6.8	
I02celadm02	DB Machine I02	174.43	0.06		52.1	0.81	0.65		55.77	6.56	
I02celadm03	DB Machine I02	174.43	0.09		52.14	0.81	0.67		57.2	6.7	
I02celadm04	DB Machine I02	174.43	0.04		644.42	10.06	0.35		0.45	0.02	

The table includes the following columns:

- **Target:** Name of the storage server. Click the name to navigate to the home page of the target.
- **Engineered System:** Name of the Engineered System to which the storage server belongs. Click the name to navigate to the home page of the target.
- **Capacity (TB):** Storage capacity in TB
- **Disk Read Latency (ms):** Latency for disk reads in milliseconds
- **Disk IOPS (Read):** Number of input/output operations per second for disk reads
- **Disk Throughput (Read):** Read throughput for disk storage
- **Disk Utilization (%):** Percentage of disk utilization
- **Flash Read Latency (ms):** Latency for flash reads in milliseconds
- **Flash IOPS (Read):** Number of input/output operations per second for flash reads
- **Flash Throughput (Read):** Read throughput for flash storage
- **Flash Utilization (%):** Percentage of flash utilization

Engineered System Autonomous Health Framework Administration

Engineered System Autonomous Health Framework (AHF) Administration provides comprehensive monitoring and management of Oracle Autonomous Health Framework and Oracle Exachk across the Engineered System fleet.

These capabilities help the administrators of Exadata and Zero Data Loss Recovery Appliance to configure and maintain their systems according to the best practices, thereby reducing downtime and maximizing performance.

Engineered System AHF Administration provides the following features:

- **Fleet-level AHF and Exachk Monitoring:** A detailed, comprehensive view of AHF installations across the Engineered System fleet (Database Machine, Exadata Infrastructure and Recovery Appliance targets), including the following details:
 - AHF upgrade status and history
 - AHF service status
 - AHF autorun and auto upgrade configuration
 - Exachk run status and history
- **Fleet-level AHF and Exachk Management:** The following management operations can be performed against multiple Database Machine targets in a single job:
 - *Upgrade AHF:* Deploy the latest AHF release on multiple Exadata compute hosts (database servers) on demand, or configure the AHF auto upgrade schedule.
 - *Stage AHF Software for Auto Upgrade:* Download the latest AHF installation zip file to the staging location on multiple Exadata hosts for later use by scheduled auto upgrades.
 - *Configure AHF Upgrade Settings:* Add to or update the existing upgrade settings on multiple Exadata clusters.
 - *Schedule Exachk Runs:* Run on-demand compliance checks or schedule Exachk autoruns on multiple Exadata clusters.

- *Run commands*: Run individual predefined or free-form `ahfctl`, `tfactl` and `exachk` commands on multiple Exadata hosts.

Topics:

- [Prerequisites](#)
- [Enabling Required Metrics](#)
- [Navigating to the AHF Administration Page](#)
- **Monitoring AHF and Exachk for the Engineered System Fleet**
 - [Monitoring AHF and Exachk](#)
 - [Monitoring AHF Jobs](#)
 - [Monitoring Exachk Runs](#)
 - [Monitoring AHF Upgrades](#)
 - [Downloading Upgrade History File](#)
- **Managing Engineered System AHF Installations**
 - [Upgrade AHF](#)
 - [Configure AHF Upgrade Settings](#)
 - [Stage AHF Software for Auto-Upgrade](#)
 - [Run AHF Commands](#)
 - [Schedule Exachk Autorun](#)

Prerequisites

- The AHF Administration page will show data only for Database Machine and Exadata Infrastructure targets that have Exadata Management Pack enabled, and for Recovery Appliance targets that have Zero Data Loss Recovery Appliance Management Pack enabled.
- The Oracle Enterprise Manager Exadata and Recovery Appliance plug-in versions must be *13.5.0.23 or above* on both the Agents and Management Server.
- Oracle Enterprise Manager Agent one-off patch *36795631* must be applied on top of *13.5 Release Update 23* for all Agents monitoring Database Machine targets.
- AHF *23.5 or above* must be installed on all Engineered Systems.
- Oracle Enterprise Manager Self Update must be configured in order to manage AHF software upgrades. See *Setting Up Self Update in Cloud Control Administrator's Guide*.
- The collection must be enabled for the following Oracle Enterprise Manager metrics on all Database Machine, Exadata Infrastructure and Recovery Appliance host targets: `AHF Autoruns Status`, `AHF Status`, `AHF Autoruns`, `AHF Installation`, `AHF Settings`, `AHF Upgrade History`, and `AHF Compliance Results`. All these metrics are contained within a single metric collection. See [Enabling Required Metrics](#).

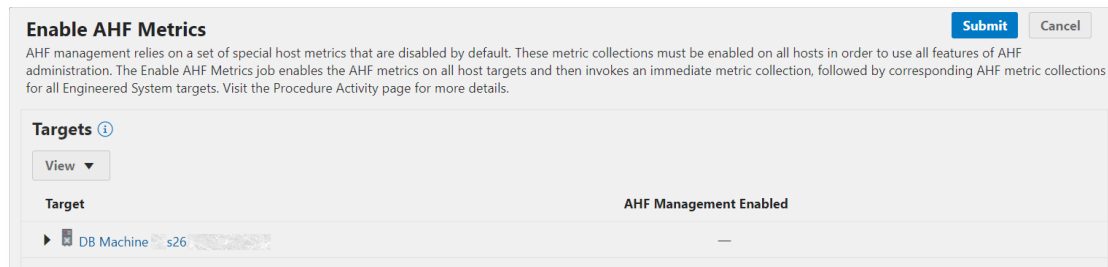
Enabling Required Metrics

AHF monitoring relies on a set of AHF-related metrics that run on Exadata and Recovery Appliance host targets. The metric collection that contains these metrics is disabled by default. In order to use the features of the AHF Administration page, this collection must be enabled on all underlying host targets associated with Database Machine, Exadata Infrastructure, and

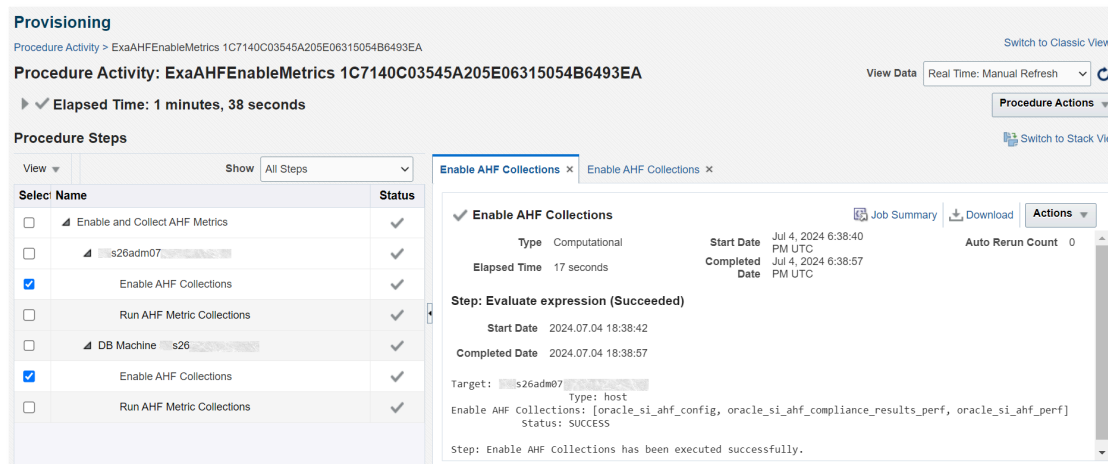
Recovery Appliance targets. This can be accomplished across all targets in one procedure using one of the following methods:

- **Warning message banner:** The AHF Administration page displays a warning message when the required collection is disabled on one or more host targets associated with a monitored Engineered System. Enable the metrics from the warning message banner by clicking **Enable Metrics for AHF Management** and follow the instructions. The list of host targets with disabled AHF metrics can be viewed by clicking **Show targets with disabled metrics**.
- **Actions menu for selected targets:** Navigate to the Engineered Systems tab, select the relevant targets in the table, click **Actions**, select **Enable AHF Metrics** and follow the instructions.

The **Enable AHF Metrics** page opens which lists all the hosts where the metric collection enablement job will be run. The job enables the collections for all host targets of the specified Engineered Systems and performs an on-demand collection on these targets in order to provide immediate functionality on the AHF Administration page.



Click **Submit**. The job ID is displayed when the job is submitted. Click the job ID to navigate to the Procedure Activity page, where the details of the job run are displayed.



Navigating to the AHF Administration Page

To access the AHF Administration page, go to **Enterprise**, click **Compliance**, and select **Engineered System AHF Administration**.


Alternatively, to access the page, from **Targets** menu, select **Exadata**. In the **Actions** menu, select **Engineered System Autonomous Health Framework Administration**.

Monitoring AHF and Exachk for the Engineered System Fleet

Topics:

- [Monitoring AHF and Exachk](#)
- [Monitoring AHF Jobs](#)
- [Monitoring Exachk Runs](#)
- [Monitoring AHF Upgrades](#)
- [Downloading Upgrade History File](#)

Monitoring AHF and Exachk



The **Engineered Systems** tab presents an overview that can be used to determine whether AHF is configured on Engineered System targets according to Oracle-recommended standards. The tab contains the following charts:

- **Total Engineered Systems:** Shows a count of total monitored Engineered System targets, broken down by whether the applicable management pack is enabled or disabled

All the other charts and monitoring data on the page omit Database Machine, Exadata Infrastructure, and Recovery Appliance targets for which the applicable management pack is disabled, and no functionality on the page will be available for such targets.

- **AHF Upgrade Status:** Shows the current upgrade status across the fleet, based on the latest AHF version available in Self Update. The possible status values are as follows:
 - **Must Upgrade:** Hosts with AHF installations older than 180 days
 - **Up-to-date:** Hosts that have the latest AHF version. This is the ideal scenario for all systems.
 - **Upgrade Needed:** Hosts where the AHF version is lower than the version present in Self Update
 - **Others:** Hosts where AHF is not installed or no AHF metrics have been collected on the target

- **AHF Service Status:** Shows the enablement status for the TFA and Exachk services. Use the menu control to specify which service to analyze. It is recommended that all systems have both the TFA and Exachk services enabled.
- **AHF Configurations:** Shows how many systems have auto-upgrade and autorun configurations enabled or disabled. Use the menu control to specify which configuration to analyze. Enabling the auto-upgrade and autorun streamlines the Engineered System maintenance.

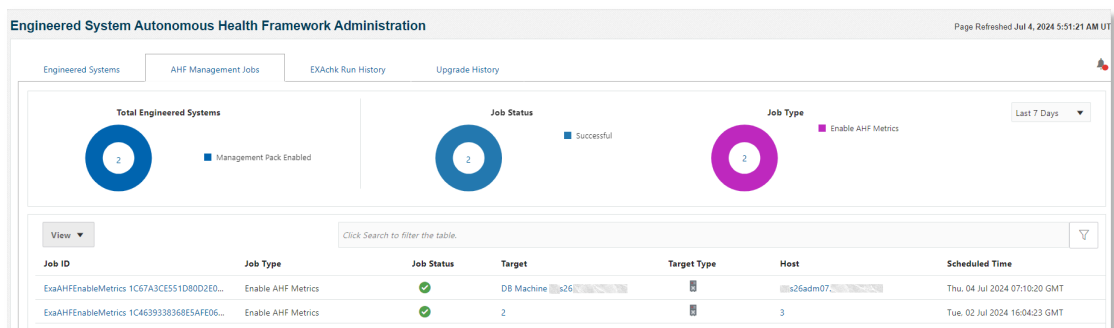
You can click the ring chart sections to filter the table based on the data in the chart.

Details regarding the health and configuration of AHF services running on host targets in the Engineered System hierarchy are aggregated in the table above. The table can be filtered either by clicking sections on the ring charts or through the filter control above the table. The following filter options are available for the table:

- **Target:** Engineered System target name. A sub-string search can be performed using the % wildcard character.
- **Target Type:** Engineered System target type. Possible values are *Oracle Exadata Database Machine, Exadata Infrastructure, and Recovery Appliance.*
- **AHF Version:** AHF version installed on hosts
- **Auto Upgrade:** Auto-upgrade state on the hosts. Possible values are *Enabled and Not Enabled.*
- **Autorun:** Autorun state on the hosts. Possible values are *Enabled and Not Enabled.*
- **Staging Location:** AHF installer staging directory on the hosts. A sub-string search can be performed using the % wildcard character.
- **Exachk Service Status:** Exachk service state on the hosts. Possible values are *Enabled and Not Enabled.*
- **TFA Service Status:** TFA service state on the hosts. Possible values are *Enabled and Not Enabled.*
- **AHF Upgrade Status:** AHF upgrade status on the hosts. Possible values are *Must Upgrade, Up-to-date, Upgrade Needed, and Others.*

Monitoring AHF Jobs

Use the **AHF Management Jobs** tab to view a summary of all the Oracle Enterprise Manager AHF management jobs run across all Engineered System hosts. You can change the job time range to the *last 24 hours, 7 days (default), or 30 days.*



The following ring charts are available in the **AHF Management Jobs** tab:

- **Total Engineered Systems:** Same as the Engineered Systems tab

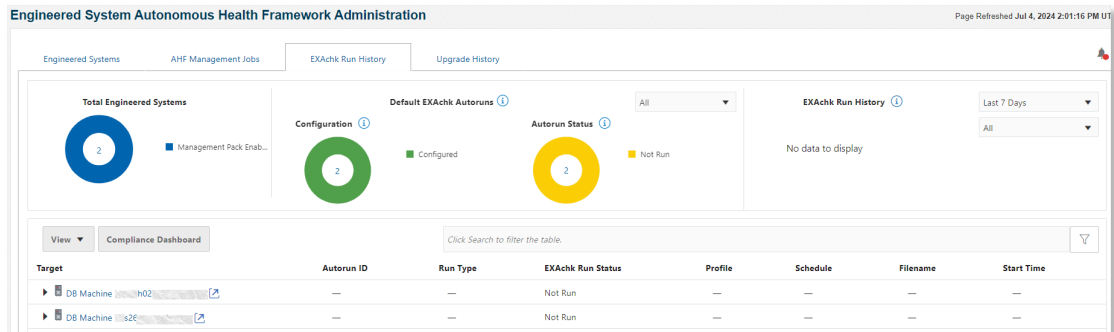
- **Job Status:** Shows the status breakdown of all AHF management jobs that have been run against the Engineered System fleet. The possible status values are:
 - *Successful*
 - *Failed*
 - *Completed with Errors*
 - *Running*
 - *Stopped*
- **Job Type:** Shows a breakdown of the type of AHF management jobs that have been run. The possible job types are:
 - *Upgrade AHF*
 - *Execute AHF Commands*
 - *Upgrade Configuration*
 - *Schedule Exachk Autorun*
 - *Enable AHF Metrics*
 - *Stage AHF Software for Auto Upgrade*

The table below the ring charts shows details about the AHF management jobs that have been run. The table can be filtered either by clicking a section of the ring charts or by using the filter above the table. The following filter options are available for the table:

- **Job ID:** Job ID of the Oracle Enterprise Manager job. A sub-string search can be performed using the % wildcard character.
- **Target:** Engineered System target name. A sub-string search can be performed using the % wildcard character.
- **Target Type:** Engineered System target type. Possible values are *Oracle Exadata Database Machine*, *Exadata Infrastructure*, and *Recovery Appliance*.
- **Host:** Host target where the job was implemented. A sub-string search can be performed using the % wildcard character.
- **Job Status:** Status of the job execution. Possible values are *Successful*, *Failed*, *Completed with Errors*, *Running*, and *Stopped*.
- **Job Type:** Type of job. Possible values are *Upgrade AHF*, *Execute AHF Commands*, *Upgrade Configuration*, *Schedule Exachk Autorun*, *Enable AHF Metrics*, and *Stage AHF Software for Auto Upgrade*.

Monitoring Exachk Runs

The **Exachk Run History** tab contains a history of all autoruns configured across the Engineered System fleet. You can change the run history time range to the *last 24 hours*, *7 days* (default), or *30 days*.



The following ring charts are available in this tab:

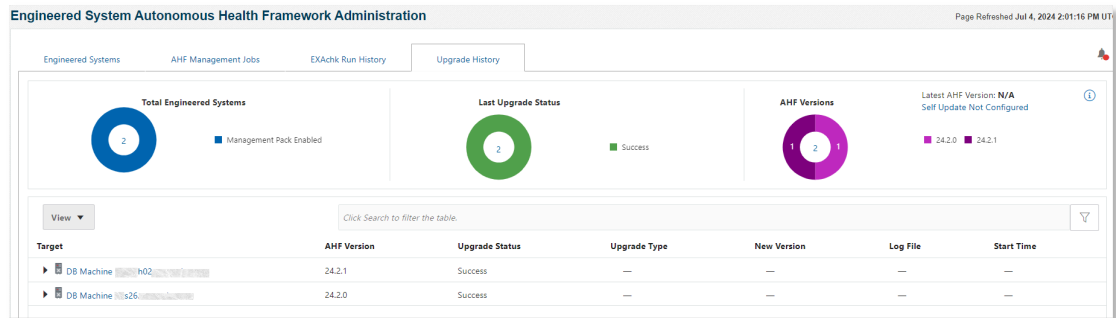
- **Total Engineered Systems:** Same as Engineered Systems tab
- **Default Exachk Autoruns:** Details about the default Oracle-defined Exachk autoruns and their latest run status. Use the menu to filter the results by autorun profile name.
 - **Configuration:** Number of autoruns aggregated at the cluster level within the respective Engineered System. Possible values are as follows:
 - * *Configured:* The selected autorun is configured for all the clusters within the Engineered System.
 - * *Partially Configured:* The selected autorun is not configured for one or more clusters within the Engineered System.
 - **Autorun Status:** Status of Exachk executions in the last 24 hours for the selected default Oracle-defined autorun
- **Exachk Run History:** Status of Exachk runs for on-demand runs, default autoruns, and user-defined autoruns for the selected time frame. Possible values are as follows:
 - *All:* All Exachk runs
 - *Default:* Exachk runs from default autoruns *autostart_client_exatier1* and *autostart_client*
 - *Default (Modified):* Exachk runs from default autoruns that have been modified
 - *User-Defined:* Exachk runs from user-defined autoruns invoked on scheduled intervals
 - *On-Demand:* Exachk runs invoked manually

The table below the ring charts shows details about Exachk runs across host targets in the Engineered System fleet. Some of the values are aggregated at the parent Engineered System target level. You can filter the table by clicking a section on the ring charts or through the filter control above the table. The following filter options are available for the table:

- **Target:** Engineered System target name. A sub-string search can be performed using the % wildcard character.
- **Target Type:** Engineered System target type. Possible values are *Oracle Exadata Database Machine*, *Exadata Infrastructure*, and *Recovery Appliance*.
- **Autorun ID:** ID assigned to the autorun by AHF. A sub-string search can be performed using the % wildcard character.
- **Exachk Run Status:** Status of the Exachk run. Possible values are *Success*, *Failed*, and *Not Run*.
- **Profile:** Name of the profile used for the Exachk run. A sub-string search can be performed using the % wildcard character.

- **Run Type:** Type of the Exachk run. Possible values are *Autoruns (Default)*, *Autoruns (Non-Default)*, and *On-Demand*.

Monitoring AHF Upgrades



The **Upgrade History** tab displays information about the history of AHF upgrades across the Engineered System fleet. The following ring charts are available in this tab:

- **Total Engineered Systems:** Same as Engineered Systems tab
- **Last Upgrade Status:** Shows a breakdown of the latest upgrade status across all Engineered Systems. Possible values are *Success* and *Failed*.
- **AHF Versions:** Shows a breakdown of AHF versions across all Engineered System hosts

The table below the ring charts contains AHF upgrade history details for each host target in the Engineered System hierarchy. Some of the values are aggregated at the parent Engineered System target level. The table can be filtered either by clicking a section on the ring charts or through the filter control above the table. The following filter options are available for the table:

- **Target:** Engineered System target name. A sub-string search can be performed using the % wildcard character.
- **Target Type:** Engineered System target type. Possible values are *Oracle Exadata Database Machine*, *Exadata Infrastructure*, and *Recovery Appliance*.
- **Upgrade Status:** Status of the upgrade operation. Possible values are *Success* and *Failed*.
- **AHF Version:** AHF version installed on hosts
- **Start Time:** Start time of the upgrade operation. Possible values are *Last 24 hours*, *Last 7 days*, and *Last 30 days*.

Downloading Upgrade History File

Upgrade history files can be downloaded for host targets in order to view details of the last AHF upgrade operation. The file for a specific host can be downloaded by clicking on the download icon in the **Log File** column, specifying the host credentials, and clicking the **OK** button. The history file will be downloaded to the device where the browser is running.

Managing Engineered System AHF Installations

Management operations can be performed for Database Machine targets, including both physical and virtual on-premises Exadata systems. Management operations cannot be performed for Exadata Infrastructure or Recovery Appliance targets.

The operations can be initiated for multiple targets from the **Actions** menu above the table in the Engineered Systems tab, or for an individual target by clicking the icon in the **Actions** column for the appropriate row in the table.

All AHF management operations require a named host credential with the user name set to root or non-root user with root privileges. The specified credential is used for all hosts selected for the operation. See Configuring and Using Target Credentials in *Cloud Control Security Guide*.

Topics:

- [Upgrade AHF](#)
- [Configure AHF Upgrade Settings](#)
- [Stage AHF Software for Auto-Upgrade](#)
- [Run AHF Commands](#)
- [Schedule Exachk Autorun](#)

Upgrade AHF

This operation performs an on-demand upgrade of AHF on the hosts of all selected Database Machine targets, using the latest AHF software version present in Oracle Enterprise Manager Self Update. Click **Actions**, select **Upgrade AHF**, and perform the following steps before submitting the operation:

- Specify fleet-level upgrade configuration settings in the **AHF Upgrade Configuration Settings** section. Click **Apply** above the table to provisionally apply the settings to all cluster targets selected in the table. By default, all the cluster targets under a Database Machine are automatically selected.
- You can customize the settings for individual cluster targets by modifying the applicable rows in the table before applying the changes.
- Click **Submit** to submit the Oracle Enterprise Manager job that will perform the configuration settings updates and AHF upgrade on all selected Exadata hosts. The upgrade is performed on all hosts of all Database Machine targets selected in the main page table before invoking the operation. The selections made in the Targets section table decide which cluster targets will have the specified fleet-level or row-level upgrade

configuration changes applied before performing the actual upgrade. Hosts within clusters that did not have their upgrade configuration changed in the table will be upgraded using their existing upgrade configuration settings.

- To view operation step details and progress, click the job ID link in the confirmation dialog to go to the Procedure Activity page.

Configure AHF Upgrade Settings

This operation configures the AHF upgrade settings on the clusters of all selected Database Machine targets, for subsequent use by scheduled auto-upgrades. Click **Actions**, select **Configure AHF**, select **Schedule Exachk Autorun**, and perform the following steps before submitting the operation:

Upgrade Configuration Submit Cancel

Set the AHF upgrade configuration on selected targets. This is a cluster-wide operation that is executed on the master node of the cluster. Select the targets on which the upgrade configuration parameters will be updated, then click the 'Apply' button to assign the settings to those targets. Click the 'Submit' button to submit the Configure Upgrade procedure, which will replace the existing settings on the selected targets with the updated values. Visit the Procedure Activity page to view procedure progress. Once the procedure is successful, execute 'Stage AHF Software for Auto upgrade' action for downloading AHF software from Self Update to configure staging location.

AHF Upgrade Configuration Settings

Auto Upgrade On Off

Staging Location *

Frequency (1-30 days)

Remove Installer Yes No

Credential *

Targets View Apply

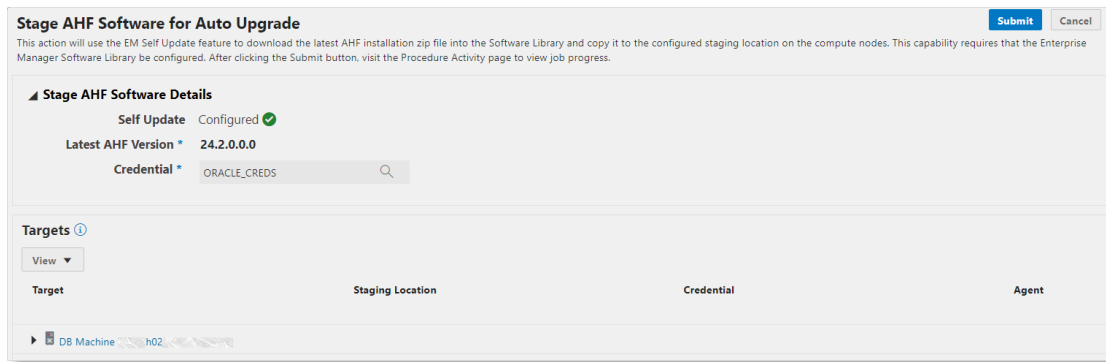
The table shows the current configuration of the targets. To assign the above inputs to selected targets, press the 'Apply' button. Updated configurations will be applied to all selected targets when the operation is submitted.

Target	Credential	Auto Upgrade	Remove Installer	Staging Location
<input checked="" type="checkbox"/> DB Machine	h02			

- Specify fleet-level upgrade settings in the **AHF Upgrade Configuration Settings** section. Click **Apply** above the table to provisionally apply the settings to all cluster targets selected in the table. By default, all the cluster targets under a Database Machine are automatically selected.
- You can customize the settings for individual cluster targets by modifying the applicable rows in the table before applying the changes.
- Click **Submit** to submit the Oracle Enterprise Manager job that will apply the upgrade settings to all the selected clusters. This operation is performed for every cluster target in the Targets table for which settings have been modified, and is run on the master node of those clusters. Clusters that did not have fleet-level settings applied or row-level settings changed will retain their existing settings.
- To view operation step details and progress, click the job ID in the confirmation dialog to go to the Procedure Activity page.
- After successful completion of the operation, the **Stage AHF Software for Auto Upgrade** operation can be invoked to stage the AHF installation in the specified staging areas.

Stage AHF Software for Auto-Upgrade

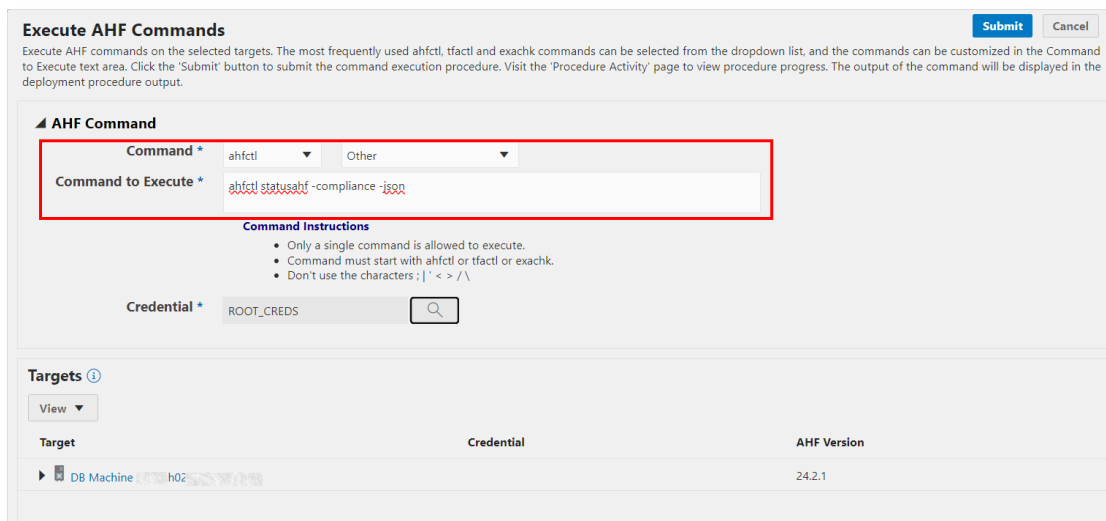
This operation downloads the AHF installation zip file for the latest available AHF version from Oracle Enterprise Manager Self Update to the configured staging location on all hosts of the selected Database Machine targets. This AHF installation zip file is used by subsequent AHF auto-upgrades. Click **Actions**, select **Configure AHF**, and select **Stage AHF Software for Auto Upgrade**.



Click **Submit** to submit the Oracle Enterprise Manager job that will perform the operation on the selected targets. To view details of the operation, click the job ID in the confirmation dialog to go to the Procedure Activity page.

Run AHF Commands

To run custom AHF commands on selected Exadata hosts, click **Actions** and select **Execute AHF Commands**.



You can select the most frequently used `ahfctl`, `tfactl` and `exachk` commands from the menu, and customize them in the **Command to Execute** text area. Click **Submit** to submit the Oracle Enterprise Manager job that will perform the operation on the selected targets. To view detailed operation progress, click the job ID in the confirmation dialog to go to the Procedure Activity page.

For information on command-line options, see [Oracle ORAchk and Oracle Exachk Command-Line Options](#).

Schedule Exachk Autorun

This operation schedules the Exachk autorun compliance checks on the clusters of all selected Database Machine targets.

Schedule EXAchk Autorun Submit Cancel

Schedule EXAchk autorun compliance checks on the selected targets. This is a cluster-wide operation that is executed on the master node of the cluster. Select the targets on which the autoruns will be scheduled, then click the 'Apply' button to assign the new autoruns to those targets. Click the 'Submit' button to submit the Schedule Autorun procedure, which will schedule the new autoruns for the the selected targets. To manage existing autoruns for any specific target, select the number in the Schedule Autoruns column. Visit the Procedure Activity page to view procedure progress.

Recurring EXAchk Details

Autorun ID * test Autorun Schedule * 0 */12 * * * Profile * dba

Tag Enter a String Collection Retention Enter number of days Notification Email Enter Email ID

Autorun Options * AUTORUN_FLAGS=-profile dba:AUTORUN_SCHEDULE=0 */12 * * *;AUTORUN_ID=test

Credential * ROOT_CREDS

Targets

View Apply

The table shows the current configuration of the targets. To assign the above inputs to selected targets, press the 'Apply' button. Updated configurations will be applied to all selected targets when the operation is submitted.

<input checked="" type="checkbox"/>	Target	Credential	Scheduled Autoruns
<input checked="" type="checkbox"/>	DB Machine	h02.	

Click **Actions**, select **Configure AHF**, select **Schedule Exachk Autorun**, and perform the following steps before submitting the operation:

- Specify fleet-level Exachk autorun settings in the **Recurring Exachk Details** section. Click **Apply** above the table to provisionally apply the settings to all cluster targets selected in the table. By default, all the cluster targets under a Database Machine are automatically selected.
- You can customize host credentials for individual cluster targets by modifying the applicable rows in the table before applying the changes.
- To modify or delete the scheduled autoruns for any cluster target, select the number of the autorun to be changed in the **Scheduled Autoruns** column of the Targets table. Follow the instructions in the dialog box below, then click **Submit** to provisionally save the scheduled autorun changes.

Manage Scheduled Autoruns

A scheduled autorun can be updated or deleted. To update an autorun, click the pencil icon in the Autorun Options column, perform the updates, then click the 'Apply' button. To delete, select the check box for the autorun, then click the trash icon at the top of the table. The total numbers of updated and deleted autoruns are noted below the table. When all updates and deletions are complete, click the 'Submit' button to commit all changes. Click the 'Cancel' button to discard all changes.

Target h020102-c2

<input type="checkbox"/>	Autorun ID	Autorun Options
<input type="checkbox"/>	AUTOSTART_CLIENT	AUTORUN_FLAGS=-usediscovery -tag autostart_client -readenvconfig:AUTORUN_SCHEDULE=
<input type="checkbox"/>	AUTOSTART_CLIENT_EXATIER1	AUTORUN_FLAGS=-usediscovery -profile exatier1 -dball -showpass -tag autostart_client_exati

Updated Rows 0
Deleted Rows 0

Submit Cancel

- Click **Submit** to submit an Oracle Enterprise Manager job that will apply the autorun settings and any scheduled autorun changes to all the selected clusters. This operation is performed for every cluster target in the Targets table for which settings have been modified, and is run on the master node of the those clusters. Clusters that did not have fleet-level settings applied or row-level settings changed will retain their existing settings.

- To view operation step details and progress, click the job ID in the confirmation dialog to go to the Procedure Activity page.

Exadata Management Pack Features

Topics:

- [Database Impact Advisor](#)
 - [AHF Configuration for Database Impact Advisor](#)
- [I/O Resource Management Advisor](#)
- [Exadata Cloud Advanced Monitoring](#)
- [Exadata Exascale Advanced Monitoring](#)
- [Exadata Infrastructure Patching Using Fleet Maintenance](#)

Database Impact Advisor

The Database Impact Advisor can be run against an individual Exadata system to perform system-wide database CPU usage *noisy-neighbor* analysis in order to identify databases whose performance is potentially impacted by other databases or other operating system processes. The analysis applies algorithms established by the Autonomous Health Framework (AHF) Balance feature to Oracle Enterprise Manager historical metric data from the past 30 days. Further, the Database Impact Advisor integrates directly with AHF Balance to generate recommendations for optimizing Database Resource Manager (DBRM) settings across all databases to minimize any CPU-based performance impacts that were found.

Database Impact Advisor is available for Database Machine targets for which the Exadata Management Pack is enabled. You can launch it by clicking **Database Machine** and selecting **Database Impact Advisor** from the menu.

Topics:

- [Prerequisites](#)
- [Key Concepts](#)
- [Using Database Impact Advisor](#)
- [Impact Analysis for Cluster Database](#)
- [Impact Analysis for Database Instance](#)
- [Related Database Management Tools](#)
- [Impact Analysis for Hosts](#)
- [Generating AHF Balance Reports](#)

Prerequisites

- Self Update should be configured in Oracle Enterprise Manager. See *Setting Up Self Update* in *Cloud Control Administrator's Guide*.
- Database Impact Advisor uses a specific AHF installation configured and managed from Oracle Enterprise Manager. This AHF installation must be configured by an Oracle Enterprise Manager administrator by following instructions in [AHF Configuration for Database Impact Advisor](#).

- In order to run AHF Balance reports from Database Impact Advisor, Oracle Enterprise Manager users must be granted one of the roles listed in [AHF Configuration for Database Impact Advisor](#).

Key Concepts

Following are some of the core concepts underlying the Database Impact Advisor analysis framework. For additional details about these and the other concepts, see [Resolve Noisy Neighbor Issues](#) in *Oracle Autonomous Health Framework User's Guide*.

- **Limit:** The maximum number of vCPUs a database instance may use simultaneously. The DBRM parameter `CPU_COUNT` implements a limit for the instance.
- **Guarantee:** The number of vCPUs a database instance is guaranteed to be able to use at any time. When a cluster is dedicated to running databases, the DBRM and the operating system cooperate to provide a guarantee. If the over-provisioning ratio $R = \text{sum}(\text{CPU_COUNT}) / \text{physical vCPUs}$, then the guarantee for a database instance is its $\text{CPU_COUNT}/R$.

For example, if we had a 64 vCPU machine running 8 database instances, all with `CPU_COUNT` set to 16, then the oversubscription ratio R would be 2, that is, $8 * 16/64$, and each individual database instance would have a guarantee of 8, that is, $16/2$.

- **Not Exposed Hour:** An hour when no database instance's CPU use exceeds its CPU guarantee. When an instance is not exposed, it cannot experience CPU-based noisy neighbor problems regardless of the CPU consumption of the other instances running on the machine.
- **Exposed Hour:** An hour when one or more database instance's CPU use exceeds its CPU guarantee. When an instance is exposed, it may experience noisy neighbor problems depending on the CPU consumption of the other instances running on the machine.
- **Impacted Hour:** An exposed hour, when the host's CPU utilization exceeded 70% during the hour. When an instance is impacted, it is likely to be experiencing noisy neighbor problems because the total CPU consumption of the machine is high.
- **Partitioned:** When a cluster is partitioned, each database instance has dedicated CPU capacity. CPU consumption by neighbors cannot interfere with a database instance. CPU resources (up to a configured limit - `CPU_COUNT`) are guaranteed to be available at all times. However, since CPU resources are dedicated to specific database instances, instances cannot take advantage of (borrow) CPU cycles that are not being used by other instances. Typically, when a cluster is partitioned, the degree of database consolidation is limited by the number of physical CPUs on each machine in the cluster, and the peak CPU consumption of each database hosted on the cluster.

A cluster is partitioned when the sum of the `CPU_COUNT` DBRM parameter values for all the database instances running on each machine in the cluster is less than or equal to the number of physical CPUs on the machine. For example, if the machines in a cluster each have 64 CPUs, and each machine is hosting 4 database instances, each with `CPU_COUNT` set to 16, the cluster is partitioned.

If the goal is to partition a cluster, then appropriate `CPU_COUNT` settings can be determined by analyzing historical CPU consumption data. AHF Balance supports this analysis.

- **Impacted Status:** Overall impact status of the database. If the database has any impacted hours within the collection then its status is **FAIL**, if it has any exposed hours its status is **WARNING**, else status is **INFO**.

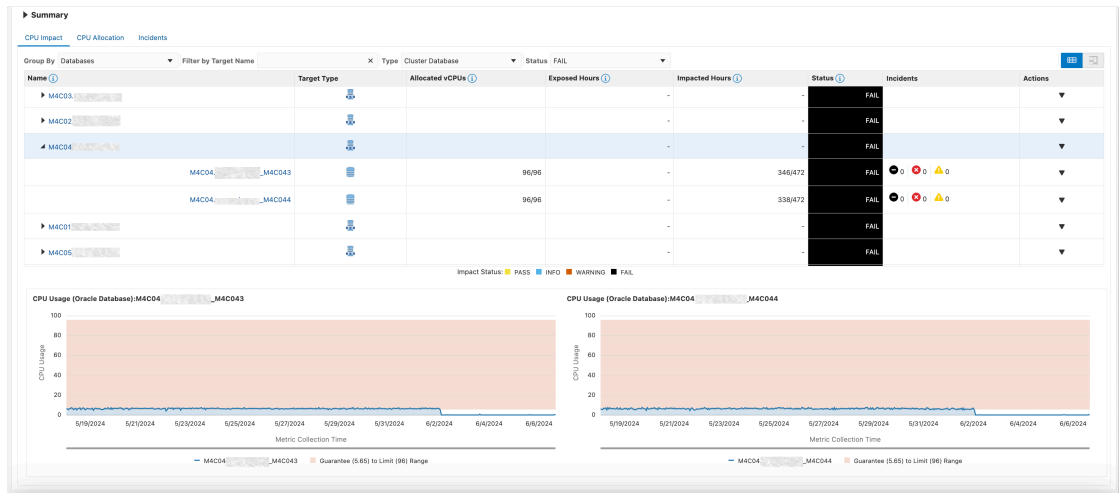
Using Database Impact Advisor

The Database Impact Advisor **CPU Impact** tab has charts that provide a top-level summary of how many clusters, hosts, databases, and database instances on the Exadata system are in

the **Exposed** (*warning*) and **Impacted** (*fail*) categories. The table below the charts enumerates the specific impact status details for each cluster, database, and database instance. Selecting a specific database or instance in the table provides in-depth historical visualization of the exposed and impacted hours for the database or instance.

Impact Analysis for Cluster Database

Click on a cluster database in the table to see its instances. The image below shows a cluster database with two instances, each of which had significant impact over the past month, with both being in **FAIL** (*Exposed*) state. The charts below the table illustrate that for each instance, the CPU usage exceeded the guarantee for a significant period of time.




Impact Analysis for Database Instance

Click on a database instance in the table to see a historical plot of the CPU usage for the instance. Zoom into a time period that needs closer analysis. Impacted or exposed time periods can be easily correlated with incidents that are reported in the same time frame by selecting the **Show Incidents** option.



Select the **Highlight Impacted Hours** option on the **CPU Usage with Guarantees and Limits (Oracle Database)** chart to see when the CPU usage of the database instance is above 70% (red dotted-line) of total host CPU utilization.

To compare the CPU usage of the selected database instance against other instances (from other databases) running on the host, select the **Show All Instance Usage Stacked** option. If the aggregate stacked CPU usage of all instances is above the 70% threshold, it may indicate the presence of noisy-neighbor instances on the host, which can be further analyzed by deselecting the **Show All Instance Usage Stacked** option in order to examine the CPU usage of each individual instance running on the host. Conversely, if the aggregate stacked CPU usage of all instances running on the host is below 70% of total host CPU utilization during a period when the instance is impacted, meaning that the 70% threshold was met only after the addition of non-database process CPU usage, it may indicate that non-database processes are an impact factor.



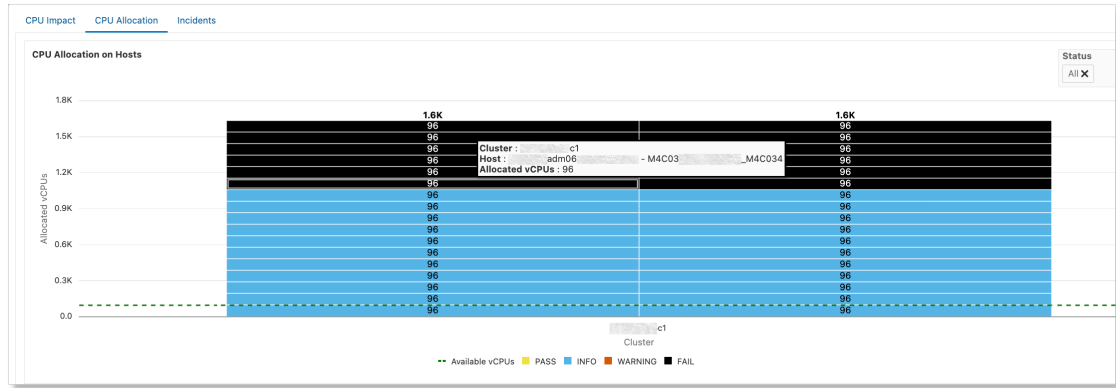
Related Database Management Tools

The **Actions** menu in the **CPU Impact** tab provides access to database management functionality that may be useful in the context of CPU impact mitigation. **Performance Hub** can be accessed to perform fine-grained analysis of performance issues for an individual database or database instance, and **Edit Database Initialization Parameters** can be accessed to tune relevant initialization parameters.

Name	Target Type	Allocated vCPUs	Exposed Hours	Impacted Hours	Status	Incidents	Actions
M4C03		96/96	-	-	FAIL	0	Generate AHF Balance Report, Open AHF Balance Report, Open Performance Hub, Edit Database Initialization Parameters
M4C033		96/96	472/472	-	FAIL	0	
M4C031		96/96	472/472	-	FAIL	0	
M4C032		96/96	472/472	-	FAIL	0	
M4C034		96/96	455/472	-	FAIL	0	

Impact Analysis for Hosts

Use the **CPU Allocation** tab to get a visual indication of how the CPU resources on the Exadata database servers are allocated. The following image shows a cluster with two hosts, each with 96 virtual CPUs (vCPUs). Several databases were deployed on the system, and each database was allocated 96 vCPUs, making it an over-subscribed system. Of these, only a few databases are actually impacted, and are depicted in the **FAIL** state.



Generating AHF Balance Reports

Analyzing the charts described above may indicate that the CPU usage of one or more databases is impacted by other databases on the same Exadata system. The Database Impact Advisor provides the ability to generate AHF Balance reports for clusters, cluster databases and database instances and identify noisy-neighbors. To generate these reports, AHF must be installed. See [AHF Configuration for Database Impact Advisor](#).

Use the **Actions** menu in the **CPU Impact** tab to generate a detailed AHF Balance report for the relevant target. Once the report is ready, a clickable icon is displayed to access the report within Oracle Enterprise Manager. The report provides a stacked time line of impacts across database instances and includes information about disaster recovery scenarios if standby databases are configured. The report can be saved for future access. The **Action Plan** section contains recommended actions that can be implemented to mitigate any CPU impacts found.

Action Plan

Summary

The recommended CPU_COUNT values differ from the current values as follows:


- 42 databases should have their CPU_COUNT reduced, then
- 18 databases should have their CPU_COUNT increased
- 27 databases do not need to have their CPU_COUNT modified

CPU_COUNT Reductions and Increases

First, reduce the CPU_COUNT for the databases listed in the left table. If this work must be done incrementally, make the changes in the order shown in the table. Next, increase the CPU_COUNT for the databases listed in the right table. Again, if the work must be done incrementally, make the changes in the order shown in the table.

Reductions		
	Current	Recommended
database-4368	32	2
database-4342	32	2
database-4372	32	5
database-4324	32	8
database-4282	24	3
database-4176	32	11
database-4208	32	12

Additions		
	Current	Recommended
database-4248	2	3
database-4364	2	3
database-4178	2	3
database-4194	2	3
database-4200	2	3
database-4228	2	3
database-4376	2	3



Database Details

Databases in cluster-1										
Scenario	Current					Disaster				
	Role	vCPUs	Exposed(hrs)	Impacted(hrs)	Status	Role	vCPUs	Exposed(hrs)	Impacted(hrs)	Status
database-309		2/88	121/792	0/792	WARNING		2/88	121/792	1/792	FAIL
database-55	Standby	2/88	58/792	0/792	WARNING	Primary	2/88	89/792	2/792	FAIL
database-333		2/88	26/792	0/792	WARNING		2/88	26/792	0/792	WARNING
database-319		24/88	18/792	0/792	WARNING		24/88	18/792	0/792	WARNING
database-312	Primary	2/88	17/792	0/792	WARNING	Primary	2/88	18/792	0/792	WARNING
database-336		10/88	14/792	0/792	WARNING		10/88	14/792	0/792	WARNING
database-311	Primary	2/88	13/792	0/792	WARNING	Primary	2/88	18/792	0/792	WARNING
database-315		2/88	10/792	0/792	WARNING		2/88	10/792	1/792	FAIL

AHF Configuration for Database Impact Advisor

The Database Impact Advisor uses a specific Autonomous Health Framework (AHF) installation configured and managed from Oracle Enterprise Manager. This AHF installation is used to run AHF Balance reports for clusters and cluster databases.

Topics:

- [Prerequisites](#)
- [Manage the AHF Installation](#)

Prerequisites

- Grant required roles:
 - There are two EM roles used for controlling the access to the AHF configuration:
 - * **EM_AHF_ADMIN** is required in order to perform and update the AHF installation. Oracle Enterprise Manager users with this role have full read and write privilege on the AHF installation and can manage it on behalf of other users.

- * **EM_AHF_USER** is required to view AHF Installation details. Oracle Enterprise Manager users with this role cannot manage the AHF installation, but can access all AHF functionality used by the Database Impact Advisor. All users that intend to use Database Impact Advisor with any Database Machine target should have this privilege granted.
 - To grant the above roles to users, go to **Enterprise**, select **Setup, Security**, and go to **Administrators** page. Both roles are automatically granted to Oracle Enterprise Manager super administrators.
- Oracle Enterprise Manager Self Update should be configured. If Self Update is configured in offline mode, the AHF installation binaries should be downloaded and imported into the Oracle Enterprise Manager Software Library. For information on configuring the Software Library and importing downloaded software in offline mode, see Understanding Software Library in *Cloud Control Extensibility Programmer's Guide*.
- AHF can be installed on any host monitored by Oracle Enterprise Manager that meets the requirement that the host has SQLNet connectivity to the management repository database. The AHF commands invoked on the host make a remote connection to the repository database..
- An Oracle Enterprise Manager default (global) or target-specific host named credential is required. The credential must specify an operating system user who can login to the selected installation host and has write permission on the specified AHF installation directory or the parent directory. For details on how to create named host credentials, see Configuring and Using Target Credentials in *Cloud Control Security Guide*.

Manage the AHF Installation

Navigate to **Enterprise**, select **Setup**, and click **Autonomous Health Framework**.

The Autonomous Health Framework page is displayed. The page is fully functional for EM_AHF_ADMIN users and read-only for EM_AHF_USER users.

The **Install** option is enabled and **Upgrade** is disabled by default when AHF is not configured, or if the configuration is reset.

Perform a new AHF installation using the following steps:

1. Select **Use Existing AHF Installation** if an AHF installation already exists on the host. This can either be installed outside of Oracle Enterprise Manager or from a previous installation performed from Oracle Enterprise Manager. The installation procedure validates the presence of an existing installation on the specified host and location, and is automatically upgraded if the version in the Oracle Enterprise Manager Software Library is of a higher version.
2. Specify a host that meets the prerequisites on which to install AHF or (if Use Existing AHF Installation was selected) one with an existing AHF installation. By default, one of the

Oracle Enterprise Manager management server hosts is selected since these hosts by definition have access to the Oracle Enterprise Manager repository database.

3. Specify the desired AHF installation location. The directory is created if it doesn't already exist.
4. Select a named host credential that meets the prerequisites.
5. Click **Submit** to install AHF on the selected host and location. A message is displayed with the name of the installation deployment procedure.
6. After the installation deployment procedure is complete, a notification is displayed at the top of the page showing the installation status (success or failure).

Upgrade an existing AHF installation using the following steps:

After AHF is successfully installed, the **Upgrade** option gets enabled. The **Submit** button is disabled if there is no change in the host named credentials or if *Latest AHF Version* matches *Installed AHF Version*.

The host credentials used for invoking AHF commands can be changed by specifying a different host named credential, in which case the **Submit** button is enabled.

I/O Resource Management Advisor

The Enterprise Manager **I/O Resource Management Advisor** facilitates effective usage of Exadata I/O Resource Management (IORM) by providing tools to analyze flash I/O utilization across all databases running on an Exadata system and to identify potential measures to optimize system I/O resource usage. It provides the following capabilities:

- **Comprehensive System-Level I/O Overview:** Different visualizations at multiple levels, including treemaps, tables and metric charts, provide a comprehensive overview of IORM settings and a summary of I/O utilization across the Exadata system. These visualizations allow identification of any container database (CDB), pluggable database (PDB), or non-container database that is experiencing flash I/O throttling due to system limits or I/O usage by neighboring databases.
- **Detailed Database-Level I/O Analysis:** Once a potentially impacted database is identified, additional drill-down visualizations allow analysis of the I/O performance history of a database and its neighbors across multiple dimensions, including I/O utilization, throughput, and latency. Analysis at this level can assist in determining whether a database is actually experiencing throttling and associated performance degradation, and whether its IORM plan needs to be updated.

Topics:

- [Prerequisites](#)
- [Reference: Concepts](#)
- [Reference: I/O Limits Calculation](#)
- [Using the I/O Resource Management Advisor](#)
- [Identifying Databases Approaching Their I/O Limit](#)
- [Determining I/O Headroom](#)
- [Analyzing Historical Usage of I/O Resources](#)
- [Analyzing Performance Data Distribution](#)
- [Refining the IORM Plan](#)

Prerequisites

The following requirements must be met in order to use the I/O Resource Management Advisor:

- Enable the Exadata Management Pack for the relevant Database Machine and/or Exadata Infrastructure targets.
- The Oracle Enterprise Manager Exadata plug-in version must be *13.5.0.23 or above* on both the Agent and Management Server.
- Every storage server in the Exadata system must have a software version *21.1.0.0.0 or above*.
- All storage servers in Exadata system should share the same interdatabase IORM plan, and should be discovered in Oracle Enterprise Manager.
- The `Exadata Database Flash Disk I/O Statistics` metric must be enabled for the Oracle Enterprise Manager *Exadata Storage Server Grid* target associated with the Exadata system.

Reference: Concepts

- **Exadata System:** Either an on-premises Exadata represented as an *Oracle Exadata Database Machine* target in Oracle Enterprise Manager or cloud-based Exadata *Oracle Exadata Database Service on Cloud@Customer* and *Oracle Exadata Database Service on Dedicated Infrastructure*, both represented by as an *Exadata Infrastructure* target in Oracle Enterprise Manager.
- **Peak Flash I/O Utilization:** The maximum of average I/O utilization for a database or PDB across all its servicing storage servers in the last 24 hours. For example, DB1 is serviced by three storage servers (S1, S2, S3). The average flash I/O utilization of DB1 over the last 24 hours on these 3 storage servers is 50%, 51% and 52% respectively. The peak flash I/O utilization for DB1 in this case is 52%.
- **Flash I/O Utilization Thresholds:** The following thresholds are built in to the I/O Resource Management Advisor:
 - *Critical:* Flash I/O utilization exceeds 75% of database I/O limit
 - *Warning:* Flash I/O utilization exceeds 50% of database I/O limit
- **Database Flash I/O Limit:** The I/O limit for a database as specified in the I/O plan that governs the database, which could be an interdatabase plan (CDBs and non-CDBs) or CDB plan (PDBs). For more information regarding I/O plan concepts, see [Managing I/O Plans](#).

Reference: I/O Limits Calculation

Note the following details about how I/O limits are calculated so that I/O Resource Management Advisor data can be interpreted correctly:

- **Database Flash I/O Limit:**
 - If the governing plan flash I/O limit value is 0, the storage server interprets this as there being no restrictions on the I/O usage, and therefore the I/O limit for that database is effectively 100% of the storage server's capacity.
 - If the governing plan flash I/O limit value is a non-zero value, the storage server uses this as the actual limit value.
- **Pluggable Database Flash I/O Limit:**

- If the governing plan flash I/O limit value is 0, the storage server interprets this as there being no restrictions on the I/O usage, and therefore the I/O limit for that database is effectively 100% of the CDB's flash I/O limit.
- If the governing plan flash I/O limit value is a non-zero value, the storage server uses this as the actual limit value, calculated as a fraction of the CDB flash I/O limit.

Also, I/O Resource Management Advisor focuses on flash I/O utilization only because from Exadata version 19.x and above, IORM plan limits are applicable to flash I/O utilization only. See [Using the limit attribute](#).

Using the I/O Resource Management Advisor

I/O Resource Management Advisor is available for Database Machine and Exadata Infrastructure targets. You can launch it by clicking **I/O Resource Management Advisor** menu from the Database Machine or Exadata Infrastructure menu.

Whether or not an IORM plan is already in place, the I/O Resource Management Advisor provides valuable insights to assist in managing and optimizing system I/O resources, including the following scenarios:

- Identify the databases with high I/O utilization and determine whether they need more I/O bandwidth.
- Identify the databases that consume more I/O than desired so that their consumption can be throttled.
- Analyze historical I/O resource usage to understand whether periods of high I/O usage are concurrent and possibly causing performance degradation.

Identifying Databases Approaching Their I/O Limit

The **Overview** section provides a summary of the I/O situation for all the databases running on the Exadata system, including the databases that are not monitored by Oracle Enterprise Manager:

- **Databases Approaching I/O Utilization Limit:** The total number of databases/CDBs for which the flash I/O utilization crossed the flash I/O limit critical threshold in the last 24 hours
- **Pluggable Databases Approaching I/O Utilization Limit:** The total number of PDBs for which the flash I/O utilization crossed the flash I/O limit critical threshold in the last 24 hours



The heatmap in the **Database Peak Flash I/O Utilization and Effective I/O Limit for Last 24 Hours** section provides a visual indication of I/O utilization for all databases, and can be used to quickly identify databases that are under or over utilizing their allocated I/O resources. The heatmap tiles, one per database, provide the following visual indicators:

- **Tile Size:** Proportional to database flash I/O utilization
- **Tile Color:** Reflects peak database flash I/O utilization over the last 24 hours, as follows:
 - Red (Critical): The I/O utilization exceeded the built-in critical threshold (75%).
 - Yellow (Warning): The I/O utilization was between the built-in warning (50%) and critical (75%) thresholds.
 - Green (Clear): The I/O utilization was under the built-in warning threshold.

The heatmap offers two sets of view options to organize the data: **View By** and **Color By**.

Following are the **View By** options:

- **Databases/CDBs:** View peak flash I/O utilization at the database or CDB level. Use this option if the Exadata system has both CDBs and non-container single-instance and/or cluster databases (non-CDBs).
- **PDB:** View peak flash I/O utilization at the PDB level. This is useful if the Exadata system has only CDBs (and therefore PDBs).

The **Color By** option controls the level at which to apply the peak flash I/O utilization comparison when rendering the heatmap. Following are the **Color By** options:

1. **Database/CDB Utilization:** The color is determined by comparing the aggregated peak flash I/O utilization at the CDB level (or non-CDB level, if applicable).
2. **PDB Utilization:** The color is determined by the highest value of the peak PDB flash I/O utilization across all PDBs in the CDB.
3. **Highest Utilization in Container:** The color is determined by the highest flash I/O utilization reached between the two scenarios above. This is the most comprehensive option, as it helps to visualize I/O hotspots, whether it is at the non-CDB/CDB level or the PDB level.

Click a tile in the heatmap to view the utilization metrics for the corresponding database in the table in the **Database Peak Flash I/O Utilization and Effective I/O Limit for Last 24 Hours** sections.

Usage Example: Container database ASM1.CDB1 has three PDBs, of which only one ASM1.CDB1.PDB2 has a lower I/O limit. The peak I/O utilization of the CDB is low but the peak I/O utilization of ASM1.CDB1.PDB2 has reached its limit at least once over the past 24 hours. The tile for this CDB would show as green if **Color By** was set to **Database/CDB Utilization**, but would show as red if set to **Highest Utilization in Container**.

A large number of databases or PDBs in critical range does not necessarily indicate a problem. The primary goal is to ensure that the most important databases (for example, those serving production applications) are displayed in green or yellow. If the important databases are in the critical (red) range, then remediation measures could be considered, such as increasing their I/O limit, re-examining how the databases are consolidated, or re-analyzing the I/O workload. From the perspective of overall resource management, it may help to understand whether the Exadata system has more *headroom*, the capacity to handle more I/O workload.

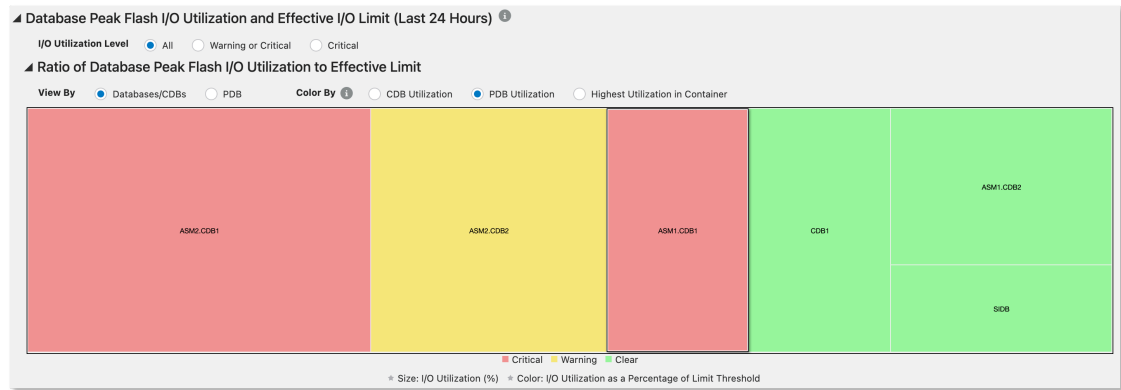
Determining I/O Headroom

Once the databases to focus on are identified, you may want to analyze if they have capacity to handle larger I/O workloads. This can be done by comparing the peak flash I/O utilization against the effective I/O limit of these databases to determine if they have sufficient headroom. In addition to the heatmap described above, the **Database Peak Flash I/O Utilization and Effective I/O Limit for Last 24 Hours** section also displays a hierarchical table view of the databases running on the Exadata system. The I/O utilization values displayed in the table are the peak values for the corresponding row within the last 24 hours.

While correlating this table with the heatmap to analyze CDBs in the critical range, this hierarchical view can be used to identify PDBs that are reaching their respective I/O limits within those CDBs. For example, consider the following scenario using the container database ASM1.CDB1:

- **View By** is specified as `Databases/CDBs` and **Color By** is specified as `CDB Utilization`. The tile for ASM.CDB1 would be green because its peak I/O utilization is 10%, which is far below the warning threshold of 50%.
- If **Color By** is changed to `PDB Utilization`, the color of the ASM.CDB1 tile would change from green to red. The table shows why the color changed - One of the PDBs (ASM1.CDB1.PDB2) in the ASM1.CDB1 reached its I/O limit, and hence crossed the critical threshold of 75%.
- If **Color By** is changed to `Highest Utilization in Container`, the ASM1.CDB1 tile would still be red because the peak I/O utilization of ASM.CDB1 is considered along with the peak I/O utilization of its three PDBs in it, and at least one of the PDBs (ASM1.CDB1.PDB2) crossed the critical threshold.

If the low I/O limit for pluggable database ASM1.CDB1.PDB2 is a concern, its I/O limit can be increased using the CDB plan.



Database Peak Flash I/O Utilization and Effective I/O Limit for Last 24 Hours

Select Database Name Synchronize Charts

Database Name	Database Type	I/O Utilization (%)	I/O Limit (%)	I/O Shares	Profile
AS1.CDB1	CDB	10.01	100	32	gold
AS1.CDB1.PDB1	PDB	5.98	100	32	—
AS1.CDB1.PDB2	PDB	4.67	5	32	—
AS1.CDB1.PDB3	PDB	5.78	100	32	—
AS1.CDB2	CDB	8.99	100	32	gold
AS2.CDB1	CDB	21.62	25	16	silver
AS2.CDB2	CDB	14.94	25	16	silver
CDB1	DB	8.99	50	16	silver
SIDB	DB	4.99	50	16	silver


Legend: ■ I/O Utilization ■ Above Warning Threshold ■ Above Critical Threshold □ I/O Limit

By default, the historical performance charts are rendered for all the databases running on the Exadata system. One or more rows in the table above can be selected to limit the charts to only the selected databases.

- Select a CDB and click **Synchronize Charts** to view the historical performance charts for all its PDBs.
- Select one or more PDBs in a CDB and click **Synchronize Charts** to view the historical performance charts for the selected PDBs.
- Select a non-CDB database and click **Synchronize Chart** to view the historical performance statistics for that database.

Analyzing Historical Usage of I/O Resources

After reviewing the table above, you can examine the historical trends of the metrics to understand whether the peak I/O utilization was seen frequently, and whether these periods of high I/O utilization coincide with high response times and/or high IORM wait times. You can select a single non-CDB, CDB or PDB to analyze its metrics, and these can be stacked against the same metrics for neighboring databases to understand how they compare with each other.



The **I/O Utilization History** section shows a set of line charts that help correlate performance metrics related to I/O. By default, the historical performance charts are displayed for last 24 hours. Zoom in and out using the time selector above the charts to focus on a specific time period.

The following controls are available in this section:

- **Select Time Frame** option allows customization of the time range used for the charts, which determines the granularity of the chart data. The available values are Last Hour, Last 24 Hours, Last 7 Days, Last 31 Days, and Custom Time Period.
- Use **Reset Charts** button to reset the filters:
 - Turning this switch on resets the charts to show data for all the databases/CDBs present in the table above.
 - Turning this switch off reverts the charts back to the state they were in prior to turning the switch on. (That is, show data only for the specific databases selected in the table.)
- **Chart Type** controls the chart style - *Line Chart* or *Stacked Chart*.
- **Show Range** controls the display of the threshold values (50% and 75%) along with the defined I/O limit on the I/O utilization chart. This option is enabled only when a single CDB or PDB is selected in the table. It is especially useful to see the range when Chart Type is specified as Stacked Chart, as it helps to see how the aggregate IOPS usage of all databases compares against these thresholds.
- **Other PDBs in Same Container** shows an aggregation of historical data values for all the PDBs that are not selected in the table. The aggregation line is plotted as the *Others* series, alongside the lines for the selected PDBs.
- **Metric Type** allows the following additional performance statistics to be displayed: *I/O Utilization*, *I/O Requests per Second*, *I/O Throughput*.

Analyzing Performance Data Distribution

For an Exadata system with a large number of databases, the **Performance Data Distribution** section provides a quick view of the I/O utilization at a specific point in time. By default, the bubble chart in this section displays data for all databases selected in the table above, with the data corresponding to the last data point in the historical charts in the I/O Utilization History section. The data point for the bubble chart can be changed by clicking on a specific time in the historical charts. Each bubble in the chart depicts three dimensions of database I/O performance:

- The size of the bubble represents the IORM wait time.
- X-axis represents small I/O response time.
- Y-axis represents I/O utilization, I/O requests per second or I/O throughput based on the selection in *Metric Type*.

For instance, if *Metric Type* is specified as I/O Utilization (%), the databases with high IORM wait time, high small I/O response time and high I/O utilization are displayed as larger bubbles on the top right quadrant of the chart. These are the databases that probably need the most attention as potential I/O performance degradation is indicated in multiple dimensions.



Refining the IORM Plan

Combining the analysis from both the table data and charts provides a comprehensive understanding of the I/O utilization of the databases running on the Exadata system. If it is determined that the I/O allocations need to change, you can do it in one of the following ways:

1. **Change the IORM plan on the storage servers:** If all the databases in the I/O Utilization history charts are stacked while showing the range and that indicates that there is significant headroom on the storage servers, you may want to increase the I/O limits and/or increase the number of I/O shares for the relevant databases. This can be accomplished by navigating to the Database Machine **IORM Management** page and editing the IORM plan. Oracle Enterprise Manager will then propagate this IORM plan to all the storage servers in the Exadata system.
2. **Reconfigure the CDB plan from the root PDB:** See [Creating a CDB Plan](#).

If all the databases are running close to their I/O limits for most of the time or are reaching peak utilization during the same time frames, there are multiple potential solutions that can be considered depending on their feasibility, including the following:

- Increase the number of storage servers on the Exadata system.
- Provision additional VMs, create additional clusters, and move databases to different clusters.
- Move PDBs to different CDBs.
- Move databases to another Exadata system.

Exadata Cloud Advanced Monitoring

The following advanced inventory and resource usage monitoring features are available for Exadata Infrastructure targets:

- **Database Topology on Exadata VM Clusters:** Hierarchical diagrams and tables to visualize the topology of the Oracle Database versions, software installations, and deployments on the VM cluster.

- **Resource Usage on Exadata Infrastructure:** Historical CPU and memory usage visualization across all VM clusters on an Exadata Infrastructure that can be used to identify periods of high resource consumption and discern patterns of usage over different time intervals.

This content is available on the Exadata Infrastructure and associated VM Cluster home pages if the Exadata Management Pack is enabled for the Exascale Infrastructure target. These features are not available for Autonomous VM clusters.

Topics:

- [Prerequisites](#)
- **[Visualizing Database Topology on Exadata VM Clusters](#)**
 - [Visualizing Database Inventory](#)
 - [Analyzing Database Inventory Details](#)
- **[Visualizing Resource Usage for Exadata Infrastructure](#)**
 - [Monitoring Database CPU Usage](#)
 - [Database Memory Usage](#)

Prerequisites

- Enable Exadata Management Pack for the relevant Exadata Infrastructure targets.
- The Oracle Enterprise Manager Exadata plug-in version must be *13.5.0.23 or above* on both the Agent and Management Server. For Exadata Cloud, see Agent Installation in *Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud*.

Visualizing Database Topology on Exadata VM Clusters

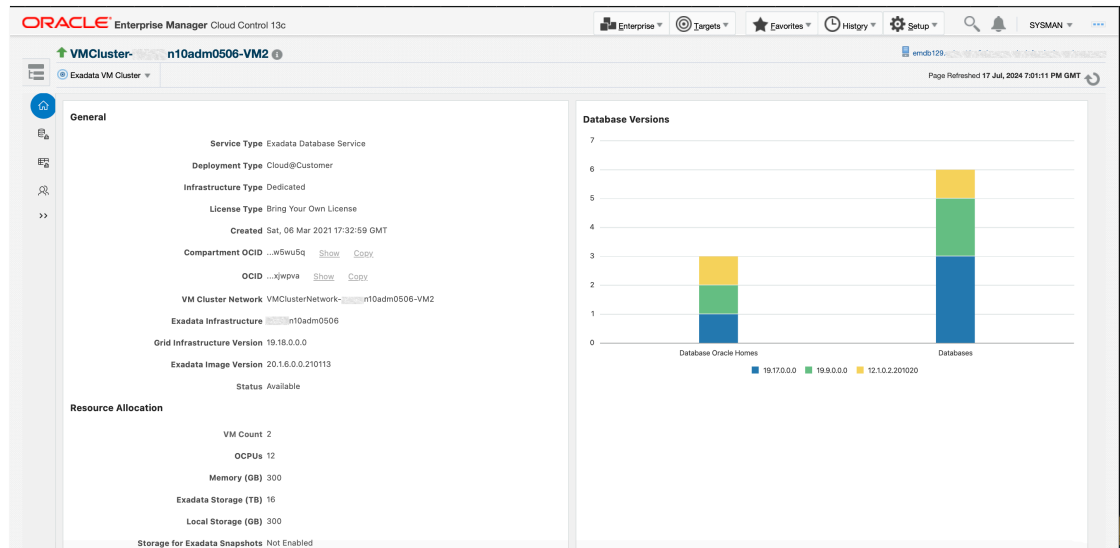
The VM Cluster target home page includes two additional database topology sections on the **Databases** tab when the Exadata Management Pack is enabled for the associated Exadata Infrastructure target:

- Database Inventory
- Database Inventory Details

Topics:

- [Visualizing Database Inventory](#)
- [Analyzing Database Inventory Details](#)

For steps to navigate to the VM Cluster home page, see Discover Exadata Cloud Targets in *Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud*.



Visualizing Database Inventory

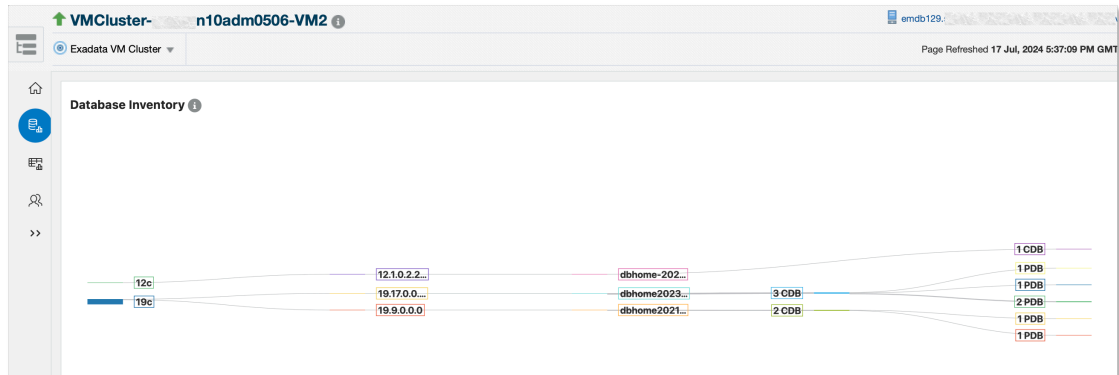
The **Database Inventory** section of the Databases tab contains a Sankey diagram that traces the database topology of the VM cluster from the database versions at the *root*, through the *Oracle homes*, down to the aggregate counts of databases residing within the Oracle homes. The data for the resource hierarchy is obtained using Oracle Cloud Infrastructure REST API. See [Manage VM Clusters](#) in *Exadata Database Service on Dedicated Infrastructure*.

The following nodes and links are available in the Sankey diagram:

- **Nodes:** Each node represents a logical database classification:
 - **Major Database Version:** Represents major releases of the Oracle Database; for example, versions *12c*, *19c*, *23ai*, and so on.
 - **Minor Database Version:** Represents Oracle Database release updates and patch levels within a major release; for example, *19.9.0.0.0*, *12.1.0.2.2*.
 - **Oracle Home:** The installed Oracle home for software binaries and related files for a specific database version; for example, *dbhome-2021*, *dbhome-2023*. For more information, see [Create Oracle Database Homes on an Exadata Cloud Infrastructure System](#).
 - **Database/CDB Count:** The aggregate number of container databases (CDBs) that can host pluggable databases (PDBs), or non-container databases, associated with their respective Oracle home. This count includes all the CDBs and non-CDBs on the VM cluster, even if they are not discovered in Oracle Enterprise Manager.
 - **PDB Count:** The number of PDBs within the associated CDBs
- **Links :** Each link represents the relationship between the nodes it connects. The width of the link is proportional to the count of components associated with the link, so a thicker link indicates that there are more components represented by the link.
 - **Initial link:** Shows the total number of databases belonging to a particular major version
 - **Major Database Version to Minor Versions:** Shows the relationship between major and minor versions of the Oracle Database
 - **Minor Database Version to Oracle Home:** Shows the association of a specific version to the corresponding Oracle homes

- **Oracle Home to Database/CDB Count:** Shows the association of Oracle homes with the databases deployed in the homes
- **CDB Count to PDB Count:** Shows the relationship from the CDBs to the PDBs within each CDB
- **Terminal link:** Shows the count of PDBs, as deployed in Oracle Cloud

The diagram in the example below illustrates that there are a large number of databases with major version *19c* as compared to databases with major version *12c*, because the initial link for *19c* is thicker. It also illustrates that two CDBs are associated to Oracle home *dbhome-2021*, and each CDB has one PDB associated with it.



Analyzing Database Inventory Details

Selecting individual nodes in the Sankey diagram described above interactively filters the table in the **Database Inventory Details** section, providing a more detailed analysis of specific sub-sections of the database topology. The table includes the following columns:

- **Hierarchy:** Represents the nested relationships of the major database versions, minor versions, Oracle homes, databases/CDBs, and PDBs, as applicable.
 - The last one or two nodes in the hierarchy display the full name of the database/CDB or PDB.
 - * If a database/CDB is discovered in EM, the name is hyper-linked to its home page.
 - * If a database/CDB is not discovered in EM, a warning icon is displayed next to its name indicating that the target is not managed in Oracle Enterprise Manager. All PDBs underneath such a CDB will be counted as unmanaged. The warning icon is bubbled up to its highest level ancestor in the hierarchy to help identify databases that are not monitored by Oracle Enterprise Manager.
 - The filter applies to the **Hierarchy** column, and allows searches by sub-strings.
- **Version:** Represents the minor version of the respective Oracle home, database/CDB and PDBs shown in the row
- **Home:** Represents the Oracle home for the database, CDB or PDB shown in the row
- **CDB:** If the row corresponds to a PDB, this column shows the name of the corresponding CDB.

The example below illustrates the hierarchy when major version node *19c* is expanded. This table complements the Sankey diagram by showing details of the two CDBs in the *19c* version hierarchy. Since one CDB *DB19DEV* and its PDB *DB19DEV_..._PDB19* are discovered in

Oracle Enterprise Manager, links are provided to the respective home pages. Since the other CDB are not discovered, its PDB is shown as *Unmanaged PDB(1)*.

Hierarchy	Version	Home	CDB
19c			
19.17.0.0.0			
19.9.0.0.0			
dbhome202103081324	19.9.0.0.0		
DB19DEV	19.9.0.0.0	dbhome202103081324	
DB19DEV_PDB19	19.9.0.0.0	dbhome202103081324	DB19DEV
DB223003_9	19.9.0.0.0	dbhome202103081324	
Unmanaged PDB(1)	19.9.0.0.0	dbhome202103081324	DB223003_9
12c			
12.1.0.2.201020			
dbhome-202404162024	12.1.0.2.201020		
DB12_	12.1.0.2.201020	dbhome-202404162024	


Visualizing Resource Usage for Exadata Infrastructure

The Exadata Infrastructure target home page includes additional monitoring features in the **Resource Usage** tab when the Exadata Management Pack is enabled for the Exadata Infrastructure target. The tables in the **CPU** and **Memory** sub-tabs provide graphical usage visualizations in the last two usage-related columns for the selected time range.

Topics:

- [Monitoring Database CPU Usage](#)
- [Database Memory Usage](#)


To ensure that the visualizations provide a holistic picture of database CPU and memory usage, ensure that all the hosts and databases on the Exadata Infrastructure are discovered, and all metric collections are enabled. These features are not available for Autonomous VM clusters. See *Monitor Exadata Cloud Service in Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud*.



Monitoring Database CPU Usage

In the table on the **CPU** sub-tab, the **Database CPU Usage** columns contain line charts corresponding to each VM Cluster that depict the historical CPU usage over the selected time interval:

- The X-axis represents the time interval.
- The Y-axis represents the CPU usage.
- The blue line represents the average CPU usage of all databases across all hosts in the VM cluster.
- The green line represents the average CPU usage of all hosts in the VM cluster.

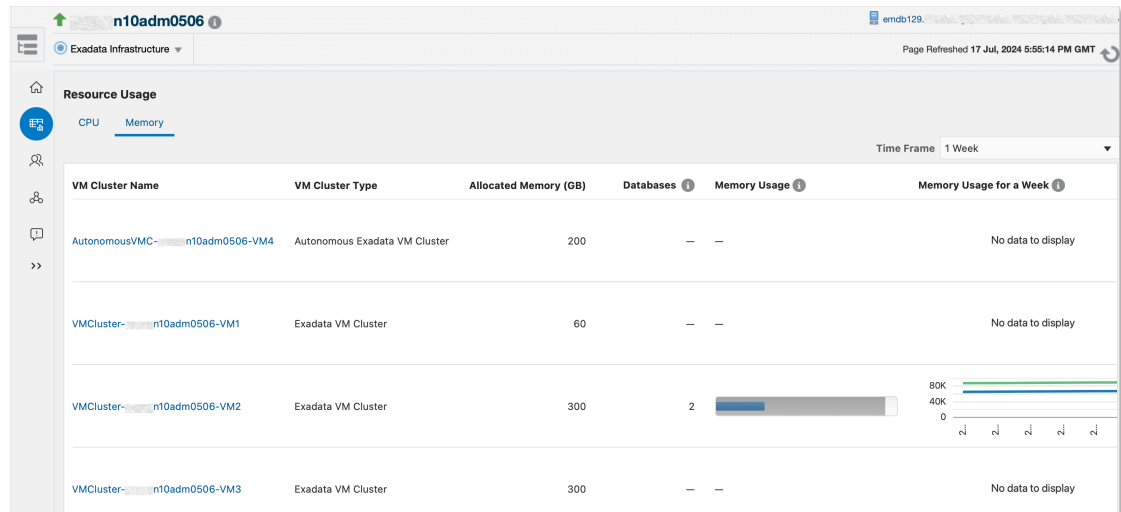


These charts can be customized to display CPU usage for the selected time period. Since the CPU usage is aggregated across all the components of the VM cluster, this chart can be used to compare CPU usage of the databases against the usage of other processes running on the system (such as application servers, business applications, and so on) that could potentially impact database performance. These charts facilitate analysis of relative historical trends and patterns in CPU consumption, providing insights into how resource usage fluctuates over time.

They can also aid consolidation planning by helping to identify periods when CPU usage is concurrently high or low across VM clusters.

Database Memory Usage


In the table on the **Memory** sub-tab, the **Memory Usage** columns depict historical memory usage through line charts with the same characteristics as the CPU table described above.



Exadata Exascale Advanced Monitoring

Advanced performance monitoring consists of an interactive Sankey-based illustration of IOPS allocations from the Exascale cluster down through vaults, storage pools and databases. The historical IOPS usage charts at the storage pool and vault levels allow granular IOPS usage analysis and identification of IOPS distribution issues at different levels. Advanced filtering options provide the ability to visually identify IOPS bottlenecks, view them in context with historical Exascale IOPS utilization, and adjust IOPS distribution for specific entities. These capabilities assist in maintaining maximum I/O performance for databases that use Exascale storage.


The Exascale home page **Performance** tab has two standard monitoring sections - **IOPS** and **IOPS Allocation**. For details, see Monitoring Exascale in *Oracle Exadata Database Machine Getting Started Guide*.



The Performance tab also offers advanced performance monitoring functionality that is part of the Exadata Management Pack. This advanced performance monitoring is available only if the pack is enabled for the Database Machine target that is the parent of the Exascale target.

The advanced performance monitoring functionality has a Sankey-based diagram that displays the Exascale IOPS distribution. It has the following nodes and links:

- **Exascale Cluster:** Shows the name of the Exascale cluster. On mouse-over, IOPS information is displayed.
- **Storage Pool:** Shows all storage pools (color-coded) that are provisioned in this cluster. On mouse-over of a storage pool, information about provisionable, provisioned, and raw IOPS for the storage pool is displayed.
- **Vaults:** Shows vaults with links to the storage pool which provisions IOPS for each vault. On mouse-over of a vault, information about provisioned IOPS for the vault is displayed.
- **Database:** The terminal nodes show the databases that are consuming the IOPS.
- **Links:** The link between two nodes shows how the IOPS are provisioned. On mouse-over, the provisioned IOPS is displayed.



You can use the following options to filter the data in the Sankey diagram:

- **Metric Period:** The metric data time span is 24 hours or 7 days.
- **Vault IOPS Exceeds:** The threshold value to use for determining whether IOPS has exceeded within the specified metric period. The values are prepopulated with the Oracle Enterprise Manager *Vault Performance* metric threshold values, if set.

- **Number of Occurrences:** The number of times IOPS has exceeded the specified threshold over the metric period. Once the values for Metric Period, Vault IOPS Exceeds, and Number of Occurrences are specified, the vaults that exceed the specified threshold for the specified number of occurrences in the specified metric period are identified.
- **View:** `All` can be specified to display all vaults and to highlight vaults that match the above filter criteria. `Highlighted` can be specified to view only the highlighted vaults.

Once the vaults of interest are identified, click a node (which is an Exascale cluster, storage pool or vault) in the diagram to obtain the chart for the aggregate IOPS consumed at the selected level over the specified time period, broken out into separate charts for the different media types that are used. The charts displays the following information:

- **IOPS (High Capacity):** Read and write IOPS for High Capacity storage.
- **IOPS (Extreme Flash):** Read and write IOPS for Extreme Flash storage.
- **IOPS (Flash Cache):** Read and write IOPS for flash cache.

You can select the following reference lines to display in the charts:

- **IOPS Reference:** The IOPS threshold selected in the filter
- **IOPS Available:** Maximum IOPS available

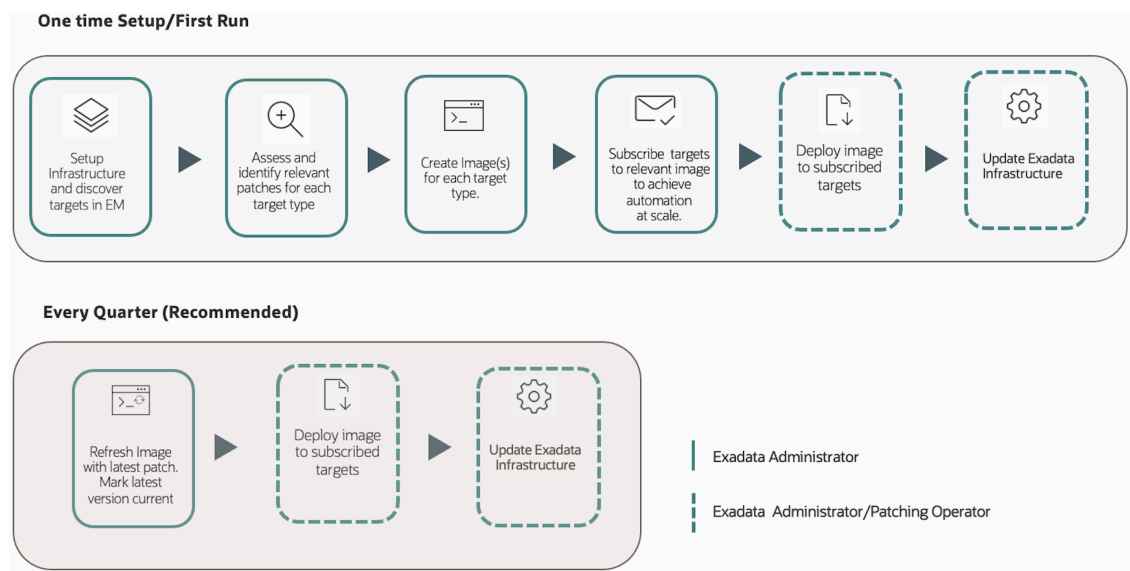
Exadata Infrastructure Patching Using Fleet Maintenance

Manually maintaining an increasing number of Exadata infrastructure components to address security vulnerabilities, bug fixes, and enhancements is time-consuming and risky. Exadata Fleet Maintenance is a solution designed to simplify and streamline the patching process for your entire Exadata infrastructure, ensuring optimal performance, security, and uptime.

Before you begin

The following is an overview diagram of the process flow to patch Exadata infrastructure

Figure 3-1 Exadata Infrastructure Patching Overview



This chapter covers the Exadata patching lifecycle with the following sections:

- [Set up Infrastructure for Exadata Patching](#)
- [Image Maintenance](#)
- [Subscribe Exadata Targets](#)
- [Exadata Fleet Maintenance Patching Operations](#)
- [Rollback Exadata Patches](#)

Set up Infrastructure for Exadata Patching

Before you begin patching your Exadata fleet with Fleet Maintenance it is important that it be properly set up.

Prerequisites for Exadata Patching

Before you begin setting up your Exadata infrastructure for Fleet Maintenance patching make sure that all prerequisites are met:

- General Enterprise Manager Prerequisites: Before you begin: [Meeting the Prerequisites](#)
- Ensure that you have properly setup your infrastructure for Enterprise Manager, for detailed information see: [Setting Up Your Infrastructure](#)
- Review and download Exadata patches using MOS Note: [888828.1](#)
- Patch manager location: [21634633](#)
- Oracle Enterprise Manager 13c Release 5 Update 23 (13.5.0.23) or above
- FPP Server and Client version 19.23 or higher

For more details for setting up and upgrading your FPP server, see [Oracle Fleet Patching and Provisioning Administrator's Guide](#).

Supported Components

Exadata Patching with Fleet Maintenance supports patching multiple Exadata components and targets simultaneously. However, each patch operation must focus on a single component type (e.g., hosts, storage servers).

The following is a list of supported Exadata types for patching:

- Database Server Bare Metal
- Kernel-based Virtual Machine (KVM) Guest
- XEN DomU
- Storage Server
- InfiniBand (IB) Switch

Topics

This chapter cover the following topics:

- [Provision an FPP Server and Client](#)
- [Register FPP Server and Client](#)

Provision an FPP Server and Client

You will need to provision and configure an FPP server and client to be able to successfully patch an Exadata target with Fleet Maintenance.

Provision an FPP Server and Client

You will need to provision both an FPP server and clients, in these examples you can see sample commands to provision both server and client in new Grid Infrastructure or into existing Grid Infrastructure through Enterprise Manager. You can review submitted provisioning activities progress in Enterprise Manager's Procedure Activity page.



Note:

This method is currently supported for Oracle 19c releases only.

- **Provision an FPP Server and Grid Infrastructure**
EM CLI

```
emcli submit_procedure -input_file=data:"/home/user_name/FPP/
gi_input.properties" -name="RACPROV2"
```

The input file must contain the following

```
ASMDBA_GROUP=oinstall
ASMOPER_GROUP=oinstall
ASM_DISK_LISTS=/mnt/emracdev/emracn0607/ASMDISK/disk1,/mnt/emracdev/
emracn0607/ASMDISK/disk2,/mnt/emracdev/emracn0607/ASMDISK/disk3
ASM_DISK_STRINGS=/mnt/emracdev/emracn0607/ASMDISK/disk*/mnt/emracdev/
emracn0607/MGMTDB/disk*
AUTOVIP_ENABLED=false
CLUSTER_NAME=emracn0607
CONFIGURE_ASMFD=false
CONFIGURE_GIMR=true
CONFIGURE_MODE=CONFIGURE_ASM
CVU_CHECK_ERRORMODE=stop
DB_COMPONENT=Fleet/OH19/linux64/DB19_18 - OH19 RU18
DB_HOST_NORMAL_CREDNAMES=CUSER:SYSMAN
DB_HOST_ROOT_CREDNAMES=CUSER_ROOT:SYSMAN
DB_INS_CS_GUID=0
DB_ORACLE_BASE_LOC=/scratch/cuser/19c/racbase
DB_ORACLE_HOME_LOC=/scratch/cuser/19c/rachome
DB_PRODUCT_VERSION=19.0.0.0.0
DEPLOY_MODE=DEPLOY_GI
DISKGROUP_AU_COMPATIBLE=DATA:4:::MGMTDB:4:::
DISK_GROUP_NAMES=DATA,MGMTDB
FLEX_ENABLED=false
GIMR_DB_ORACLE_BASE_LOC=/scratch/cuser/19c/racbase
GIMR_DB_ORACLE_HOME_LOC=/scratch/cuser/19c/rachome
GIMR_DISK_GROUP=MGMTDB
GIMR_TYPE=NOGIMR
GI_AU_SIZE=4
GI_CLUSTER_TYPE=1
```

```
GI_COMPONENT=Fleet/GI1919/linux64/GI_19_19 - GI 19 RU19
GI_HOME_SHARED=false
GI_HOST_NORMAL_CREDNAMES=CUSER:SYSMAN
GI_HOST_ROOT_CREDNAMES=CUSER_ROOT:SYSMAN
GI_ORACLE_BASE_LOC=/scratch/gridbase
GI_ORACLE_HOME_LOC=/scratch/gridhome
GI_PRODUCT_VERSION=19.0.0.0.0
GI_STORAGE_TYPE=1
GNS_ENABLED=true
GNS_VIP_ADDRESS=emracp1011-gns.subnet1rg.emsubnet.companyvcn.com
GRID_DISK_FLAG=true
GROUP_VERSION=19.0.0.0.0
INTERFACE_PARAM_VALUE=ens3:100.75.0.0:1,ens4:192.168.1.0:5
NOT_RELINK_DB_HOME_SELECTED=false
NOT_RELINK_GIMR_DB_HOME_SELECTED=false
NOT_RELINK_GI_HOME_SELECTED=false
OCR_VDSK_DISK_GROUP=DATA
OINSTALL_GROUP=oinstall
OSASM_GROUP=oinstall
OSBACKUPDBA_GROUP=oinstall
OSDBA_GROUP=oinstall
OSDGDBA_GROUP=oinstall
OSKMDBA_GROUP=oinstall
OSOPER_GROUP=oinstall
OSRACDBA_GROUP=oinstall
PAUSE_AFTER_PREREQ=false
PLATFORM_ARUID=226
PREREQ_CHECK_ERRORMODE=continue
RAC_DBCA_MODE=CreateRACDatabase
RAC_HOME_SHARED=false
RAC_STORAGE_TYPE=1
REDUNDANCYS=NORMAL,NORMAL
ROLE_NODENAMES=HUB,HUB
RUN_CVU_FROM_OH=true
SCAN_NAME=emracn0607-r.racdb-scan.emsubnet.companyvcn.com
SCAN_PORT=1521
SKIP_CVU_CHECK=false
SKIP_PREREQUISITE_CHECKS=false
SOURCE_TYPE=SOFTWARE_LIBRARY
SSH_SETUP=false
TARGET_HOST_LIST=emracn07.subnet1rg.emsubnet.companyvcn.com,emracn06.subnet
1rg2.emsubnet.companyvcn.com
VERSION_GREATER_12102=true
VIP_NODENAMES=emracn07-vip.subnet1rg.emsubnet.companyvcn.com,emracn06-
vip.subnet1rg.emsubnet.companyvcn.com
WORK_DIR_LOC=/tmp/
ASM_SYS_PASSWORD=oracle
FPP_CONFIGURATION_TYPE=server
FPP_DISK_GROUP=MGMTDB
FPP_IMAGES_MOUNTPOINT=/scratch/cuser/mount
REST_USER_NAME=restuser
REST_USER_PASSWORD=<password>
REST_USER_EMAIL=user@example.com
```

REST API:

Request Method

POST

URI

https://<OMS_CONSOLE_HOST>:<OMS_CONSOLE_PORT>/em/api/deploymentProcedures/
<IMAGE_ID>/actions/submit

Request Headers

Authorization: Basic
Content-Type:
 application/json

Payload

```
{
  "targetList": [

  ],
  "variables": {
    "items": [
      {
        "variableName": "ASMDBA_GROUP",
        "variableValue": "oinstall"
      },
      {
        "variableName": "ASMOPER_GROUP",
        "variableValue": "oinstall"
      },
      {
        "variableName": "ASM_DISK_LISTS",
        "variableValue": "/mnt/emracdev/emracn0607/ASMDISK/disk1,/mnt/emracdev/
emracn0607/ASMDISK/disk2,/mnt/emracdev/emracn0607/ASMDISK/disk3"
      },
      {
        "variableName": "ASM_DISK_STRINGS",
        "variableValue": "/mnt/emracdev/emracn0607/ASMDISK/disk*/mnt/emracdev/
emracn0607/MGMTDB/disk*"
      },
      {
        "variableName": "AUTOVIP_ENABLED",
        "variableValue": "false"
      },
      {
        "variableName": "CLUSTER_NAME",
        "variableValue": "phxemracp1011"
      },
      {
        "variableName": "CONFIGURE_ASMFD",
        "variableValue": "false"
      },
    ]
  }
}
```

```
"variableName": "CONFIGURE_GIMR",
"variableValue": "true"
},
{
"variableName": "CONFIGURE_MODE",
"variableValue": "CONFIGURE_ASM"
},
{
"variableName": "CVU_CHECK_ERRORMODE",
"variableValue": "stop"
},
{
"variableName": "DB_COMPONENT",
"variableValue": "Fleet/OH19/linux64/DB19_18 - OH19 RU18"
},
{
"variableName": "DB_HOST_NORMAL_CREDNAMES",
"variableValue": "CUSER:SYSMAN"
},
{
"variableName": "DB_HOST_ROOT_CREDNAMES",
"variableValue": "CUSER_ROOT:SYSMAN"
},
{
"variableName": "DB_INS_CS_GUID",
"variableValue": "0"
},
{
"variableName": "DB_ORACLE_BASE_LOC",
"variableValue": "/scratch/cuser/19c/racbase"
},
{
"variableName": "DB_ORACLE_HOME_LOC",
"variableValue": "/scratch/cuser/19c/rachome"
},
{
"variableName": "DB_PRODUCT_VERSION",
"variableValue": "19.0.0.0.0"
},
{
"variableName": "DEPLOY_MODE",
"variableValue": "DEPLOY_GI"
},
{
"variableName": "DISKGROUP_AU_COMPATIBLE",
"variableValue": "DATA:4:::;MGMTDB:4:::"
},
{
"variableName": "DISK_GROUP_NAMES",
"variableValue": "DATA,MGMTDB"
},
{
"variableName": "FLEX_ENABLED",
"variableValue": "false"
},
},
{
```

```
"variableName": "GIMR_DB_ORACLE_BASE_LOC",
"variableValue": "/scratch/cuser/19c/racbase"
},
{
"variableName": "GIMR_DB_ORACLE_HOME_LOC",
"variableValue": "/scratch/cuser/19c/rachome"
},
{
"variableName": "GIMR_DISK_GROUP",
"variableValue": "MGMTDB"
},
{
"variableName": "GIMR_TYPE",
"variableValue": "NOGIMR"
},
{
"variableName": "GI_AU_SIZE",
"variableValue": "4"
},
{
"variableName": "GI_CLUSTER_TYPE",
"variableValue": "1"
},
{
"variableName": "GI_COMPONENT",
"variableValue": "Fleet/OH19/linux64/DB19RU - OH19RU"
},
{
"variableName": "GI_HOME_SHARED",
"variableValue": "false"
},
{
"variableName": "GI_HOST_NORMAL_CREDNAMES",
"variableValue": "CUSER:SYSMAN"
},
{
"variableName": "GI_HOST_ROOT_CREDNAMES",
"variableValue": "CUSER_ROOT:SYSMAN"
},
{
"variableName": "GI_ORACLE_BASE_LOC",
"variableValue": "/scratch/gridbase"
},
{
"variableName": "GI_ORACLE_HOME_LOC",
"variableValue": "/scratch/gridhome"
},
{
"variableName": "GI_PRODUCT_VERSION",
"variableValue": "19.0.0.0.0"
},
{
"variableName": "GI_STORAGE_TYPE",
"variableValue": "1"
},
{
```

```
"variableName": "GNS_ENABLED",
"variableValue": "true"
},
{
"variableName": "GNS_VIP_ADDRESS",
"variableValue": "emracp1011-gns.subnet1rg.emsubnet.companyvcn.com"
},
{
"variableName": "GRID_DISK_FLAG",
"variableValue": "true"
},
{
"variableName": "GROUP_VERSION",
"variableValue": "19.0.0.0.0"
},
{
"variableName": "INTERFACE_PARAM_VALUE",
"variableValue": "ens3:100.75.0.0:1,ens4:192.168.1.0:5"
},
{
"variableName": "NOT_RELINK_DB_HOME_SELECTED",
"variableValue": "false"
},
{
"variableName": "NOT_RELINK_GIMR_DB_HOME_SELECTED",
"variableValue": "false"
},
{
"variableName": "NOT_RELINK_GI_HOME_SELECTED",
"variableValue": "false"
},
{
"variableName": "OCR_VDSK_DISK_GROUP",
"variableValue": "DATA"
},
{
"variableName": "OINSTALL_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSASM_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSBACKUPDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSDGDBA_GROUP",
"variableValue": "oinstall"
},
{

```

```
"variableName": "OSKMDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSOPER_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSRACDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "PAUSE_AFTER_PREREQ",
"variableValue": "false"
},
{
"variableName": "PLATFORM_ARUID",
"variableValue": "226"
},
{
"variableName": "PREREQ_CHECK_ERRORMODE",
"variableValue": "continue"
},
{
"variableName": "RAC_DBCA_MODE",
"variableValue": "CreateRACDatabase"
},
{
"variableName": "RAC_HOME_SHARED",
"variableValue": "false"
},
{
"variableName": "RAC_STORAGE_TYPE",
"variableValue": "1"
},
{
"variableName": "REDUNDANCYS",
"variableValue": "NORMAL,NORMAL"
},
{
"variableName": "ROLE_NODENAMES",
"variableValue": "HUB,HUB"
},
{
"variableName": "RUN_CVU_FROM_OH",
"variableValue": "true"
},
{
"variableName": "SCAN_NAME",
"variableValue": "emracn0607-r.racdb-scan.emsubnet.companyvcn.com"
},
{
"variableName": "SCAN_PORT",
"variableValue": "1521"
},
{
```

```
"variableName": "SKIP_CVU_CHECK",
"variableValue": "false"
},
{
"variableName": "SKIP_PREREQUISITE_CHECKS",
"variableValue": "false"
},
{
"variableName": "SOURCE_TYPE",
"variableValue": "SOFTWARE_LIBRARY"
},
{
"variableName": "SSH_SETUP",
"variableValue": "false"
},
{
"variableName": "TARGET_HOST_LIST",
"variableValue":
"emracn07.subnet1rg.emsubnet.companyvcn.com,emracn06.subnet1rg2.emsubnet.co
mpanyvcn.com"
},
{
"variableName": "VERSION_GREATER_12102",
"variableValue": "true"
},
{
"variableName": "VIP_NODENAMES",
"variableValue": "emracn07-vip.subnet1rg.emsubnet.companyvcn.com,emracn06-
vip.subnet1rg.emsubnet.companyvcn.com"
},
{
"variableName": "WORK_DIR_LOC",
"variableValue": "/tmp/"
},
{
"variableName": "ASM_SYS_PASSWORD",
"variableValue": "oracle"
},
{
"variableName": "ASM_MONITOR_PASSWORD",
"variableValue": "oracle"
},
{
"variableName": "FPP_CONFIGURATION_TYPE",
"variableValue": "client"
},
{
"variableName": "FPP_SERVER_TARGET_NAME",
"variableValue": "phxemracn0607"
},
{
"variableName": "FPP_SERVER_HOST_USER",
"variableValue": "CUSER:SYSMAN"
},
{
"variableName": "PREREQS_ONLY",
```

```
"variableValue": "false"
}
]
}
}
```

- **Provision an FPP Server on Existing Grid Infrastructure
EM CLI**

```
emcli submit_procedure -input_file=data:"/home/user_name/FPP/
input.properties" -name="RACPROV2"
```

The input file must contain the following

```
TARGET_NAME=emracn0607
FPP_DISK_GROUP=MGMTDB
FPP_IMAGES_MOUNTPOINT=/scratch/cuser/fppImages
GNS_VIP_ADDRESS=emracn0607-gns.subnet1rg.emsubnet.companyvcn.com
NORMAL_USER=CUSER:SYSMAN
ROOT_USER=CUSER_ROOT:SYSMAN
MIN_HEAP_SIZE=3G
MAX_HEAP_SIZE=5G
REST_USER_NAME=restuser
REST_USER_PASSWORD=<password>
REST_USER_EMAIL=user@example.com
```

REST API:

Request Method

POST

URI

```
https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT/em/api/deploymentProcedures/
<IMAGE_ID>/actions/submit
```

Request Headers

```
Authorization: Basic
Content-Type:
  application/json
```

Payload

```
{
  "targetList": [
  ],
  "variables": {
    "items": [
      {
        "variableName": "TARGET_NAME",
```

```

        "variableValue": "emracn0607"
    },
    {
        "variableName": "FPP_DISK_GROUP",
        "variableValue": "MGMTDB"
    },
    {
        "variableName": "FPP_IMAGES_MOUNTPOINT",
        "variableValue": "/scratch/cuser/fppImages"
    },
    {
        "variableName": "NORMAL_USER",
        "variableValue": "CUSER:SYSMAN"
    },
    {
        "variableName": "ROOT_USER",
        "variableValue": "CUSER_ROOT:SYSMAN"
    },
    {
        "variableName": "REST_USER_NAME",
        "variableValue": "restuser"
    },
    {
        "variableName": "REST_USER_PASSWORD",
        "variableValue": "userdefined"
    },
    {
        "variableName": "REST_USER_EMAIL",
        "variableValue": "user@example.com"
    }
    ]
}
}
}

```

- **Provision an FPP Client and Grid Infrastructure EM CLI**

```
emcli submit_procedure -input_file=data:"/home/user_name/FPP/
gi_client_input.properties"-name="RAC_PROV"
```

The input file must contain the following

```

ASMDBA_GROUP=oinstall
ASMOPER_GROUP=oinstall
ASM_DISK_LISTS=/mnt/emracdev/emracp1011/ASMDISK/disk1,/mnt/emracdev/
emracp1011/ASMDISK/disk2,/mnt/emracdev/emracp1011/ASMDISK/disk3
ASM_DISK_STRINGS=/mnt/emracdev/emracp1011/ASMDISK/disk*
ASM_FAILURE_AND_QUORUM_GROUPS=/mnt/emracdev/emracp1011/ASMDISK/disk1||
false,/mnt/emracdev/emracp1011/ASMDISK/disk2||false,/mnt/emracdev/
emracp1011/ASMDISK/disk3||false
ASM_OCRVDSK_FAILURE_AND_QUORUM_GROUPS=/mnt/emracdev/emracp1011/ASMDISK/
disk1||false,/mnt/emracdev/emracp1011/ASMDISK/disk2||false,/mnt/emracdev/
emracp1011/ASMDISK/disk3||false
ASM_SYS_PASSWORD=sample1
AUTOVIP_ENABLED=false

```



```
CLUSTER_NAME=emracp1011
CONFIGURE_ASMFD=false
CONFIGURE_GIMR=false
CONFIGURE_MODE=CONFIGURE_ASM
DB_ORACLE_BASE_LOC=/u02/app/oracle
DB_ORACLE_HOME_LOC=/u02/app/oracle/app/product/19.0.0.0.0/db
DB_PRODUCT_VERSION=19.0.0.0.0
DEPLOY_MODE=DEPLOY_GI
DISKGROUP_AU_COMPATIBLE=DATA:4:::
DISK_GROUP_NAMES=DATA
FLEX_ENABLED=false
GIMR_DB_ORACLE_BASE_LOC=/u03/app/oracle
GIMR_DB_ORACLE_HOME_LOC=/u03/app/oracle/app/product/21.0.0.0/gimrdb
GIMR_DISK_GROUP=DATA
GIMR_TYPE=NOGIMR
GI_AU_SIZE=4
GI_CLUSTER_TYPE=1
GI_COMPONENT=Fleet/OH19/linux64/DB19 - RU19
GI_HOME_SHARED=false
GI_HOST_NORMAL_CREDNAMES=CUSER:SYSMAN
GI_HOST_ROOT_CREDNAMES=CUSER_ROOT:SYSMAN
GI_ORACLE_BASE_LOC=/scratch/gridbase
GI_ORACLE_HOME_LOC=/scratch/gridhome
GI_PRODUCT_VERSION=19.0.0.0.0
GI_STORAGE_TYPE=1
GNS_ENABLED=false
GRID_DISK_FLAG=true
GROUP_VERSION=19.0.0.0.0
INTERFACE_PARAM_VALUE=ens3:100.75.0.0:1,ens4:192.168.1.0:5
NOT_RELINK_DB_HOME_SELECTED=false
NOT_RELINK_GIMR_DB_HOME_SELECTED=false
NOT_RELINK_GI_HOME_SELECTED=false
OCR_VDSK_DISK_GROUP=DATA
OINSTALL_GROUP=oinstall
OSASM_GROUP=oinstall
PAUSE_AFTER_PREREQ=false
PLATFORM_ARUID=226
PREREQ_CHECK_ERRORMODE=continue
RAC_DBCA_MODE=CreateRACDatabase
RAC_HOME_SHARED=false
RAC_STORAGE_TYPE=1
REDUNDANCYS=NORMAL
ROLE_NODENAMES=HUB,HUB
RUN_CVU_FROM_OH=true
SCAN_NAME=emracp1011-r.racdb-scan.emsubnet.companyvcn.com
SCAN_PORT=1521
SKIP_CVU_CHECK=true
SKIP_PREREQUISITE_CHECKS=false
SOURCE_TYPE=SOFTWARE_LIBRARY
SSH_SETUP=false
TARGET_HOST_LIST=emracp10.subnet1rg2.emsubnet.companyvcn.com,emracp11.subne
t1rg2.emsubnet.example.com
VERSION_GREATER_12102=true
VIP_NODENAMES=emracp10-vip.subnet1rg2.emsubnet.companyvcn.com,emracp11-
vip.subnet1rg2.emsubnet.companyvcn.com
FPP_SERVER_TARGET_NAME=emracn0607
```

```
FPP_SERVER_HOST_USER=CUSER:SYSMAN
FPP_CONFIGURATION_TYPE=client
#PREREQS_ONLY=true
WORK_DIR_LOC=/tmp/
```

REST API:

Request Method

POST

URI

```
https://<OMS_CONSOLE_HOST>:<OMS_CONSOLE_PORT>/em/api/deploymentProcedures/
<IMAGE_ID>/actions/submit
```

Request Headers

```
Authorization: Basic
Content-Type:
    application/json
```

Payload

```
},
{
  "variableName": "GIMR_DB_ORACLE_BASE_LOC",
  "variableValue": "/scratch/cuser/19c/racbase"
},
{
  "variableName": "GIMR_DB_ORACLE_HOME_LOC",
  "variableValue": "/scratch/cuser/19c/rachome"
},
{
  "variableName": "GIMR_DISK_GROUP",
  "variableValue": "MGMTDB"
},
{
  "variableName": "GIMR_TYPE",
  "variableValue": "NOGIMR"
},
{
  "variableName": "GI_AU_SIZE",
  "variableValue": "4"
},
{
  "variableName": "GI_CLUSTER_TYPE",
  "variableValue": "1"
},
{
  "variableName": "GI_COMPONENT",
  "variableValue": "Fleet/OH19/linux64/DB19RU - OH19RU"
},
{
```

```
"variableName": "GI_HOME_SHARED",
"variableValue": "false"
},
{
"variableName": "GI_HOST_NORMAL_CREDNAMES",
"variableValue": "CUSER:SYSMAN"
},
{
"variableName": "GI_HOST_ROOT_CREDNAMES",
"variableValue": "CUSER_ROOT:SYSMAN"
},
{
"variableName": "GI_ORACLE_BASE_LOC",
"variableValue": "/scratch/gridbase"
},
{
"variableName": "GI_ORACLE_HOME_LOC",
"variableValue": "/scratch/gridhome"
},
{
"variableName": "GI_PRODUCT_VERSION",
"variableValue": "19.0.0.0.0"
},
{
"variableName": "GI_STORAGE_TYPE",
"variableValue": "1"
},
{
"variableName": "GNS_ENABLED",
"variableValue": "true"
},
{
"variableName": "GNS_VIP_ADDRESS",
"variableValue": "emracp1011-gns.subnet1rg2.examplevcn.com"
},
{
"variableName": "GRID_DISK_FLAG",
"variableValue": "true"
},
{
"variableName": "GROUP_VERSION",
"variableValue": "19.0.0.0.0"
},
{
"variableName": "INTERFACE_PARAM_VALUE",
"variableValue": "ens3:100.75.0.0:1,ens4:192.168.1.0:5"
},
{
"variableName": "NOT_RELINK_DB_HOME_SELECTED",
"variableValue": "false"
},
{
"variableName": "NOT_RELINK_GIMR_DB_HOME_SELECTED",
"variableValue": "false"
},
{
```

```
"variableName": "NOT_RELINK_GI_HOME_SELECTED",
"variableValue": "false"
},
{
"variableName": "OCR_VDSK_DISK_GROUP",
"variableValue": "DATA"
},
{
"variableName": "OINSTALL_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSASM_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSBACKUPDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSDGDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSKMDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSOPER_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "OSRACDBA_GROUP",
"variableValue": "oinstall"
},
{
"variableName": "PAUSE_AFTER_PREREQ",
"variableValue": "false"
},
{
"variableName": "PLATFORM_ARUID",
"variableValue": "226"
},
{
"variableName": "PREREQ_CHECK_ERRORMODE",
"variableValue": "continue"
},
{
"variableName": "RAC_DBCA_MODE",
"variableValue": "CreateRACDatabase"
},
{
```

```
"variableName": "RAC_HOME_SHARED",
"variableValue": "false"
},
{
"variableName": "RAC_STORAGE_TYPE",
"variableValue": "1"
},
{
"variableName": "REDUNDANCYS",
"variableValue": "NORMAL,NORMAL"
},
{
"variableName": "ROLE_NODENAMES",
"variableValue": "HUB,HUB"
},
{
"variableName": "RUN_CVU_FROM_OH",
"variableValue": "true"
},
{
"variableName": "SCAN_NAME",
"variableValue": "emracp1011-r.racdb-scan.emsubnet.companyvcn.com"
},
{
"variableName": "SCAN_PORT",
"variableValue": "1521"
},
{
"variableName": "SKIP_CVU_CHECK",
"variableValue": "false"
},
{
"variableName": "SKIP_PREREQUISITE_CHECKS",
"variableValue": "false"
},
{
"variableName": "SOURCE_TYPE",
"variableValue": "SOFTWARE_LIBRARY"
},
{
"variableName": "SSH_SETUP",
"variableValue": "false"
},
{
"variableName": "TARGET_HOST_LIST",
"variableValue":
"emracp10.subnet1rg2.emsubnet.companyvcn.com,emracp11.subnet1rg2.emsubnet.e
xample.com"
},
{
"variableName": "VERSION_GREATER_12102",
"variableValue": "true"
},
{
"variableName": "VIP_NODENAMES",
"variableValue": "emracp10-vip.subnet1rg2.emsubnet.companyvcn.com,emracp11-
```

```

vip.subnet1rg2.emsubnet.companyvcn.com"
},
{
"variableName": "WORK_DIR_LOC",
"variableValue": "/tmp/"
},
{
"variableName": "ASM_SYS_PASSWORD",
"variableValue": "oracle"
},
{
"variableName": "ASM_MONITOR_PASSWORD",
"variableValue": "oracle"
},
{
"variableName": "FPP_CONFIGURATION_TYPE",
"variableValue": "client"
},
{
"variableName": "FPP_SERVER_TARGET_NAME",
"variableValue": "phxemracn0607"
},
{
"variableName": "FPP_SERVER_HOST_USER",
"variableValue": "CUSER:SYSMAN"
},
{
"variableName": "PREREQS_ONLY",
"variableValue": "false"
}
]
}
}

```

- **Provision an FPP Client on Existing Grid Infrastructure**
EMCLI:

```

emcli submit_procedure -input_file=data:"/home/user_name/FPP/
input.properties" -name="SETUP_FPP_CLIENT"

```

The input file must contain the following

```

TARGET_NAME=emracp1011
NORMAL_USER=CUSER:SYSMAN
ROOT_USER=CUSER_ROOT:SYSMAN
FPP_SERVER_TARGET_NAME=emracn0607
FPP_SERVER_HOST_USER=CUSER:SYSMAN
WORKING_DIRECTORY=/scratch/cuser/tmp

```

REST API:

Request Method

POST

URI

```
https://<OMS_CONSOLE_HOST>:OMS_CONSOLE_PORT>/em/api/deploymentProcedures/  
<IMAGE_ID>/actions/submit
```

Request Headers

```
Authorization: Basic  
Content-Type:  
    application/json
```

Payload

```
{  
  "targetList": [  
  
  ],  
  "variables": {  
    "items": [  
      {  
        "variableName": "TARGET_NAME",  
        "variableValue": "emracp1011"  
      },  
      {  
        "variableName": "NORMAL_USER",  
        "variableValue": "CUSER:SYSMAN"  
      },  
      {  
        "variableName": "ROOT_USER",  
        "variableValue": "CUSER_ROOT:SYSMAN"  
      },  
      {  
        "variableName": "FPP_SERVER_TARGET_NAME",  
        "variableValue": "emracn0607"  
      },  
      {  
        "variableName": "FPP_SERVER_HOST_USER",  
        "variableValue": "CUSER:SYSMAN"  
      }  
    ]  
  }  
}
```

Register FPP Server and Client

Once the FPP server and client have been provisioned, the next step is to register them to Oracle Enterprise Manager. These steps are for existing FPP servers and clients that need to be registered in Enterprise Manager.



Note:

Registering the server is not required when provisioning when provisioning via automation using the steps outlined in [Provision an FPP Server and Client](#)

Register FPP server and FPP Client with Enterprise Manager

Discover the FPP server and client, running on Grid Infrastructure, as cluster targets in Enterprise Manager. If the Grid Infrastructure where the FPP Server or FPP Client are running are already monitored by Enterprise Manager, no additional discovery of targets is required. You can discover targets in EM by configuring autodiscovery, or by using either a guided or a declarative discovery process. For more details, see [Discovering and Adding Host Targets](#).

Register an FPP Server Using EM CLI

To register an FPP server use EM CLI command

```
emcli db_software_maintenance
-setup -registerFPPServer
-input_file="data:<fileLocation>"
```

with the input file containing:

```
{
  "name": "myCluster0203",
  "properties": {
    "restUrl": "https://fppserver.example.com:8894/rhp-restapi/rhp",
    "agentCredential": "USER:SYSMAN"
  },
  "restUser": {
    "username": "restUser",
    "password": "restPassword"
  },
  "schedule": {
    ["timeToStart": "2024-09-01T07:10:50.266Z",]
    ["timeToEnd": "2024-09-01T07:10:50.266Z",]
    ["isTargetZone": true,]
    ["scheduleType": "string",]
    ["recurring": "Yearly/Monlthy/Weekly/Daily",]
    ["startGracePeriodInMinutes": 60]
  }
}
```


**Note:**

[] Denotes an optional parameter

Register an FPP Client Using EM CLI

To register an FPP client use EM CLI command use command

```
emcli db_software_maintenance
-setup -registerFPPClient
-input_file="data:/<fileLocation>"
```

with the following input file:

```
{
  "serverName": "<ServerName>",
  "clientName": "<Client Name>"
}
```

Image Maintenance

Exadata Infrastructure Patching using Fleet Maintenance is a gold-image based patching solution.

This Image Maintenance chapter guides you on the process of opting in an image (optional), create a Gold Image, and finally update a Gold Image.

- [Opt in an Image](#)
- [Create a Gold Image](#)
- [Mark a Gold Image as Current](#)

Opt in an Image

Fleet Maintenance offers the flexibility to leverage an existing image stored in Fleet Provisioning and Patching (FPP) for patching all supported Exadata components.

**Note:**

Opt in is an optional process, you can skip the opt in if you fulfill one of the following conditions:

- You don't plan to utilize an existing FPP image
- You intend to create/update a new image

To opt in your FPP patching server follow these steps:

1. Review available FPP Server series, by querying the FPP server and listing the image series. Use the following command:

```
emcli db_software_maintenance -setup -getFPPSeries
-input_file="data:/scratch/<user_name>/fpp/payload.properties"
```

Where the input file contains the following:

```
{
"targetName": "<host, oracle_exadata, or oracle_si_netswitch target
name>",
"targetType": "<host, oracle_exadata, or oracle_si_netswitch>"
}
```

2. List all available FPP images. Use the following command:

```
emcli db_software_maintenance -setup -getFPPImages
-input_file="data:/scratch/<user_name>/fpp/payload.properties"
```

Where the input file contains the following:

```
{
"seriesId": "<Series ID>",
"seriesName": "<Series Name>",
"serverName": "<Server Name>"
}
```

3. Opt in the FPP series using the command outlined below. In this step, the image series information is stored in Enterprise Manager.

```
emcli db_software_maintenance -setup -optInFPPSeries -
input_file="data:<fileLocation>"
```

Where the input file contains the following:

```
{
"seriesId": "<Image series ID>",
"currentVersionImageId": "<Image ID>"
}
```

 **Note:**

"currentVersionImageId": "<Image ID>" is an optional parameter, if it is not specified the latest image created in the series is marked as *current*.

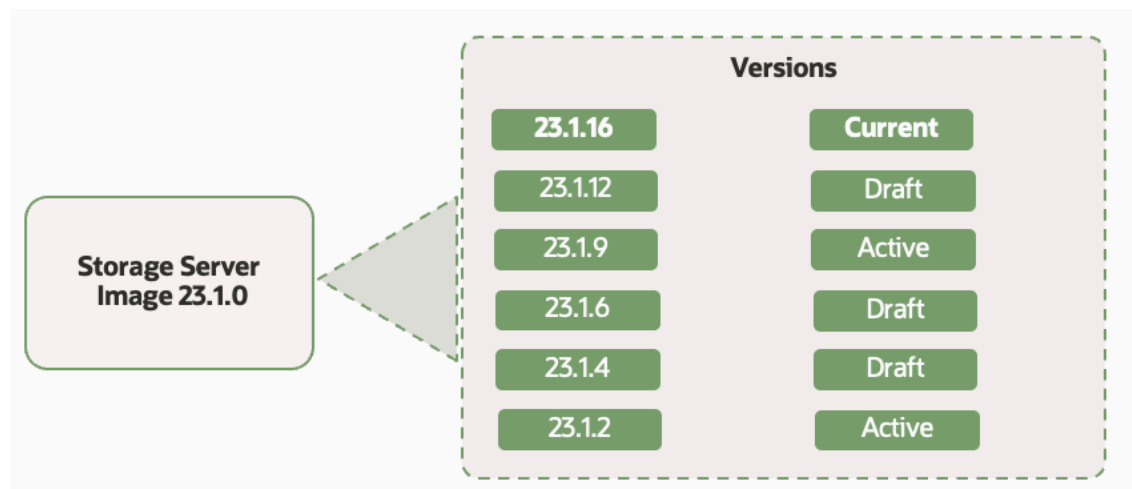
For REST API Opt in see: [Database Maintenance- Using Oracle Fleet Patching And Provisioning \(FPP\) REST Endpoints](#)

Create a Gold Image

A Gold Image represents a specific Exadata release for a particular target type. It can include multiple versions aligned with corresponding Exadata software release. Each time a new Exadata version is released, a new Gold Image version is created.

Note:

To maintain efficiency and storage optimization, we recommend limiting each Gold Image to three versions.



Note:

For FPP patching the Gold Image will reside in the FPP server and not in Enterprise Manager. Keep this in mind for FPP server storage space considerations.

Follow these steps to create and FPP Gold Image:

Download Patches

Review and download Exadata patches using MOS Note: Exadata Database Machine and Exadata Storage Server Supported Versions [888828.1](#)

Create a Gold Image

Run the following EM CLI command to create the Gold Image: `emcli db_software_maintenance -createSoftwareImage -input_file="data:<location>/fpp_image".`

The input file will vary depending on the Exadata resource type to be patched:

- Storage Server:

```
IMAGE_NAME=EMCLI_INSTALL_FEST_STORAGE_SERIES
VERSION_NAME=EMCLI_INSTALL_FEST_STORAGE_IMAGE_v01
IMAGE_VERSION=22.1.3.0.0
```

```
FPP_SERVER_HOST=emracr06.subnet.example.com
IMAGE_LOCATION=/mnt/emrac/patches/patches_22.1.3.0.0/
storage_22.1.3.0.0.220914
WORKING_DIRECTORY=/tmp
REF_HOST_CREDENTIALS=FPP_SERVER_CUSER:SYSMAN
IMAGE_DESCRIPTION=Image description for metadata
TARGET_TYPE=oracle_exadata
STORAGE_FILE=p34568888_221300_Linux-x86-64.zip
```

- **Switch:**

```
IMAGE_NAME=Exadata_switch_Series_2317_1
VERSION_NAME=Exadata_switch_Series_2317_1
IMAGE_VERSION=22.1.3.0.0.221207
FPP_SERVER_HOST=emracr06.subnet.example.com
IMAGE_LOCATION=/scratch/switch_231700
WORKING_DIRECTORY=/tmp
REF_HOST_CREDENTIALS=FPP_SERVER_CUSER:SYSMAN
IMAGE_DESCRIPTION=Image description for metadata
TARGET_TYPE=oracle_si_netswitch
SWITCH_FILE=p36261780_2311200_Linux-x86-64.zip
```

- **Database Node:**

```
IMAGE_NAME=Exadata_dbnode_Series_2317_1
VERSION_NAME=Exadata_dbnode_Series_2317_1
IMAGE_VERSION=22.1.3.0.0.221207
FPP_SERVER_HOST=emracr06.subnet.example.com
IMAGE_LOCATION=/scratch/dbnode_231700
WORKING_DIRECTORY=/tmp
REF_HOST_CREDENTIALS=FPP_SERVER_CUSER:SYSMAN
IMAGE_DESCRIPTION=Image description for metadata
TARGET_TYPE=host
DBNODE_FILE=p36261778_2311200_Linux-x86-64.zip
PATCH_MANAGER_LOCATION=dbserver_patch_220810
```

 **Note:**

- IMAGE_VERSION must contain the full version name including patch date. For example 22.1.3.0.0.221207.
- Patch manager location file must be unzipped.

Create an Image Version

In a recurring patch cycle you can create a new image version of an existing Gold Image, using the image ID you can create an up to date image version that will have the latest patches without having to create a new Gold Image at every cycle.

Run the following EM CLI command to create the Gold Image: `emcli`

```
db_software_maintenance -createSoftwareImage -input_file="data:/<location>/
fpp_image".
```

The input file will vary depending on the Exadata resource type to be patched:

- **Storage Server:**

```
FPP_IMAGE_ID=1C3BAEBE500B3F67E063DF034B64620
VERSION_NAME=EMCLI_INSTALL_FEST_STORAGE_IMAGE_v01
IMAGE_VERSION=22.1.3.0.0
FPP_SERVER_HOST=emracr06.subnet.example.com
IMAGE_LOCATION=/mnt/emrac/patches/patches_22.1.3.0.0/
storage_22.1.3.0.0.220914
WORKING_DIRECTORY=/tmp
REF_HOST_CREDENTIALS=FPP_SERVER_CUSER:SYSMAN
IMAGE_DESCRIPTION=Image description for metadata
TARGET_TYPE=oracle_exadata
STORAGE_FILE=p34568888_221300_Linux-x86-64.zip
```

- **Switch:**

```
FPP_IMAGE_ID=1C3BAEBE500B3F67E063DF034B64620
VERSION_NAME=Exadata_switch_Series_2317_1
IMAGE_VERSION=22.1.3.0.0.221207
FPP_SERVER_HOST=emracr06.subnet.example.com
IMAGE_LOCATION=/scratch/switch_231700
WORKING_DIRECTORY=/tmp
REF_HOST_CREDENTIALS=FPP_SERVER_CUSER:SYSMAN
IMAGE_DESCRIPTION=Image description for metadata
TARGET_TYPE=oracle_si_networkswitch
SWITCH_FILE=p36261780_2311200_Linux-x86-64.zip
```

- **Database Node:**

```
FPP_IMAGE_ID=1C3BAEBE500B3F67E063DF034B64620
VERSION_NAME=Exadata_dbnode_Series_2317_1
IMAGE_VERSION=22.1.3.0.0.221207
FPP_SERVER_HOST=emracr06.subnet.example.com
IMAGE_LOCATION=/scratch/dbnode_231700
WORKING_DIRECTORY=/tmp
REF_HOST_CREDENTIALS=FPP_SERVER_CUSER:SYSMAN
IMAGE_DESCRIPTION=Image description for metadata
TARGET_TYPE=host
DBNODE_FILE=p36261778_2311200_Linux-x86-64.zip
PATCH_MANAGER_LOCATION=dbserver_patch_220810
```

 **Note:**

- **IMAGE_VERSION** must contain the full version name including patch date. For example 22.1.3.0.0.221207.
- Patch manager location file must be unzipped.

Mark a Gold Image as Current

When an updated version of the image has been created, this new version must be marked as Current. This indicates that the upcoming patching operation will apply patches based on this version.

The mark a Gold Image cycle can be broken down in two parts

- Initial patching cycle:
 1. [Create a Gold Image](#) that serves as the baseline for patching
 2. Subscribe the targets to the initially created Gold Image
 3. Patch the targets with the Gold Image
- Recurring patching cycle:
 1. [Create a new version](#) of the existing the Gold Image
 2. Mark the latest version as *Current*
 3. Patch the targets with the latest version (Gold Image)

To set the Gold Image to current using EM CLI use the following commands based on Exadata type:



Note:

Patch zip file must be extracted before running the command.

Storage Server:

```
emcli db_software_maintenance
-updateVersionStatus
-version_id="166E5A0FA92D6D46E063DF034B64E21E"
-status="CURRENT"
-image_id="15314B07900969F6E063DF034B64F40C"
-targetType="oracle_exadata"
-imageFile="/home/<user name>/FPP/p34568888_221300_Linux-x86-64";
```

Switch

```
emcli db_software_maintenance
-updateVersionStatus
-version_id="166E5A0FA92D6D46E063DF034B64E21E"
-status="CURRENT"
-image_id="15314B07900969F6E063DF034B64F40C"
-targetType="oracle_si_netswitch"
-imageFile="/home/<user name>/FPP/p34568888_221300_Linux-x86-64";
```

Database Node:

```
emcli db_software_maintenance
-updateVersionStatus
-version_id="166E5A0FA92D6D46E063DF034B64E21E"
```

```
-status="CURRENT"  
-image_id="15314B07900969F6E063DF034B64F40C"  
-targetType="host"  
-imageFile="/home/<user name>/FPP/p34568888_221300_Linux-x86-64"  
-patchManagerLocation="dbserver_patch_220810";
```

Subscribe Exadata Targets

To achieve standardization, Exadata targets must be subscribed to the image. An Exadata target should be subscribed to only one gold image at a given time.

To subscribe Exadata targets to the updated (latest) Gold Image run the following EM CLI commands based on Exadata type:

Storage Server

```
emcli db_software_maintenance  
-subscribeTarget  
-target_type="oracle_exadata"  
-target_list="cell1.subdom.example.com, cell2.subdom.example.com"  
-image_id="148AFC66D2523DDDE063DF034B64E36D"
```

Switch

```
emcli db_software_maintenance  
-subscribeTarget  
-target_type="oracle_si_netswitch"  
-target_list="switch4.subdom.example.com, switch05.subdom.example.com"  
-image_id="148AFC66D2523DDDE063DF034B64E36D"
```

Database Node

```
emcli db_software_maintenance  
-subscribeTarget  
-target_type="host"  
-target_list="exadatabase4.subdom.example.com, exadatabse5.subdom.example.com"  
-image_id="148AFC66D2523DDDE063DF034B64E36D"
```

Once the targets have been subscribed, you can now proceed to patching.

Exadata Fleet Maintenance Patching Operations

Having completed the essential infrastructure setup (configuring and registering an FPP server and client within EM), creating a gold image, and subscribing targets to the image, you're now ready to deploy the image, update the targets and perform an optional validation step before the update.

Patching Operations

The following is a list of operations required to successfully patch an Exadata target. First you deploy, then perform an optional validation, and finally perform an update. The following links are the EM CLI samples, while the entire UI flow can be executed in a single operational flow and is outlined in full here: [Patch Exadata Infrastructure with the Fleet Maintenance UI](#).

- [Deploy a Patch](#)

- [Update a Patch](#)

For REST API Exadata Patching operations see: [Database Maintenance- Using Oracle Fleet Patching And Provisioning \(FPP\) REST Endpoints](#).

Patch Exadata Infrastructure with the Fleet Maintenance UI

1. In Enterprise Manager click on **Targets** and select **Exadata**. In the Engineered Systems page click on the **Actions** button and select **Fleet Maintenance**.
2. In the Targets page select the desired FPP image to which the Exadata component are subscribed to. You can filter by **Image** to select the Gold Image to be used. There are three supported components Exadata storage cells, network switches, and hosts. Once you select an image, **Target Type** will be automatically populated with the corresponding type (host, switch, Exadata cell). Verify that Update appears under **Operation**.
3. Select all Exadata infrastructure targets requiring a patch deployment and/or update and click **Next**.
4. In the Tasks page **Software Deployment** will be automatically selected, enter the **Deployment Path**, where the patch will be staged, and select a **Start** time. If you wish to both deploy and update at the same time, ensure that **Update Exadata** task is selected with a check mark, and select a **Start** time.

Note:

The scheduled time for the update can be either Immediately/Auto or Later. If you select **Immediately** the patching task will begin immediately after the deployment task is successfully completed. When selecting Later make sure there is ample time between operations in case of issues.

In the **Credentials** section, specify the FPP Rest User credential to be used.

Under **Options**, enter the **Work Directory** and add any additional FPP options for the deployment and update operations.

Once all fields have been entered and verified click **Next**.

5. In the Review page, verify all information entered is correct, and click on **Validate**. Select **Quick** to perform a quick validation or **Full** to perform a full validation (full validations submit a deployment procedure). Once the validation completes successfully, submit the operation by clicking **Submit**. Enter a name for the procedure and click **Submit** once again.
6. Click on **Monitor Progress** to view the operation in the Procedure Activity Page. Once the procedure completes successfully, you have updated your Exadata infrastructure.

Deploy a Patch

Deploy operations are performed to simplify Exadata patch deployments. They can be performed in a short period of time because the Gold Image (the version of the image that is marked as current) is copied to a parallel Exadata type which is then turned into a newly patched Exadata type.

To deploy a patch use the following EM CLI command based on the Exadata type you wish to deploy a patch on:

**Note:**

[] Denotes an optional parameter.

Storage Server

```
emcli db_software_maintenance
-performOperation
-name="Deploy Exadata Storage Cell node"
-purpose=DEPLOY_EXADATA_SOFTWARE
-target_type=oracle_exadata -target_list="cell1,cell2"
-input_file="data:/scratch/user_name/fpp/deployExadata.prop"
[-start_schedule="2024/04/20 06:10"]
```

Switch

```
emcli db_software_maintenance
-performOperation
-name="Deploy Exadata IB Switch node"
-purpose=DEPLOY_EXADATA_SOFTWARE
-target_type=oracle_si_netswitch
-target_list="Switch1,Switch2"
-input_file="data:/scratch/user_name/fpp/deployExadata.prop"
[-start_schedule="2024/04/20 06:10"]
```

Database Node

```
emcli db_software_maintenance
-performOperation
-name="Deploy Exadata DB node"
-purpose=DEPLOY_EXADATA_SOFTWARE
-target_type=host
-target_list="DB1"
-input_file="data:/scratch/user_name/fpp/deployExadata.prop"
[-start_schedule="2024/04/20 06:10"]
```

Where the input file will contain the following for all Exadata types:

```
WORKING_DIRECTORY=/tmp
PATCH_EXTRACT_LOCATION=/<path>
```

Update a Patch**Validate a Patch**

It is recommended that before you perform a patch update, you first perform a validation. Use the following EM CLI commands to perform a validation based on Exadata type patch:

Storage Server

```
emcli db_software_maintenance
-performOperation
```

```
-name="Update Storage Cell node"  
-purpose="UPDATE_EXADATA_SOFTWARE"  
-target_list="Cell,Cell2"  
-target_type="oracle_exadata"  
-input_file="data:/user_name/fppinstallfest/updateExadata.prop"
```

Switch

```
emcli db_software_maintenance  
-performOperation  
-name="Update IB Switch node"  
-purpose="UPDATE_EXADATA_SOFTWARE"  
-target_list="IB1,IB2"  
-target_type="oracle_si_netswitch"  
-input_file="data:/user_name/fppinstallfest/updateExadata.prop"
```

Database Node

```
emcli db_software_maintenance  
-performOperation  
-name="Update DB node"  
-purpose="UPDATE_EXADATA_SOFTWARE"  
-target_list="exadatabase2.subdom.example.com"  
-target_type="host"  
-input_file="data:/user_name/fppinstallfest/updateExadata.prop"
```

Where the input file will contain the following:

```
WORKING_DIRECTORY=/tmp  
fppOptions={"eval": "true"}
```

Update a Patch

To perform a patch update use the following EM CLI command:
Image Cell

```
emcli db_software_maintenance  
-performOperation  
-name="Update Storage Cell node"  
-purpose="UPDATE_EXADATA_SOFTWARE"  
-target_list="Cell,Cell2"  
-target_type="oracle_exadata"  
-input_file="data:/user_name/fppinstallfest/updateExadata.prop"
```

Switch

```
emcli db_software_maintenance  
-performOperation  
-name="Update IB Switch node"  
-purpose="UPDATE_EXADATA_SOFTWARE"  
-target_list="IB1,IB2"  
-target_type="oracle_si_netswitch"  
-input_file="data:/user_name/fppinstallfest/updateExadata.prop"
```

Database Node

```
emcli db_software_maintenance
-performOperation
-name="Update DB node"
-purpose="UPDATE_EXADATA_SOFTWARE" -
target_list="exadatabase2.subdom.example.com"
-target_type="host"
-input_file="data:/user_name/fppinstallfest/updateExadata.prop"
```

Where the input file will contain the following:

```
WORKING_DIRECTORY=/tmp
```

Rollback Exadata Patches

Should a rollback be required after successfully deploying a patch on your Exadata targets, it is possible to rollback and retry patches using EM CLI commands.

Before you consider performing a rollback verify if the update procedure was successful or failed. You can verify via the UI by following these instructions: Track, Fix and Resume Operations.

Evaluate if a Rollback is Possible

To verify if an update can be rolled back use the following EM CLI commands depending on Exadata type:

Storage Server

```
emcli db_software_maintenance
-performOperation
-name="Rollback Storage Cell node"
-purpose="ROLLBACK_EXADATA_SOFTWARE"
-target_list="Cell1,Cell2"
-target_type="oracle_exadata"
-input_file="data://home/user_name/fppinstallfest/updateExadata.prop"
```

Switch

```
emcli db_software_maintenance
-performOperation
-name="Rollback IB Switch node"
-purpose="ROLLBACK_EXADATA_SOFTWARE"
-target_list="IB1,IB2"
-target_type="oracle_si_netswitch"
-input_file="data://home/user_name/fppinstallfest/updateExadata.prop"
```

Database

```
emcli db_software_maintenance
-performOperation -name="Rollback DB node"
-purpose="ROLLBACK_EXADATA_SOFTWARE"
-target_list="exadatadatabase3.subnet.example.com"
```

```
-target_type="host"  
-input_file="data://home/user_name/fppinstallfest/updateExadata.prop"
```

Where the input file will contain:

```
WORKING_DIRECTORY=/  
fppOptions={"eval": "true"}
```

Rollback an Update

To rollback an update operation use the following EM CLI commands depending on Exadata type:

Storage Server

```
emcli db_software_maintenance  
-performOperation  
-name="Rollback Storage Cell node"  
-purpose="ROLLBACK_EXADATA_SOFTWARE"  
-target_list="Cell1,Cell2"  
-target_type="oracle_exadata"  
-input_file="data://home/user_name/fppinstallfest/updateExadata.prop"
```

Switch

```
emcli db_software_maintenance  
-performOperation  
-name="Rollback IB Switch node"  
-purpose="ROLLBACK_EXADATA_SOFTWARE"  
-target_list="IB1,IB2"  
-target_type="oracle_si_netswitch"  
-input_file="data://home/user_name/fppinstallfest/updateExadata.prop"
```

Database

```
emcli db_software_maintenance  
-performOperation -name="Rollback DB node"  
-purpose="ROLLBACK_EXADATA_SOFTWARE"  
-target_list="exadatadatabase3.subnet.example.com"  
-target_type="host"  
-input_file="data://home/user_name/fppinstallfest/updateExadata.prop"
```

Where the input file will contain:

```
WORKING_DIRECTORY=/  
fppOptions={"eval": "true"}
```

To retry an operation follow the steps outlined in [Exadata Fleet Maintenance Patching Operations](#)

Zero Data Loss Recovery Appliance Management Pack Features

Topics:

- [Database Fleet Backup Configuration for Recovery Appliance](#)
- [Database Fleet Backup Scheduling for Recovery Appliance](#)
- [Touchless Automated End-to-End Database Fleet Backup Management](#)
- [Recovery Appliance Protected Database Fleet Archival Backup Management](#)
- [Database Fleet Advanced Backup Configuration for Recovery Appliance Using EMCLI](#)

Database Fleet Backup Credential Prerequisites

The database fleet backup configuration and scheduling operations require both database user and host user credentials in order to perform configuration and backup tasks at the database and operating system levels. Either Enterprise Manager named or preferred credentials can be used, and they must be configured prior to performing these operations. To manage named and preferred credentials, see *Configuring and Using Target Credentials in Cloud Control Security Guide*.

The specific user requirements for backup configuration and scheduling operations are as follows:

- The database user must have SYSDBA or SYSBACKUP role.
- The host (operating system) user must have write and execute permission for the database Oracle home.

To use named credentials, create the following:

- **Database credentials:** A **database** named credential with global scope (that is, applicable to all database targets) that specifies a database user meeting the above requirements for all databases in the operation.
- **Host credentials:** A **database host** named credential with global scope (that is, applicable to all database targets) that specifies an operating system user meeting the above requirements for all databases in the operation.

Alternatively, preferred credentials can be used. Since preferred credentials can be configured globally or for each individual database, they are recommended if a single credential cannot be used for all databases. To use preferred credentials, configure the following:

- **Database credentials:** Default and/or target (per-database) preferred credentials of credential set type *SYSDBA Database Credentials*
- **Host credentials:** Default and/or target (per-database) preferred credentials of credential set type *Database Host Credentials*

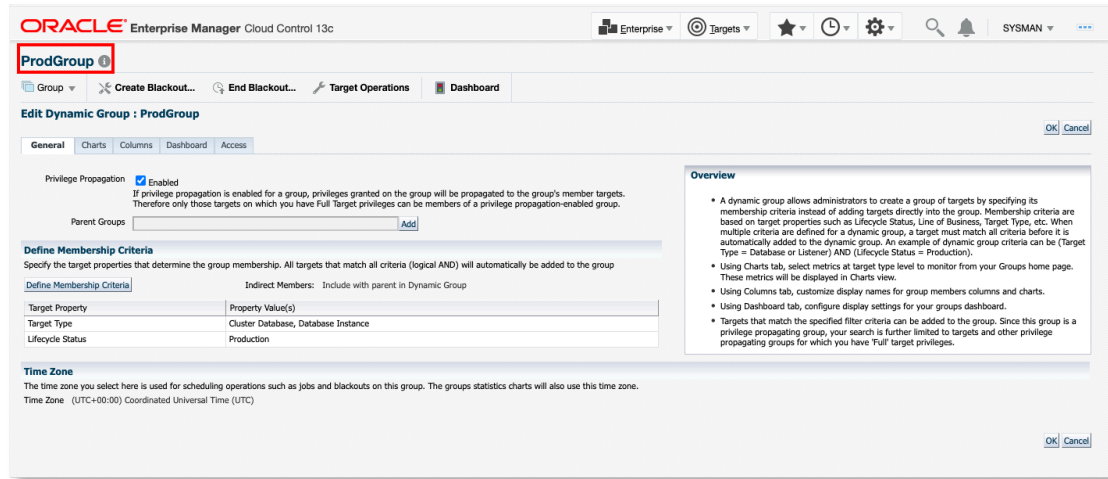
Use Groups for Configuring and Scheduling Fleet Backup to Recovery Appliance

The **Configure Fleet Backup to Recovery Appliance** wizard can be run against an Oracle Enterprise Manager *group* target. All databases within the group are configured to backup to the selected Recovery Appliance. This is particularly useful if the Oracle Enterprise Manager groups are already being used to group and manage the database targets. You can use the same group to schedule backups for the databases. Additionally, if the group is a dynamic group, the fleet backup configuration deployment procedure can be scheduled to run on a recurring interval, and new databases that subsequently join the group will be automatically

processed and configured for backup during the next scheduled procedure implementation. Similarly, after the databases are configured, they will be automatically backed during the next scheduled implementation of the fleet backup deployment procedure. See [Touchless Automated End-to-End Database Fleet Backup Management](#).

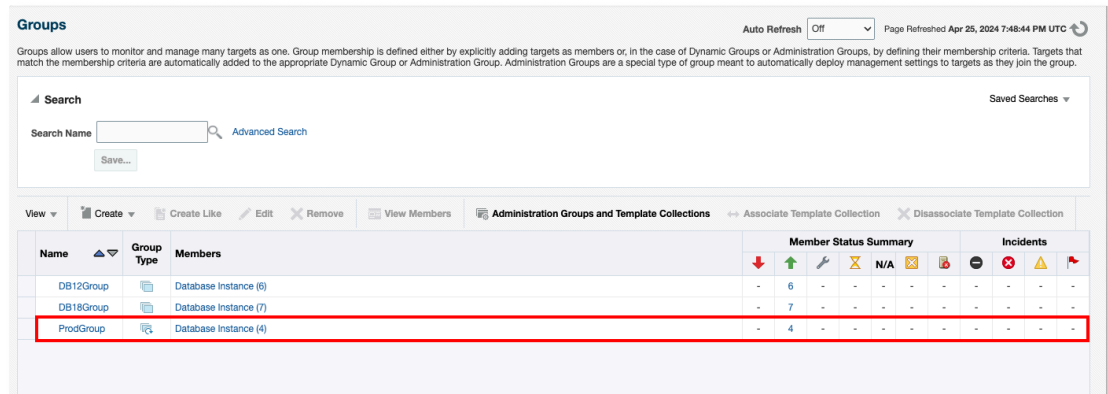
Group Membership Criteria:

The following image shows the group *ProdGroup*, a dynamic group which contains all cluster databases or single-instance databases that have their Lifecycle Status set to *Production*.



Group Members:

The following image shows that *ProdGroup* currently has 4 database members:



Database Fleet Backup Configuration for Recovery Appliance

You can use **Configure Fleet Backup to Recovery Appliance** UI to configure a fleet of databases to send backups and redo to a Recovery Appliance. The databases can be either non-Data Guard or Data Guard databases. A second Recovery Appliance can also be configured as an alternate backup destination, referred to as backup and recovery continuity.

Topics:

- [Prerequisites](#)

- [Launch the Configure Fleet Backup to Recovery Appliance Wizard](#)
- [Configure Non-Data Guard Databases to Send Backups and Redo to a Recovery Appliance](#)
- [Configure Non-Data Guard Databases for Recovery Appliance Backup and Recovery Continuity](#)
- [Configure Data Guard Databases to Send Backups and Redo to a Recovery Appliance](#)
- [Set Up Automatic Database Fleet Backup Configuration for Recovery Appliance](#)

Prerequisites

- Ensure that the Recovery Appliances that will be configured as the backup destinations have the *Zero Data Loss Recovery Appliance Management Pack* license enabled.
- Optionally, create a group with all the databases that you want to configure backups for.
- Identify the **Virtual Private Catalog User** who has access (permissions) to configure the backups for the selected databases.
- **Credentials:** [Database Fleet Backup Credential Prerequisites](#)

Launch the Configure Fleet Backup to Recovery Appliance Wizard

To launch the wizard, from the **Targets** menu, click **Recovery Appliance**. In the **Recovery Appliances** page, under **Actions**, click **Configure Fleet Backup to Recovery Appliance**.

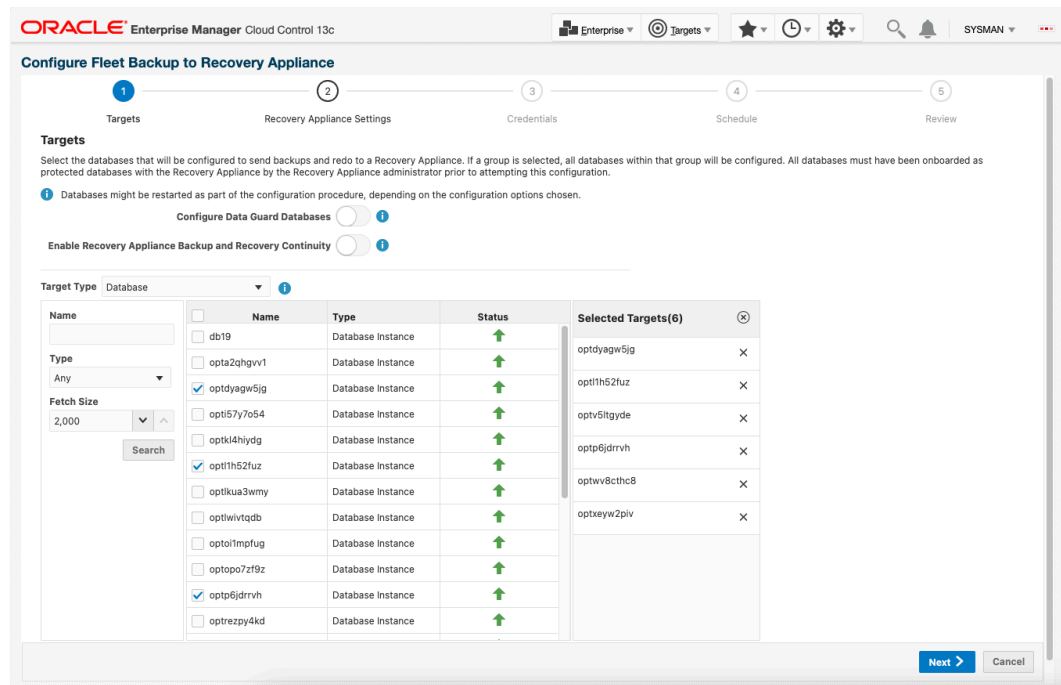
Configure Non-Data Guard Databases to Send Backups and Redo to a Recovery Appliance

In the **Configure Fleet Backup to Recovery Appliance** wizard:

1. *Targets* page:

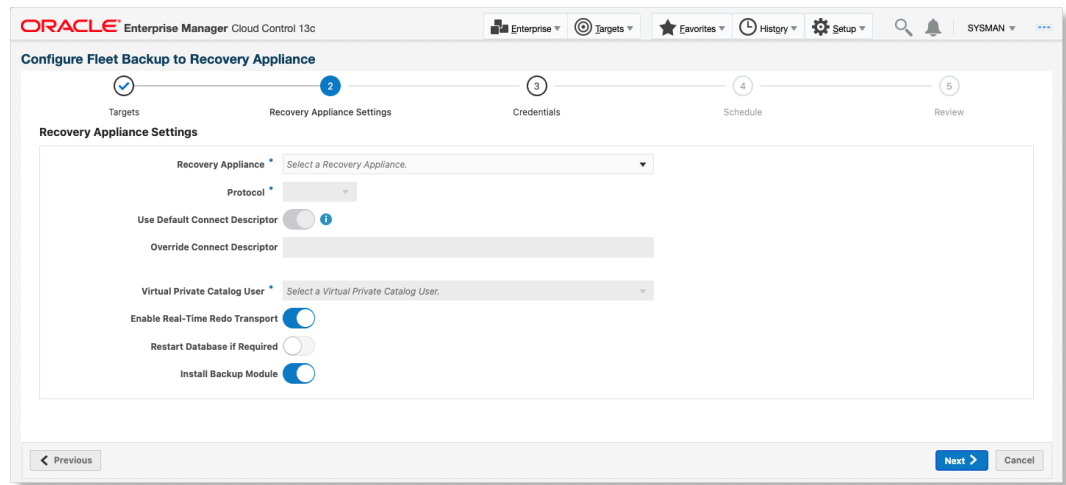
- If the databases to be configured are Data Guard primary databases, enable **Configure Data Guard Databases**. In the following example, this option is not enabled.
- Select **Enable Recovery Appliance Backup and Recovery Continuity** if required. In the following example, this option is not selected. For more information about this option, see [Configure Non-Data Guard Databases for Recovery Appliance Backup and Recovery Continuity](#). Both the Recovery Appliances must have the *Zero Data Loss Recovery Appliance Management Pack* enabled. This feature can be used only if at least two Recovery Appliances have the pack enabled.
- Select the **Target Type**. Select `Database` if the configuration is for individual databases. Select `Group` if the configuration is for a fleet of databases.
- Select the targets to configure the backups for. Select individual databases or groups containing databases.

The target selector filters the targets displayed depending on the above selection. The following example shows the selection of non-Data Guard databases and individual databases.



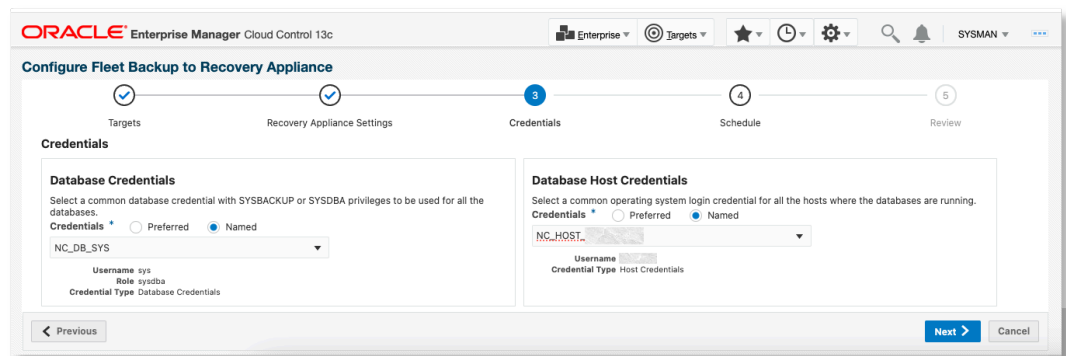
2. *Recovery Appliance Settings* page:

- Select the **Recovery Appliance** to send the backups to. All the Recovery Appliances recognized by Oracle Enterprise Manager are available in the menu. You can only configure backups to a Recovery Appliance that has the *Zero Data Loss Recovery Appliance Management Pack* enabled.
- For **Protocol**, select TCP or TCPS. If your Recovery Appliance supports only one protocol, then that is selected by default.
- Optionally, you can specify **Override Connect Descriptor** to override the one used by Oracle Enterprise Manager to connect to the Recovery Appliance.
- Select the **Virtual Private Catalog User** to use for sending the backups. All the virtual private catalog users for this Recovery Appliance that are known to Oracle Enterprise Manager are listed here. Select a virtual private catalog user who has access to the recovery catalog metadata for all the selected databases.
- Optionally, you can opt to **Enable Real-Time Redo Transport** to the Recovery Appliance.
- Optionally, select whether to install the Recovery Appliance backup module (recommended for *pre-23ai database versions*).



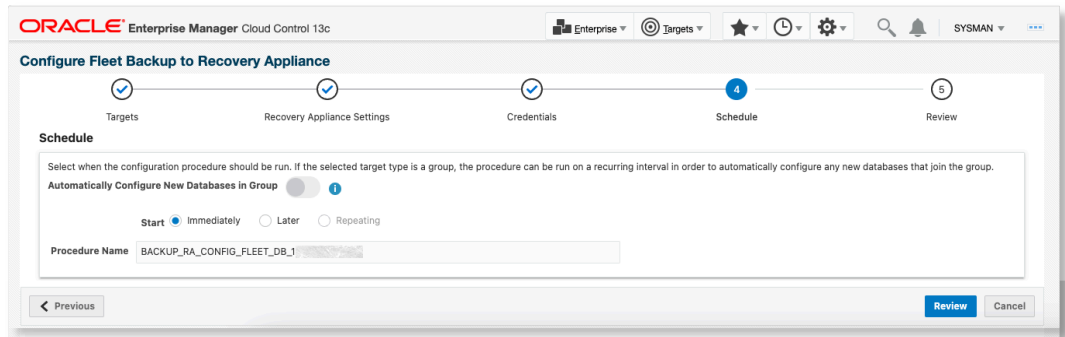
3. Credentials page:

- Select either preferred or named database credentials as per the requirements in [Prerequisites](#).
- Select either preferred or named database host credentials as per the requirements in [Prerequisites](#).

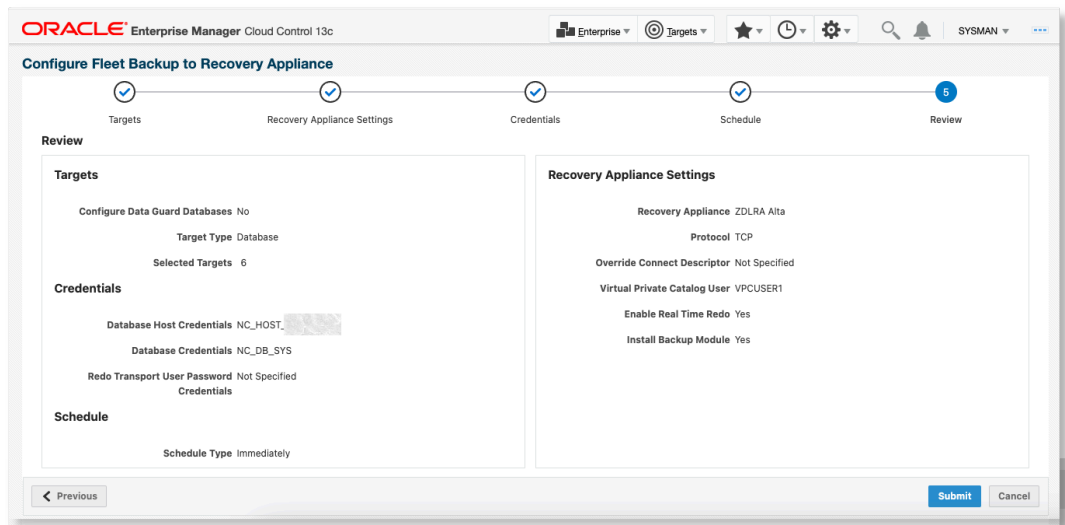


4. Schedule page:

- Enable **Automatically Configure New Databases in Group** if required. For details, see [Set Up Automatic Database Fleet Backup Configuration for Recovery Appliance](#).
- Specify when to run the deployment procedure. In the case of individual databases, select from `Immediately` or `Later`. In the case of a group of databases, you can also select `Repeating` if you want the deployment procedure to run on a recurring interval to automatically configure any new databases that join the group. In the following example, this option is not available because the use case is for individual databases.
- Specify a uniquely identifiable name for the deployment procedure that's submitted to perform the configuration.



5. **Review page:** Verify the selection made in the wizard. Click **Submit**.



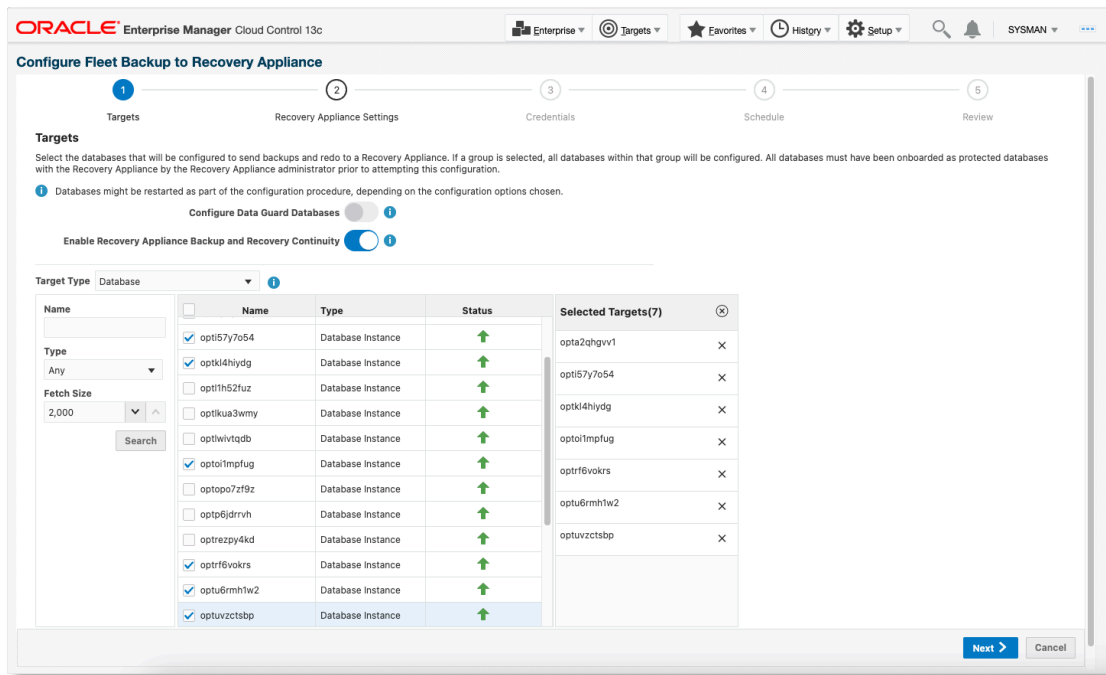
The deployment procedure to configure these databases to backup to the Recovery Appliance is submitted. A link to the deployment procedure is displayed in the confirmation dialog.

Configure Non-Data Guard Databases for Recovery Appliance Backup and Recovery Continuity

For details about the wizard options, see the section [Configure Non-Data Guard Databases to Send Backups and Redo to a Recovery Appliance](#). For the options that are unique to configuring non-Data Guard databases for Recovery Appliance backup and recovery continuity, see below:

In the *Targets* page, select **Enable Recovery Appliance Backup and Recovery Continuity**.

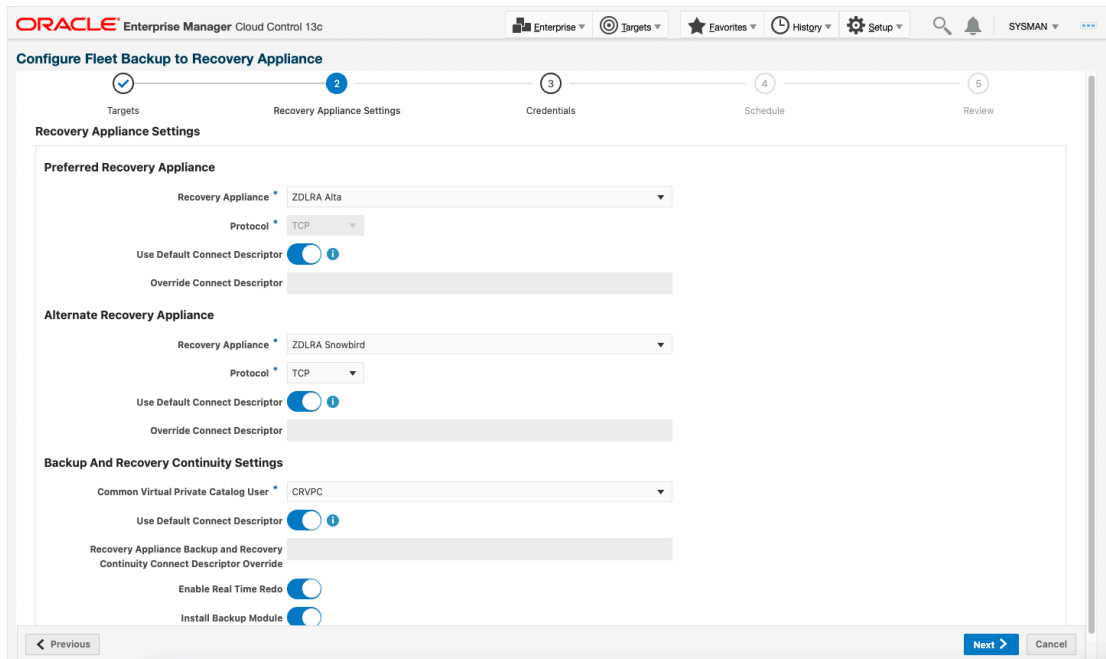
The Recovery Appliance Backup and Recovery Continuity option allows you to specify a preferred Recovery Appliance for backups and also provide an alternate Recovery Appliance to use if the preferred appliance is down or unreachable. This option is currently available only with non-Data Guard databases. Both the Recovery Appliances must have the *Zero Data Loss Recovery Appliance Management Pack* enabled. This feature can be used only if at least two Recovery Appliances have the pack enabled.



In the *Recovery Appliance Settings* page:

- Select the **Preferred Recovery Appliance**, specify the **Protocol**, and optionally provide the **Override Connect Descriptor**.
- Select the **Alternate Recovery Appliance**, specify the **Protocol**, and optionally provide the **Override Connect Descriptor**.
- Optionally, provide the **Override Connect Descriptor**.

The Common Virtual Private Recovery User must be common and have the same password for both Recovery Appliances.



Configure Data Guard Databases to Send Backups and Redo to a Recovery Appliance

For the details about the wizard options, see the section [Configure Non-Data Guard Databases to Send Backups and Redo to a Recovery Appliance](#). For the options that are unique to configuring Data Guard databases for Recovery Appliance backup, see below:

In the *Targets* page:

- Enable **Configure Data Guard Databases**. With this option, only individual databases can be selected.
- The target selector will only display the databases that are in a Data Guard configuration and are currently in a *primary* Data Guard role.
- Note that although only primary databases are shown and selected here, all the primary and standby databases in each respective Data Guard configuration are configured to backup to the selected Recovery Appliance.

Configure Fleet Backup to Recovery Appliance

1 Targets 2 Recovery Appliance Settings 3 Credentials 4 Schedule 5 Review

Targets

Select the databases that will be configured to send backups and redo to a Recovery Appliance. If a group is selected, all databases within that group will be configured. All databases must have been onboarded as protected databases with the Recovery Appliance by the Recovery Appliance administrator prior to attempting this configuration.

ⓘ Databases might be restarted as part of the configuration procedure, depending on the configuration options chosen.

Configure Data Guard Databases ⓘ

Enable Recovery Appliance Backup and Recovery Continuity ⓘ

Target Type Database ⓘ

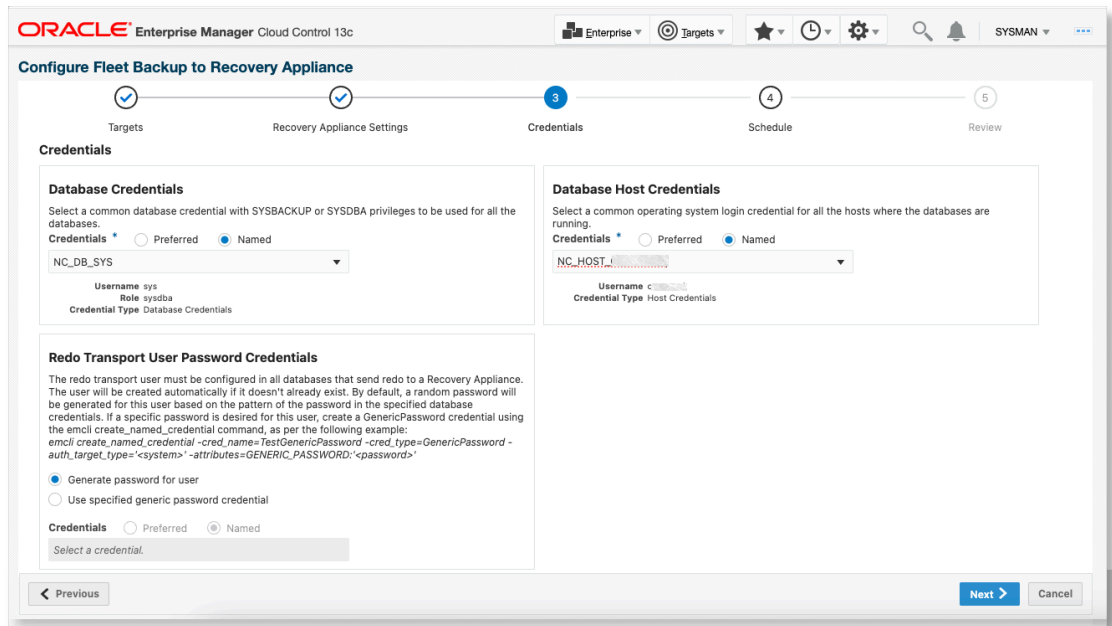
ⓘ Only primary databases in a Data Guard configuration are shown below.

Name	<input type="checkbox"/>	Name	Type	Status	Selected Targets(7)	<input type="checkbox"/>
<input checked="" type="checkbox"/>	dghfghna46_...	-targ	Database Instance	↑	dguu2lygv0_...-targ...	×
<input checked="" type="checkbox"/>	dgh4rsdtbs_...	-targ	Database Instance	↑	dgtkbwix36_...-targ...	×
<input checked="" type="checkbox"/>	dgp8z0lx9q_...	-targ	Database Instance	↑	dsgsd5bkmq_...-ta...	×
<input checked="" type="checkbox"/>	dgpy03raq5_...	-targ	Database Instance	↑	dgpy03raq5_...-tar...	×
<input checked="" type="checkbox"/>	dsgsd5bkmq_...	-tar	Database Instance	↑	dgp8z0lx9q_...-tar...	×
<input checked="" type="checkbox"/>	dgtkbwix36_...	-targ	Database Instance	↑	dgh4rsdtbs_...-targ...	×
<input checked="" type="checkbox"/>	dguu2lygv0_...	-targ	Database Instance	↑	dghfghna46_...-targ...	×
<input type="checkbox"/>	multidgb762gnzy_...		Database Instance	↑		
<input type="checkbox"/>	multidguaxw85mm_...		Database Instance	↑		

Next Cancel

In the *Credentials* page:

- The credentials specified here will also be used for all standby databases in the Data Guard configurations.
- With the addition of the Recovery Appliance to the respective Data Guard configurations, the primary database redo transport user will be changed to the Recovery Appliance virtual private catalog (VPC) user that was specified in the previous step. This user must be created in the primary database. Oracle Enterprise Manager can either automatically generate a password for this user, or an Oracle Enterprise Manager generic password credential containing this password can be specified. For the latter choice, the generic password credential must first be created using EMCLI `create_named_credential` command.

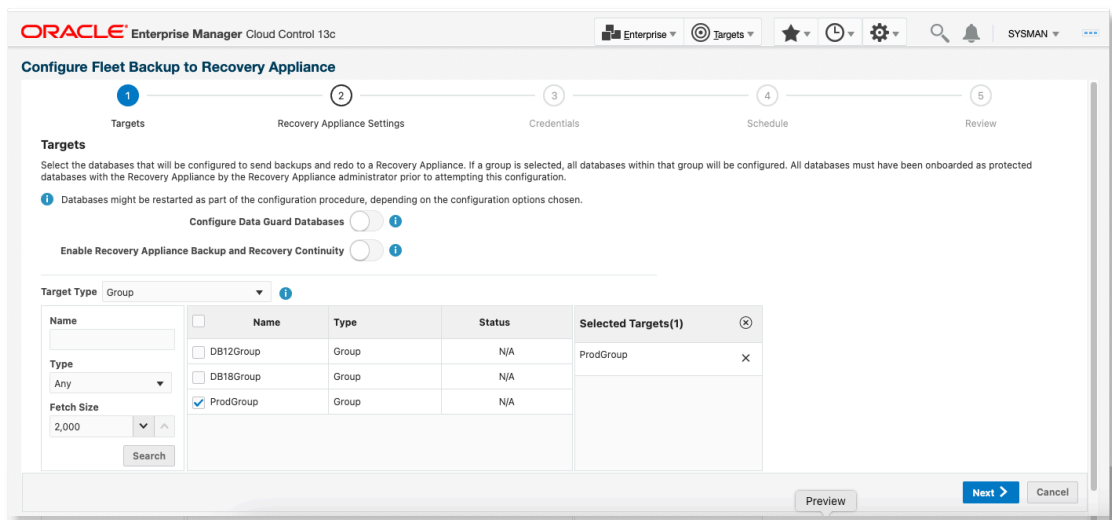


Set Up Automatic Database Fleet Backup Configuration for Recovery Appliance

For the details about the wizard options, see the section [Configure Non-Data Guard Databases to Send Backups and Redo to a Recovery Appliance](#). For the options that are unique to configuring automatic database backup configuration using a group, see below.

In *Targets* page:

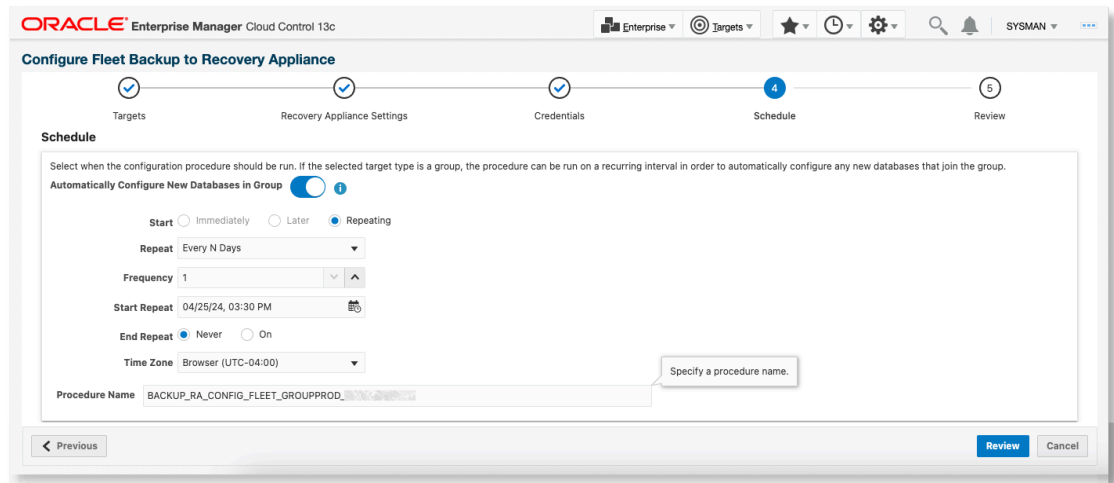
- Select the **Target Type** as *Group* for the fleet of databases.
- Select the group target to configure; for example, *ProdGroup*.



In the *Schedule* page:

- **Automatically Configure New Databases in Group** is enabled by default.

- Since the schedule is configured for a group, the default **Start** is `Repeating` so that the deployment procedure can run on a recurring interval to automatically configure any new databases that join the group. If you disable **Automatically Configure New Databases in Group**, then the options to start the procedure are **Immediately** or **Later**.
- With the fields **Repeat**, **Frequency**, **Start Repeat**, and **Time Zone**, select the frequency at which to run the deployment procedure. Optionally, using the field **End Repeat**, you can also specify an end time for this recurring execution.
- Specify a uniquely identifiable name for the deployment procedure to run for the configuration.



After the deployment procedure is submitted, it is displayed in the Procedure Activity page as *Scheduled*.

Database Fleet Backup Scheduling for Recovery Appliance

With the **Schedule Oracle-Suggested Fleet Backup to Recovery Appliance** UI, you can schedule the backup to Recovery Appliance of one or more databases that are already configured for Recovery Appliance protection.

Topics:

- [Prerequisites](#)
- [Launch the Schedule Oracle-Suggested Fleet Backup to Recovery Appliance Wizard](#)
- [Schedule Database Fleet Backup to Recovery Appliance](#)
- [Schedule Backup for a Group of Databases to Recovery Appliance](#)

Prerequisites

- The *Zero Data Loss Recovery Appliance Management Pack* must be enabled for all the Recovery Appliances involved.
- All the databases for which you want to schedule backup must already be configured to backup to a Recovery Appliance.
- **Credentials:** [Database Fleet Backup Credential Prerequisites](#)

Launch the Schedule Oracle-Suggested Fleet Backup to Recovery Appliance Wizard

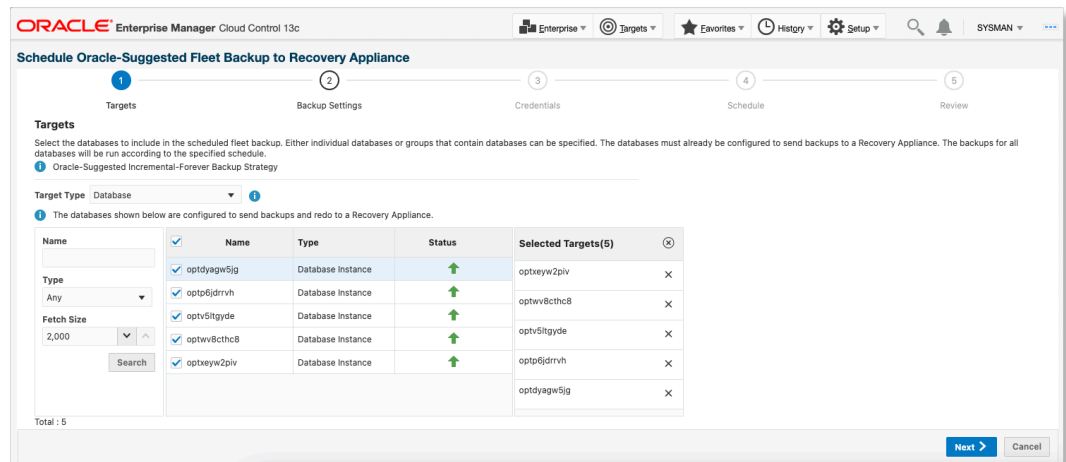
To launch the wizard, from the **Targets** menu, click **Recovery Appliance**. In the **Recovery Appliances** page, under **Actions**, click **Schedule Oracle-Suggested Fleet Backup to Recovery Appliance**.

Schedule Database Fleet Backup to Recovery Appliance

In the **Schedule Oracle-Suggested Fleet Backup to Recovery Appliance** wizard:

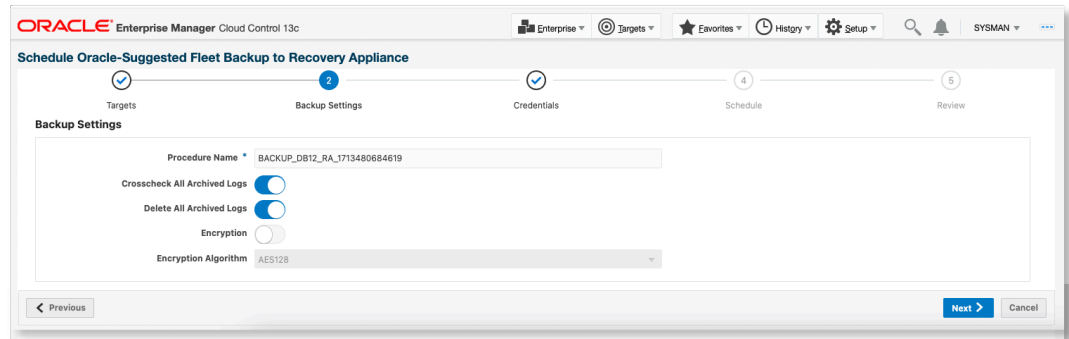
1. *Targets* page:

- Select the **Target Type**. Select **Database** if the schedule is for individual databases. Otherwise select **Group** to include all the databases within that group that are configured to backup to a Recovery Appliance in the scheduled procedure. In the current example, the schedule is for individual databases, and hence the target type is **Database**.
- Select the databases for backup. The target selector lists only those databases that are configured to backup to a Recovery Appliance. Note that the selected databases might be configured to backup to different Recovery Appliances.



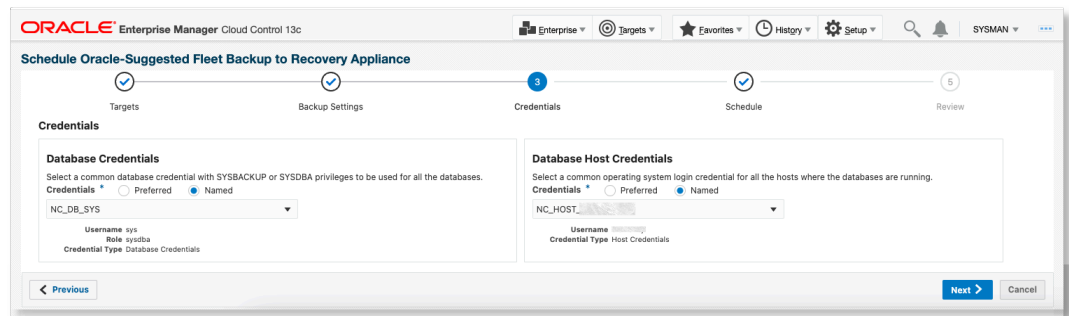
2. *Backup Settings* page:

- Specify a unique **Procedure Name** for the procedure to track it in the Procedure Activity page.
- Make archived log and encryption selections for the backup as needed.



3. **Credentials** page:

- Select either preferred or named database credentials as per the requirements in [Prerequisites](#).
- Select either preferred or named database host credentials as per the requirements in [Prerequisites](#).




4. **Schedule** page:

All the databases will be backed up to the Recovery Appliance on the schedule specified on the Schedule page.

The Oracle-recommended backup strategy for databases backing up to the Recovery Appliance is to perform an incremental backup every day (the default schedule selection).

- You can change the **Frequency** of the backup, if required.
- Using the fields **Start Repeat** and **Time Zone**, specify the time to run the backup. Optionally, using the field **End Repeat**, you can also specify an end time for this recurring execution.



5. *Review* page: Verify the selection made in the wizard. Click **Submit**.

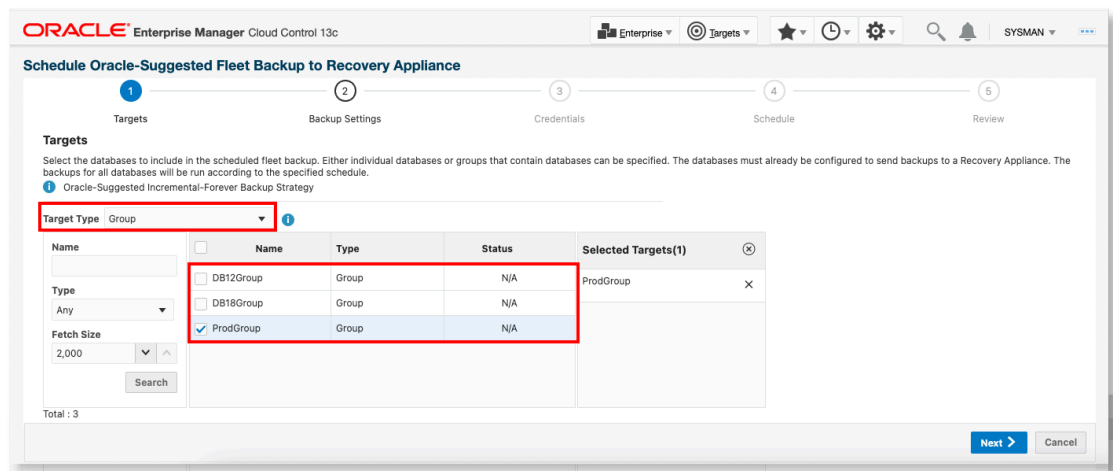
The deployment procedure to backup the databases is scheduled and the confirmation message is displayed.

Schedule Backup for a Group of Databases to Recovery Appliance

For the details about the wizard options, see the section [Schedule Database Fleet Backup to Recovery Appliance](#). For the options that are unique to scheduling backup for a group of databases to Recovery Appliance, see below:

Targets page:

Select **Group** for the **Target Type** and select one or more groups. All databases within the specified groups that are already configured to backup to a Recovery Appliance will be backed up. Databases that subsequently join the groups and that are configured will be backed up in the next scheduled procedure implementation.



Touchless Automated End-to-End Database Fleet Backup Management

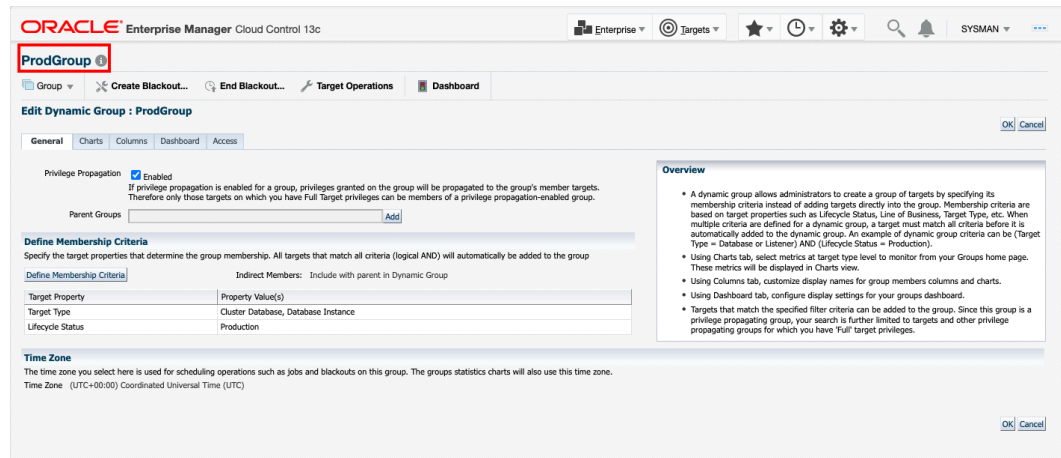
Touchless, self-managing fleet backup automation capability can be leveraged by using the protected database management features with EM dynamic groups. After the initial scheduling of the configuration and backup procedures, subsequently discovered databases that join the group are automatically configured to backup to the designated Recovery Appliance and backups scheduled.

Protected Database Fleet Configuration: Select a group as the target for the operation, then select the Auto-Configure option to schedule the procedure to run periodically. Any new databases that join the group will automatically be configured to send backups and redo to the Recovery Appliance on the next execution of the procedure. Databases that have already been configured will be skipped.

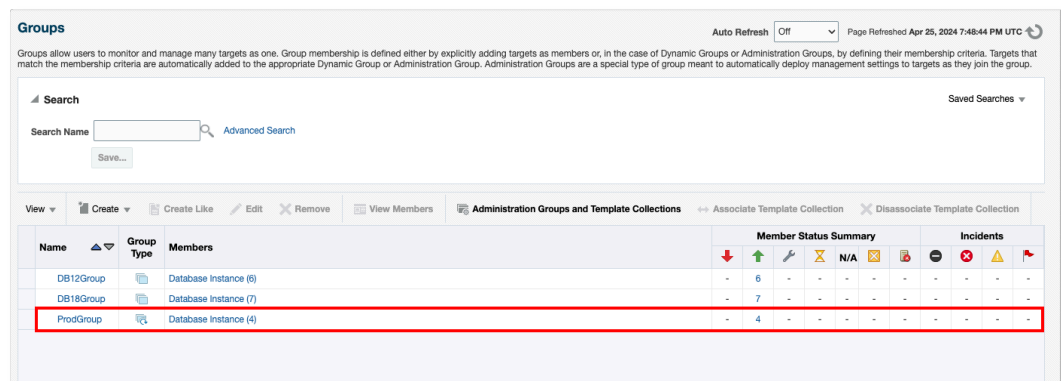
Protected Database Fleet Backup Scheduling: Select a group as the target for the operation. Scheduled backup executions against the group will automatically pick up new databases that join the group.

To set up touchless automated end-to-end backup management for a group of databases:

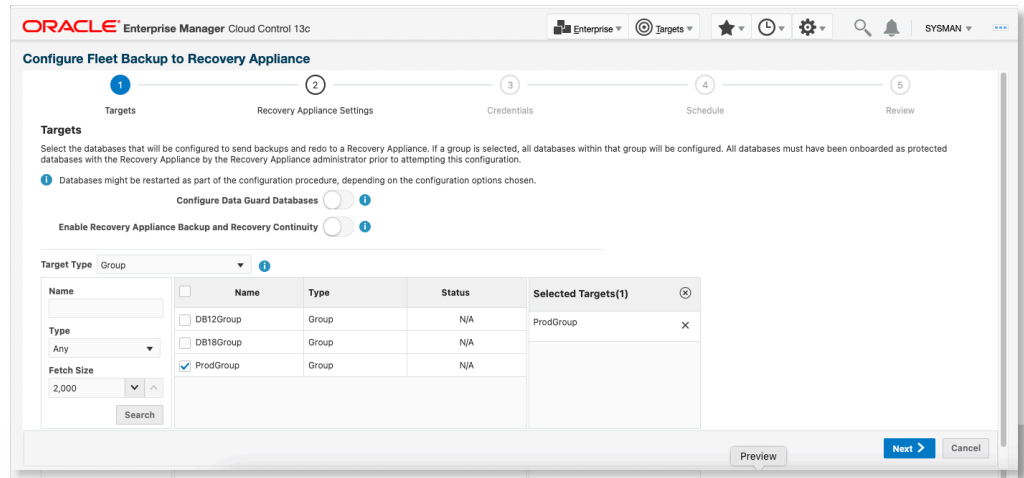
1. **Create a dynamic group in Oracle Enterprise Manager:** Specify the membership criteria for databases to join this group. In this example, the criteria is that **Lifecycle Status** must be set to `Production` value. For more information about groups, see [Use Groups for Configuring and Scheduling Fleet Backup to Recovery Appliance](#).



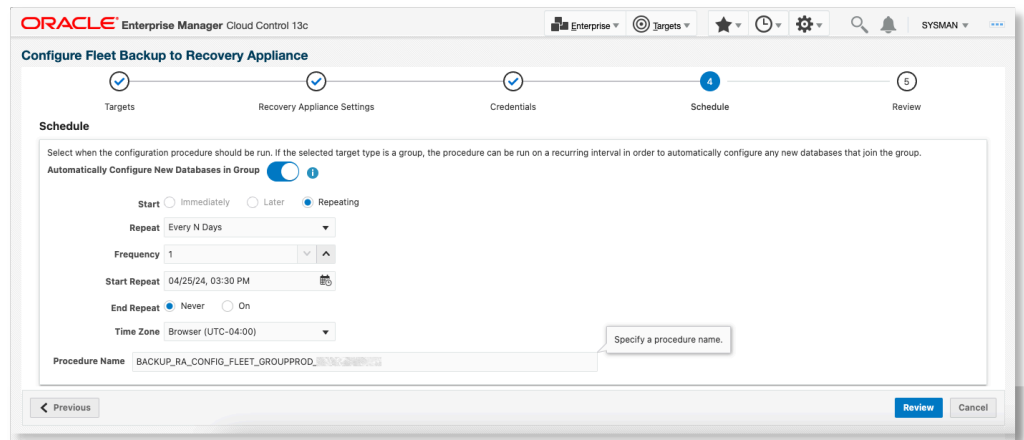
The initial status of the group `ProdGroup` shows that it has 4 members.



2. In the **Configure Fleet Backup to Recovery Appliance** wizard:
 - In the *Targets* page, select the group (for example, `ProdGroup`) as the target for the fleet configuration operation.

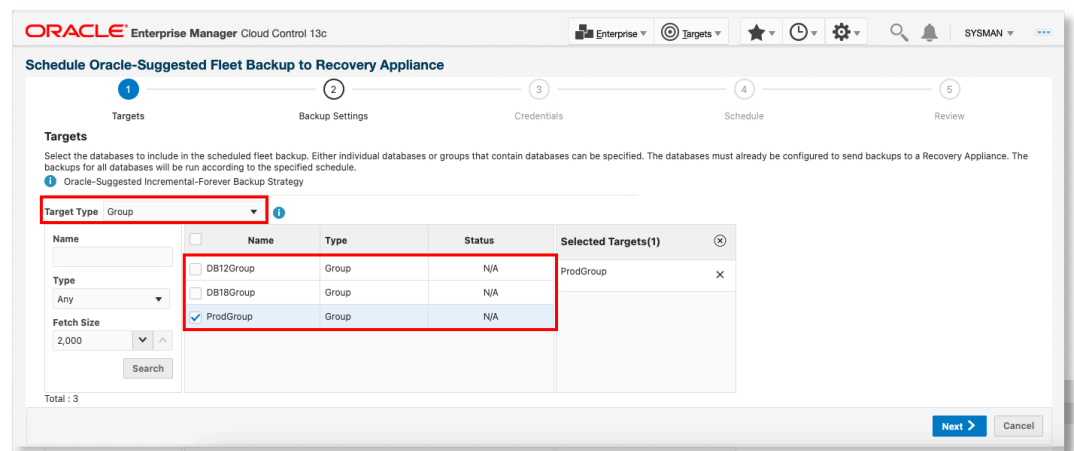


- In the *Schedule* page, enable **Automatically Configure New Databases in Group**.

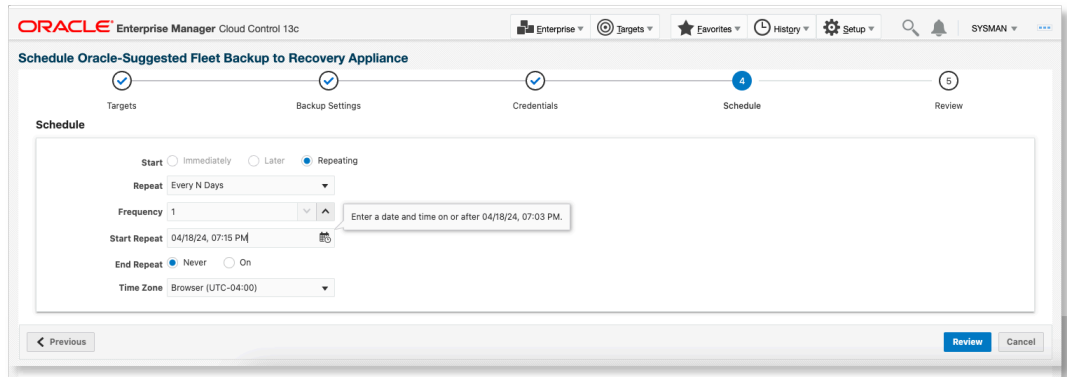


When the deployment procedure is submitted, the first execution of the deployment procedure shows that the four databases that are part of this group are configured to backup to the Recovery Appliance.

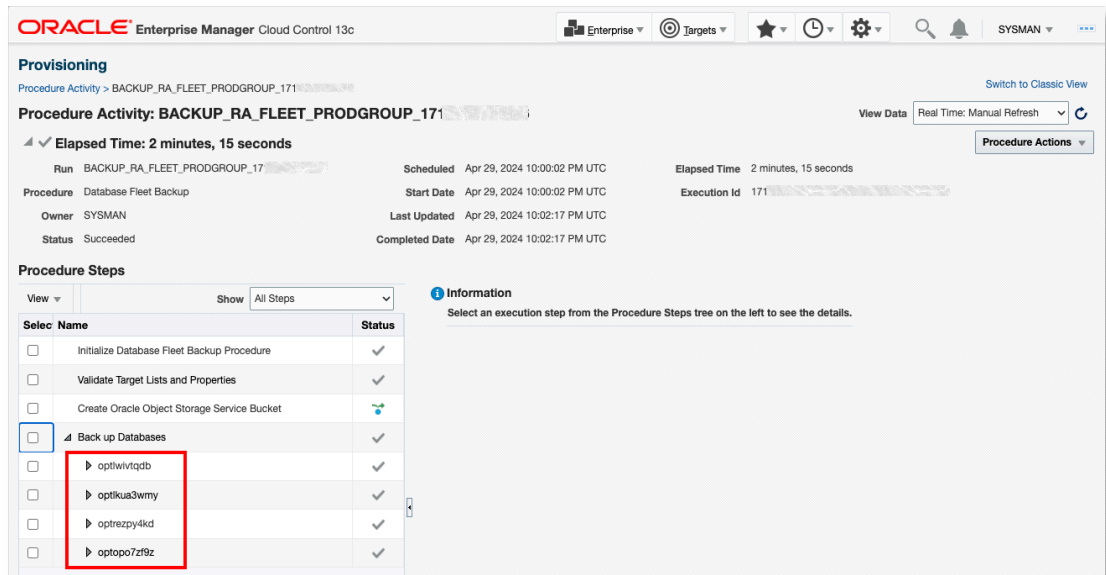
3. In the **Schedule Oracle-Suggested Fleet Backup to Recovery Appliance** wizard, select this group (for example, *ProdGroup*) as the target.



Specify to run the backup for this group of databases every night.

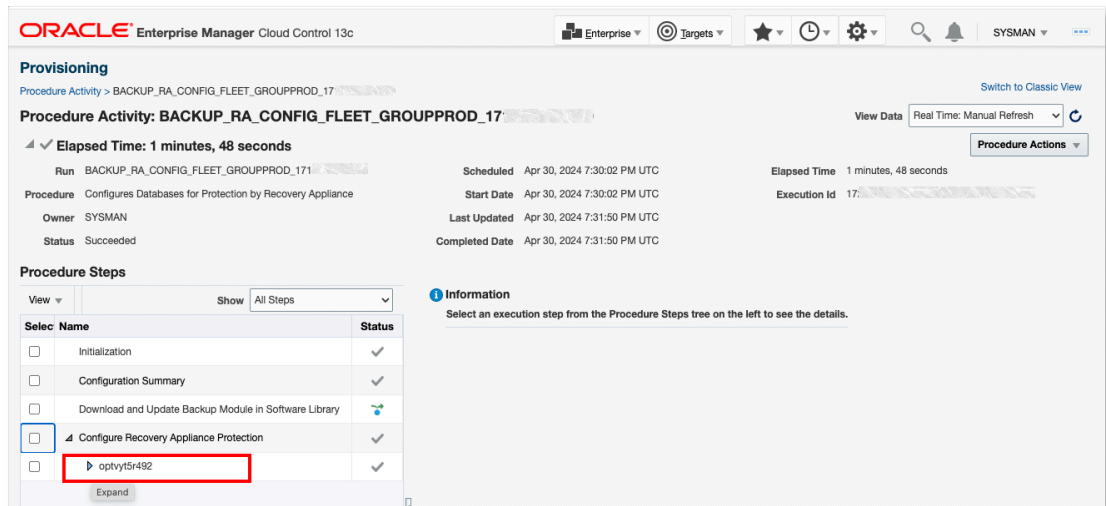


The Fleet Schedule backup procedure that runs shows that the four databases were indeed backed up.




At a later point, a fifth database is discovered with **Lifecycle status** set to **Production**. This database automatically joins the group **ProdGroup**.

The next run of the Fleet Configure Backup procedure picks up this one new database and configures it to backup to the Recovery Appliance.



Once this is done, the next run of the Fleet Schedule Backup procedure automatically backs up this database (along with the four original databases) to the Recovery Appliance.



Recovery Appliance Protected Database Fleet Archival Backup Management

Using the **Create Archival Backup** UI, you can set up a one-time archival backup or a recurring archival backup for a single protected database or all the protected databases in a specific protection policy. The archival backup is created from the database backups already existing on the Recovery Appliance.

Topics:

- [Prerequisites](#)
- [Navigate to Archival Backups Page](#)
- [Create Archival Backup for the Selected Database](#)

- [Create Recurring Archival Backup for a Single Database](#)
- [Create Archival Backups for a Database Fleet](#)

Prerequisites

- To create the archival backup for a single database, the enablement of *Zero Data Loss Recovery Appliance Management Pack* is not required. However, if you want to create archival backups for all the protected databases in a specific protection policy, then ensure to enable the *Zero Data Loss Recovery Appliance Management Pack* first.
- Before creating the archival backup, create the media manager that controls where the backup is stored (either in the Oracle Cloud, tape, or other media). See [Zero Data Loss Recovery Appliance Administrator's Guide](#).

Navigate to Archival Backups Page

To open the Archival backups page, in your Recovery Appliance page, click **Recovery Appliance**, and click **Archival Backups**. Alternatively, click **Recovery Appliance**, click **Protected Databases**, and then click **Archival Backups**.

The Archival Backups page lists all the archival backups found for the selected database. Click **Create Archival Backup** to create a new archival backup.

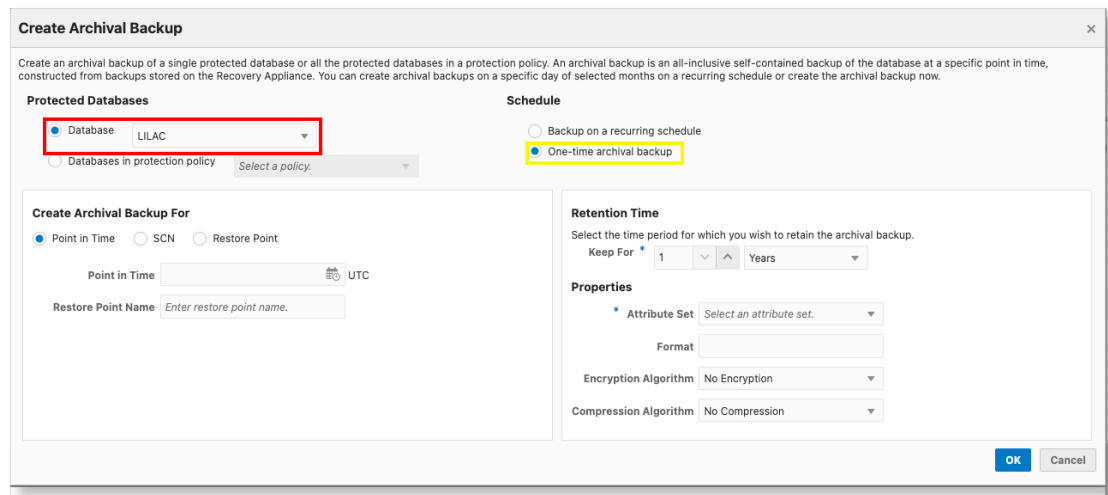
Restore Point Name	Status	SCN	Restore Tag	Recovery Point Time	Retention Time	Created
KEEP_BACKUP_20230906151356	Available	112890330	KEEP_BACKUP_20230906151356	9/5/23, 9:00:00 PM EDT	10/6/23, 5:13:00 PM EDT	9/5/23, 5:13:57 PM EDT
KEEP_BACKUP_20230906143143	Available	112890330	KEEP_BACKUP_20230906143143	9/5/23, 9:00:00 PM EDT	10/6/23, 4:31:00 PM EDT	9/5/23, 4:31:43 PM EDT

The Create Archival Backup page opens. If **Databases in Protection Policy** toggle is disabled, then the Recovery Appliance doesn't have the management pack enabled.

Create Archival Backup for the Selected Database

1. Select the database to create the archival backup for.
2. Under **Schedule**, select **One-time archival backup**.
3. Specify the **Point in Time** to create the backup for. Alternatively, specify an **SCN** or a **Restore Point** created earlier.
4. Select the **Retention Time**. This is the time period to retain this archival backup.
5. In the **Attribute Set** field, specify the media manager which controls where the backup is stored.
6. Specify the encryption and compression attributes for the archival backup, if required.

Oracle Enterprise Manager launches a deployment procedure to create the archival backup, as specified.



Create Recurring Archival Backup for a Single Database

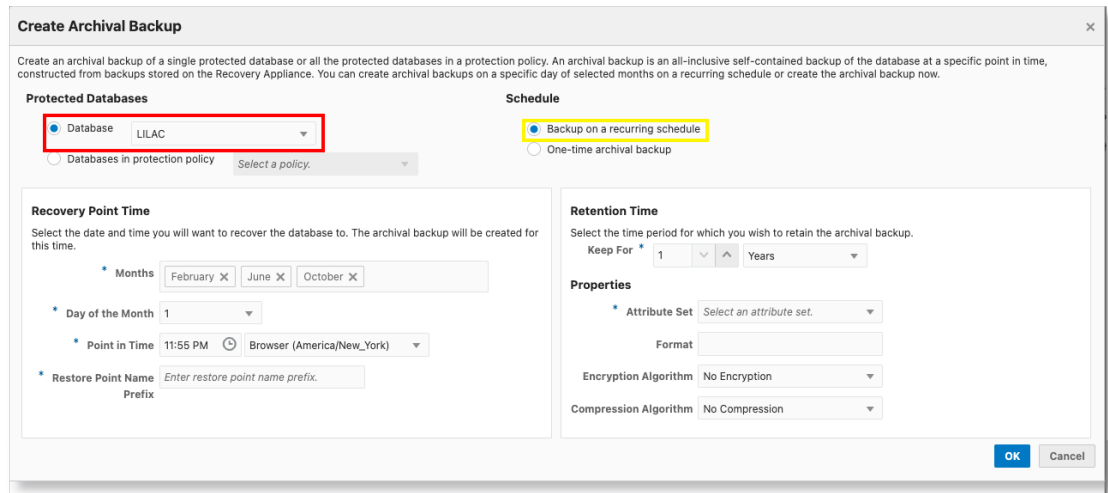
You may want to create a recurring archival backup; for example, at the end of every month, every quarter, or every year. This can be scheduled for a single database using the same **Create Archival Backup** dialog.

1. Select the database to create the archival backup for.
2. Under **Schedule**, select **Backup on a recurring schedule**.
3. Under **Recovery Point Time**, specify the months, day of the month, and time for the archival backup.
4. Specify the **Restore Point Name Prefix**. This is the prefix for the name of the archival backup restore point that is created. The resulting name is the prefix followed by the timestamp when it's created.

The remaining fields are the same as those in the one-time archival backup.

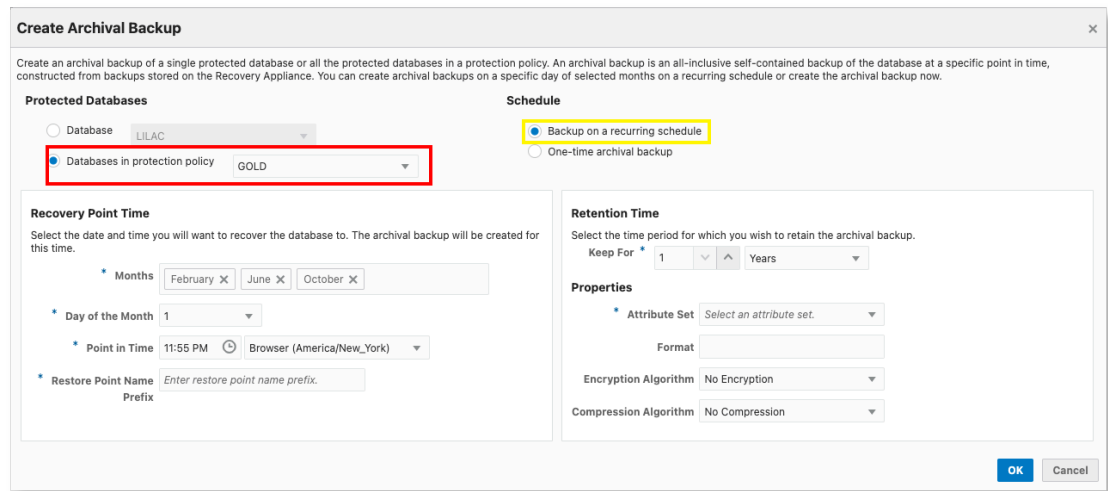
Oracle Enterprise Manager launches a deployment procedure to create an archival backup at the first date and time specified. The procedure polls the Recovery Appliance and calls the Create Archival Backup procedure when all the information required to create an optimized archival backup is available, thereby ensuring that the archival backup contains the least possible number of archived log backups.

In the following example, the first backup runs after Feb 1st. The backup created can be used to restore the database to its state on Feb 1st. The next backup runs shortly after June 1st and the following one after October 1st. It can take up to two days to create the archival backup.



Create Archival Backups for a Database Fleet

Select a protection policy for the field **Databases in protection policy** in the **Create Archival Backup** dialog box. Archival backups are created for all databases within this protection policy with the schedule you specify. For details of other fields in the dialog box, see previous sections.



Database Fleet Advanced Backup Configuration for Recovery Appliance Using EMCLI

The EMCLI `configure_db_ha` command offers highly granular per-database customization of backup and configuration options and also supports several advanced options. It allows configuration of Data Guard databases to send backup and redo to multiple Recovery Appliances, in accordance with several Maximum Availability Architecture (MAA) recommended database protection configurations.

Discussed in this topic are examples of using this EMCLI command to configure several advanced database fleet backup topologies that involve a combination of one or more of the following:


1. Data Guard primary and standby databases, either or both of which may be backing up to one or multiple corresponding Recovery Appliances.
2. Two Recovery Appliances, which may or may not be configured for two-way (Backup Anywhere) replication.
3. One or more Recovery Appliances configured as alternate backup/redo destinations in order to provide backup and recovery continuity.

Topics:

- [Basic Backup Configuration](#)
- [Basic Backup Configuration with Backup and Recovery Continuity](#)
- [MAA Gold Backup Configuration](#)
- [MAA Gold Backup Configuration with Backup and Recovery Continuity](#)
- [MAA Gold Backup Configuration with ZDLRA Backup Anywhere Replication](#)
- [MAA Gold Backup Configuration with ZDLRA Backup Anywhere Replication and Backup and Redo Continuity](#)

Basic Backup Configuration

Configuration: Non-Data Guard database sending backups and (optionally) redo to one Recovery Appliance




Configure non-Data Guard database *SalesDB* to send backups and redo to Recovery Appliance *Chicago_ZDLRA* using the virtual private catalog user *vpcuser1*. Named credentials are provided for the database and the database host.

EMCLI Command:

```
emcli configure_db_ha -configureBackupToRA -backup_config=NO_DG -
target_name="SalesDB"
    -target_type="oracle_database" -ra_target_name="Chicago_ZDLRA"
    -ra_vpc_username="vpcuser1" -db_cred="DB_USER" -
db_host_cred="DB_HOST_USER"
    -ship_redo=YES
```

Basic Backup Configuration with Backup and Recovery Continuity

Configuration: Non-Data Guard database sending backups and (optionally) redo to one Recovery Appliance, with an alternate Recovery Appliance configured for backup and redo continuity



Configure a non-Data Guard database *SalesDB* to send backups and redo to Recovery Appliance *Chicago_ZDLRA*. Backup and recovery continuity is desired and the alternate Recovery Appliance is *Bombay_ZDLRA*. The common virtual private catalog user for both Recovery Appliances is *commonvpc*. Named credentials are provided for the database and the database host.

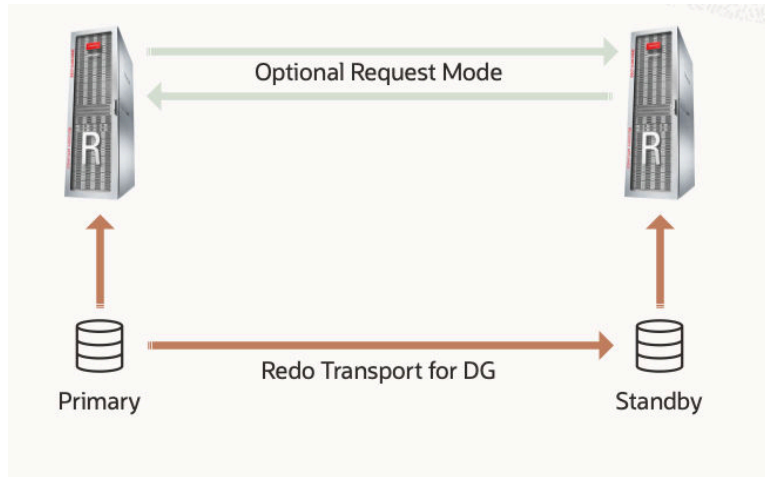
EMCLI Command:

```
emcli configure_db_ha -configureBackuptoRA -backup_config=NO_DG -
target_name="SalesDB"
    -target_type="oracle_database" -ra_target_name="Chicago_ZDLRA"
    -ra_vpc_username="commonvpc" -db_cred="DB_USER" -
db_host_cred="DB_HOST_USER"
```

```
-ship_redo=YES -br_continuity
-alternate_ra_target_name="Bombay_ZDLRA"
```

MAA Gold Backup Configuration

Configuration: Data Guard primary and standby databases sending backups and (optionally) redo to different Recovery Appliances that are not configured for Backup Anywhere Replication



FinanceDB and *FinanceStby* are the primary and standby databases respectively in a Data Guard configuration, onboarded with a pair of Recovery Appliances *Montreal_ZDLRA* and *Vancouver_ZDLRA* using a non-replicating protection policy. Configure these databases to send backups and redo to their respective Recovery Appliance using a common vpc user *commonvpc*. A global named credential *DB_SYSDBA_USER* is specified for the databases and target-specific named credential specified for each host, *DB_HOST1_CRED* and *DB_HOST2_CRED*.

EMCLI Command:

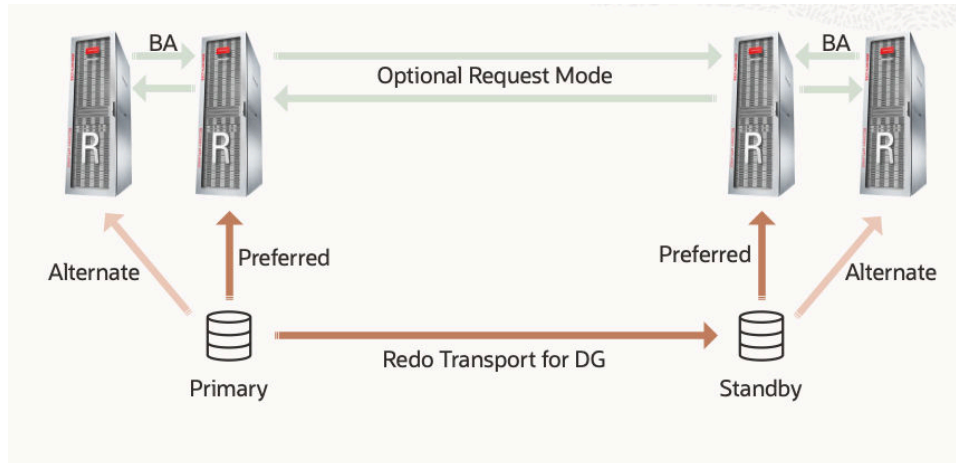
```
emcli configure_db_ha
    -configureBackupToRA -backup_config="GOLD" -input_file="target_list:
/tmp/dblist"
```

The content of */tmp/dblist* file:

```
target.0.target_name="FinanceDB"
target.0.target_type="oracle_database"
target.0.ra_target_name="Montreal_ZDLRA"
target.0.ra_vpc_username="commonvpc"
target.0.db_cred="DB_SYSDBA_USER"
target.0.db_host_cred="DB_HOST1_CRED"
target.1.target_name="FinanceStby"
target.1.target_type="oracle_database"
target.1.ra_target_name="Vancouver_ZDLRA"
target.1.ra_vpc_username="commonvpc"
target.1.db_cred="DB_SYSDBA_USER"
target.1.db_host_cred="DB_HOST2_CRED"
```

MAA Gold Backup Configuration with Backup and Recovery Continuity

Configuration: Data Guard primary and standby databases sending backups and (optionally) redo to different Recovery Appliances that are not configured for Backup Anywhere Replication, with the primary and standby databases each configured with an alternate Recovery Appliance (backup and redo continuity)



FinanceDB and *FinanceStby* are the primary and standby databases respectively in a Data Guard configuration, onboarded with a pair of Recovery Appliances *Montreal_ZDLRA* and *Vancouver_ZDLRA* using a non-replicating protection policy. Configure these databases to send backups (but not redo in this example) to their respective Recovery Appliance using a common vpc user *commonvpc*. A global named credential *DB_SYSDBA_USER* is specified for the databases and target-specific named credential specified for each host, *DB_HOST1_CRED* and *DB_HOST2_CRED*. Since backup and recovery continuity is required, alternate Recovery Appliances *Quebec_ZDLRA* and *Victoria_ZDLRA* are specified for the primary and standby databases, respectively. A replication server is set up between *Montreal_ZDLRA* and *Quebec_ZDLRA* with the protection policy being used for the *FinanceDB*. Similarly, replication is set up between *Vancouver_ZDLRA* and *Victoria_ZDLRA*.

EMCLI Command:

```
emcli configure_db_ha
    -configureBackupToRA -backup_config="GOLD" -input_file="target_list:
    /tmp/dblist" -br_continuity
```

The content of */tmp/dblist* file:


```
target.0.target_name="FinanceDB"
target.0.target_type="oracle_database"
target.0.ra_target_name="Montreal_ZDLRA"
target.0.ra_vpc_username="commonvpc"
target.0.db_cred="DB_SYSDBA_USER"
target.0.db_host_cred="DB_HOST1_CRED"
target.0.alternate_ra.0.target_name="Quebec_ZDLRA"
target.1.target_name="FinanceStby"
target.1.target_type="oracle_database"
target.1.ra_target_name="Vancouver_ZDLRA"
target.1.ra_vpc_username="commonvpc"
```

```
target.1.db_cred="DB_SYSDBA_USER"
target.1.db_host_cred="DB_HOST2_CRED"
target.1.alternate_ra.0.target_name="Victoria_ZDLRA"
```

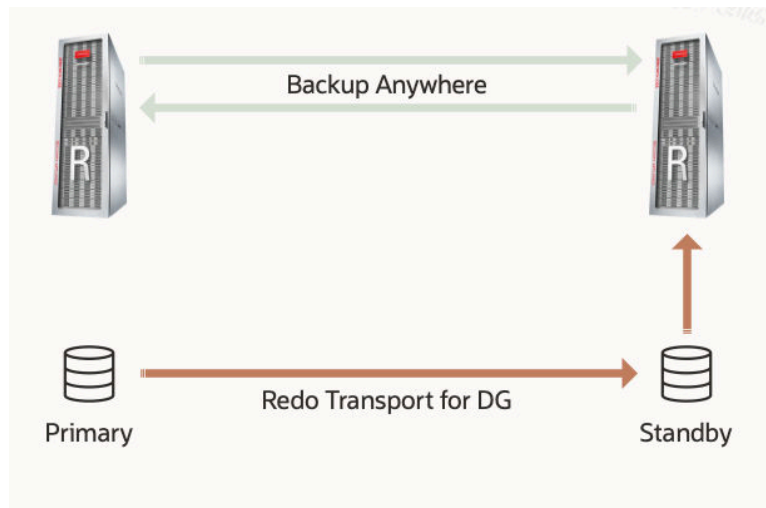
MAA Gold Backup Configuration with ZDLRA Backup Anywhere Replication

Configuration: Data Guard primary and standby databases configured to send backups and (optionally) redo to different Recovery Appliances that are configured for Backup Anywhere Replication, but with backups and redo only actually being sent by one of the databases to its respective Recovery Appliance

Backups scheduled from primary database:



Backups scheduled from standby database:



MktingDB and *MktingStby* are the primary and standby RAC databases respectively in a Data Guard configuration, onboarded with a pair of Recovery Appliances *Boston_ZDLRA* and *NY_ZDLRA* using a replicating protection policy. Backup Anywhere replication is set up for this database's protection policy between *Boston_ZDLRA* and *NY_ZDLRA*. Configure these databases to back up to their respective Recovery Appliance using a common vpc user *vpcuser1*. A global named credential *DB_SYSDBA_USER* is specified for the database(s) and target-specific named credential specified for each host, *DB_BOSHOST_CRED* and

DB_NYHOST_CRED. The *-ship_redo_from_standby* argument is not used, so redo shipping is set up from the primary database *MktingDB* to its Recovery Appliance *Boston_ZDLRA*.

EMCLI Command:

```
emcli configure_db_ha
    -configureBackupToRA -backup_config="GOLD_WITH_REP" -
input_file="target_list:
    /tmp/dblist"
```


The content of */tmp/dblist* file:

```
target.0.target_name="MktingDB"
target.0.target_type="rac_database"
target.0.ra_target_name="Boston_ZDLRA"
target.0.ra_vpc_username="vpcuser1"
target.0.db_cred="DB_SYSDBA_USER"
target.0.db_host_cred="DB_BOSHOST_CRED"
target.1.target_name="MktingStby"
target.1.target_type="rac_database"
target.1.ra_target_name="NY_ZDLRA"
target.1.ra_vpc_username="vpcuser1"
target.1.db_cred="DB_SYSDBA_USER"
target.1.db_host_cred="DB_NYHOST_CRED"
```


MAA Gold Backup Configuration with ZDLRA Backup Anywhere Replication and Backup and Redo Continuity

Configuration: Data Guard primary and standby databases configured to send backups and (optionally) redo to different Recovery Appliances that are configured for Backup Anywhere Replication, but with backups and redo only actually being sent by one of the databases to its respective Recovery Appliance, and with the primary and standby databases each configured with an alternate Recovery Appliance (backup and redo continuity)

Backups scheduled from primary database:



Backups scheduled from standby database:



MktingDB and *MktingStby* are the primary and standby RAC databases respectively in a Data Guard configuration, onboarded with a pair of Recovery Appliances *Boston_ZDLRA* and *NY_ZDLRA* using a replicating protection policy. Backup Anywhere replication is set up for this database's protection policy between *Boston_ZDLRA* and *NY_ZDLRA*. Configure these databases to back up to their respective Recovery Appliance using a common vpc user *vpcuser1*. A global named credential *DB_SYSDBA_USER* is specified for the databases and target-specific named credential specified for each host, *DB_BOSHOST_CRED* and *DB_NYHOST_CRED*. The *-ship_redo_from_standby* argument is not used, so redo shipping will be set up from the primary database *MktingDB* to its Recovery Appliance *Boston_ZDLRA*.

Since backup and recovery continuity is required, the input file also lists the alternate Recovery Appliances *Concord_ZDLRA* and *Hudson_ZDLRA* for the primary and standby databases, respectively. Backup Anywhere replication is set up for this database's protection policy between *Boston_ZDLRA* and *Concord_ZDLRA* and between *NY_ZDLRA* and *Hudson_ZDLRA*.

EMCLI Command:

```
emcli configure_db_ha
    -configureBackupToRA -backup_config="GOLD_WITH_REP" -
input_file="target_list:
    /tmp/dblist" -br_continuity
```

The content of */tmp/dblist* file:

```
target.0.target_name="MktingDB"
target.0.target_type="rac_database"
target.0.ra_target_name="Boston_ZDLRA"
target.0.ra_vpc_username="vpcuser1"
target.0.db_cred="DB_SYSDBA_USER"
target.0.db_host_cred="DB_BOSHOST_CRED"
target.0.alternate_ra.0.target_name="Concord_ZDLRA"
target.1.target_name="MktingStby"
target.1.target_type="rac_database"
target.1.ra_target_name="NY_ZDLRA"
target.1.ra_vpc_username="vpcuser1"
target.1.db_cred="DB_SYSDBA_USER"
target.1.db_host_cred="DB_NYHOST_CRED"
target.1.alternate_ra.0.target_name="Hudson_ZDLRA"
```

4

Troubleshooting

Topics:

- [Issues with Database Impact Advisor](#)

Issues with Database Impact Advisor

Topics:

- [No data in the Database Impact Advisor page](#)
- [Generate AHF Balance Report button disabled](#)
- [Generate AHF Balance Report fails](#)
- [AHF installation details invalidated](#)

No data in the Database Impact Advisor page

Issue: No data in the Database Impact Advisor page

Resolution: Ensure the Database CPU Impact Advisor metric is enabled on the Database Machine target.

Generate AHF Balance Report button disabled

Issue: Generate AHF Balance Report button disabled

Resolution: Check that the EM_AHF_USER role has been granted and AHF is installed.

Generate AHF Balance Report fails

Issue: Generate AHF Balance Report fails

Resolution: From **Enterprise**, navigate to **Provision and Patching, Procedure Activity**, and click the first *AHFBalanceRunReport* deployment procedure implementation to view errors in the procedure steps.

AHF installation details invalidated

Issue: AHF installation invalidation resulting in an error message appearing at the top of the Autonomous Health Framework page. AHF installation details can be invalidated in the following scenarios:

- The provided host credentials become invalid after the initial installation; for example, password reset, user account deleted, and so on.
- The AHF installation is manually deleted.
- The ownership or permissions on the AHF installation are changed.

Resolution: Fixing the issue requires reconfiguration of the AHF installation by a user with the EM_AHF_ADMIN role.